# vSphere Monitoring and Performance

Update 1

VMware vSphere 6.0
VMware ESXi 6.0
vCenter Server 6.0

**vm**ware®

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

https://docs.vmware.com/

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

# Contents

# About vSphere Monitoring and Performance

VMware provides several tools to help you monitor your virtual environment and to locate the source of potential issues and current problems.

**Performance charts**

Allow you to see performance data on a variety of system resources including CPU, Memory, Storage, and so on.

**Performance monitoring command-line utilities**

Allow you to access detailed information on system performance through the command line.

**Host health**

Allows you to quickly identify which hosts are healthy and which are experiencing problems.

**Events, alerts, and alarms**

Allow you to configure alerts and alarms and to specify the actions the system should take when they are triggered.

**System Log Files**

System logs contain additional information about activities in your vSphere environment.

## Intended Audience

The content in this section is intended for vSphere administrators who perform the following tasks:

- Monitor the health and performance of physical hardware backings for the virtual environment.

- Monitor the health and performance of virtual devices in the virtual environment.

- Troubleshoot problems in the system.

- Configure alarms.

- Configure SNMP messages.

Virtual machine administrators also might find the section on Chapter 3 Monitoring Guest Operating System Performance helpful.

# Updated Information

This *vSphere Monitoring and Performance* documentation is updated with each release of the product or when necessary.

This table provides the update history of the *vSphere Monitoring and Performance* documentation.

| Revision | Description |
| --- | --- |
| 11 AUG 2020 | At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we are replacing some of the terminology in our content. We have updated this guide to remove instances of non-inclusive language. |
| EN-001901-02 | Updated the required privileges to Export System Log Files. |
| EN-001901-01 | Adding information to Memory Panel about CNSM statistic that is part of the memory utilization statistics that esxtop utility provides. |
| EN-001901-00 | Initial release. |

# Monitoring Inventory Objects with Performance Charts

<span style="font-size:3em; color:#cccccc; float:right;">1</span>

The vSphere statistics subsystem collects data on the resource usage of inventory objects. Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter Server database. You can access statistical information through command-line monitoring utilities or by viewing performance charts in the vSphere Web Client.

## Counters and Metric Groups

vCenter Server systems and hosts use data counters to query for statistics. A data counter is a unit of information relevant to a given inventory object or device. Each counter collects data for a different statistic in a metric group. For example, the disk metric group includes separate data counters to collect data for disk read rate, disk write rate, and disk usage. Statistics for each counter are rolled up after a specified collection interval. Each data counter consists of several attributes that are used to determine the statistical value collected.

For a complete list and description of performance metrics, see the *vSphere API Reference*.

**Note**   Counters that are introduced in later versions might not contain data from hosts of earlier versions. For details, see the VMware Knowledge Base.

## Collection Levels and Collection Intervals

Collection levels determine the number of counters for which data is gathered during each collection interval. Collection intervals determine the time period during which statistics are aggregated, calculated, rolled up, and archived in the vCenter Server database. Together, the collection interval and collection level determine how much statistical data is collected and stored in your vCenter Server database.

## Data Availability

Real-time data appears in the performance charts only for hosts and virtual machines that are powered on. Historical data appears for all supported inventory objects, but might be unavailable during certain circumstances.

This chapter includes the following topics:

- Performance Chart Types

- Data Counters

- Metric Groups in vSphere

- Data Collection Intervals

- Data Collection Levels

- View Performance Charts

- Performance Charts Options Available Under the View Menu

- Overview Performance Charts

- Working with Advanced and Custom Charts

- Troubleshoot and Enhance Performance

## Performance Chart Types

Performance metrics are displayed in different types of charts, depending on the metric type and object.

Table 1-1. Performance Chart Types

| Chart Type | Description |
| --- | --- |
| Line chart | Displays metrics for a single inventory object. The data for each performance counter is plotted on a separate line in the chart. For example, a network chart for a host can contain two lines: one showing the number of packets received, and one showing the number of packets transmitted. |
| Bar chart | Displays storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, snapshots, swap files, and other files. |
| Pie chart | Displays storage metrics for a single object, based on the file types or virtual machines. For example, a pie chart for a datastore can display the amount of storage space occupied by the virtual machines taking up the largest space. |
| Stacked chart | Displays metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term **Other**. For example, a host's stacked CPU usage chart displays CPU usage metrics for the ten virtual machines on the host that are consuming the most CPU. The **Other** amount contains the total CPU usage of the remaining virtual machines. The metrics for the host itself are displayed in separate line charts. Stacked charts are useful in comparing resource allocation and usage across multiple hosts or virtual machines. By default, the ten child objects with the highest data counter values are displayed. |

## Data Counters

Each data counter includes several attributes that are used to determine the statistical value collected. See the *vSphere API Reference* for a complete list and description of supported counters.

## Table 1-2. Data Counter Attributes

| Attribute | Description |
|---|---|
| Unit of Measurement | Standard in which the statistic quantity is measured.<br>■ Kilobytes (KB) – 1024 bytes<br>■ Kilobytes per second (KBps) – 1024 bytes per second<br>■ Kilobits (kb) – 1000 bits<br>■ Kilobits per second (kbps) – 1000 bits per second<br>■ Megabytes (MB)<br>■ Megabytes per second (MBps)<br>■ Megabits (Mb), megabits per second (Mbps)<br>■ Megahertz (MHz)<br>■ Microseconds (μs)<br>■ Milliseconds (ms)<br>■ Number (#)<br>■ Percent (%)<br>■ Seconds (s) |
| Description | Text description of the data counter. |
| Statistics Type | Measurement used during the statistics interval. Related to the unit of measurement.<br>■ Rate – Value over the current statistics interval<br>■ Delta – Change from previous statistics interval.<br>■ Absolute – Absolute value (independent of the statistics interval). |
| Rollup Type | Calculation method used during the statistics interval to roll up data. Determines the type of statistical values that are returned for the counter.<br>■ Average – Data collected during the interval is aggregated and averaged.<br>■ Minimum – The minimum value is rolled up.<br>■ Maximum – The maximum value is rolled up.<br><br>The Minimum and Maximum values are collected and displayed only in statistics level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the average minimum or average maximum.<br><br>For example, the following information for the CPU usage chart shows that the average is collected at statistics level 1 and the minimum and maximum values are collected at statistics level 4.<br>■ Counter: usage<br>■ Unit: Percentage (%)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4)<br>■ Summation – Data collected is summed. The measurement displayed in the chart represents the sum of data collected during the interval.<br>■ Latest – Data collected during the interval is a set value. The value displayed in the performance charts represents the current value. |
| Collection level | Number of data counters used to collect statistics. Collection levels range from 1 to 4, with 4 having the most counters.<br><br>**Note** Be careful when you set a higher collection level, as the process requires significant increase of resource usage. For more information, see Data Collection Levels . |

# Metric Groups in vSphere

The performance data collection subsystem for vSphere collects performance data on a variety of inventory items and their devices. Data counters define individual performance metrics. Performance metrics are organized into logical groups based on the object or object device. Statistics for one or more metrics can be displayed in a chart.

Table 1-3. Metric Groups

| Metric group | Description |
| --- | --- |
| Cluster Services | Performance statistics for clusters configured by using vSphere Distributed Resource Scheduler, vSphere High Availability, or both. |
| CPU | CPU utilization per host, virtual machine, resource pool, or compute resource. |
| Datastore | Statistics for datastore utilization. |
| Disk | Disk utilization per host, virtual machine, or datastore. Disk metrics include I/O performance, such as latency and read/write speeds, and utilization metrics for storage as a finite resource. |
| Memory | Memory utilization per host, virtual machine, resource pool, or compute resource. The value obtained is one of the following:<br><br>■ For virtual machines, memory refers to the guest physical memory. Guest physical memory is the amount of physical memory presented as a virtual-hardware component to the virtual machine, at creation time, and made available when the virtual machine is running.<br>■ For hosts, memory refers to the machine memory. Machine memory is the RAM that is installed on the hardware that comprises the host. |
| Network | Network utilization for both physical and virtual network interface controllers (NICs) and other network devices, such as the virtual switches that support connectivity among all components, such as hosts, virtual machines, VMkernel, and so on. |
| Power | Energy usage statistics per host. |
| Storage Adapter | Data traffic statistics per host bus adapter (HBA). |
| Storage Path | Data traffic statistics per path. |
| System | Overall system availability, such as system heartbeat and uptime. These counters are available directly from hosts and from vCenter Server. |
| Virtual Disk | Disk utilization and disk performance metrics for virtual machines. |
| Virtual Flash | Virtual flash counters. |
| Virtual Machine Operations | Virtual machine power and provisioning operations in a cluster or data center. |
| vSphere Replication | Statistics for virtual machine replication performed by VMware vCenter Site Recovery Manager. |

# Data Collection Intervals

Collection intervals determine the duration for which statistics are aggregated, calculated, rolled up, and archived. Together, the collection interval and collection level determine how much statistical data is gathered and stored in your vCenter Server database.

Table 1-4. Collection Intervals

| Collection Interval/ Archive Length | Collection Frequency | Default Behavior |
|---|---|---|
| 1 Day | 5 Minutes | Real-time statistics are rolled up to create one data point every 5 minutes. The result is 12 data points every hour and 288 data points every day. After 30 minutes, the six data points collected are aggregated and rolled up as a data point for the 1 Week time range.<br><br>You can change the interval duration and archive length of the 1 Day collection interval by configuring the statistics settings. |
| 1 Week | 30 Minutes | 1 Day statistics are rolled up to create one data point every 30 minutes. The result is 48 data points every day and 336 data points every week. Every 2 hours, the 12 data points collected are aggregated and rolled up as a data point for the 1 Month time range.<br><br>You cannot change the default settings of the 1 Week collection interval. |
| 1 Month | 2 Hours | 1 Week statistics are rolled up to create one data point every 2 hours. The result is 12 data points every day and 360 data points every month (assuming a 30-day month). After 24 hours, the 12 data points collected are aggregated and rolled up as a data point for the **1 Year** time range.<br><br>You cannot change the default settings of the 1 Month collection interval. |
| 1 Year | 1 Day | 1 Month statistics are rolled up to create one data point every day. The result is 365 data points each year.<br><br>You can change the archive length of the 1 Year collection interval by configuring the statistics settings. |

**Note** If you change the duration of data collection intervals you might need to allocate more storage resources.

## Data Collection Levels

Each collection interval has a default collection level that determines the amount of data gathered and which counters are available for display in charts. Collection levels are also referred to as statistics levels.

Table 1-5. Statistics Levels

| Level | Metrics | Best Practice |
|---|---|---|
| Level 1 | ■ Cluster Services (VMware Distributed Resource Scheduler) – all metrics<br>■ CPU – cpuentitlement, totalmhz, usage (average), usagemhz<br>■ Disk – capacity, maxTotalLatency, provisioned, unshared, usage (average), used<br>■ Memory – consumed, mementitlement, overhead, swapinRate, swapoutRate, swapused, totalmb, usage (average), vmmemctl (balloon)<br>■ Network – usage (average), IPv6<br>■ System – heartbeat, uptime<br>■ Virtual Machine Operations – numChangeDS, numChangeHost, numChangeHostDS | Use for long-term performance monitoring when device statistics are not required.<br><br>Level 1 is the default Collection Level for all Collection Intervals. |
| Level 2 | ■ Level 1 metrics<br>■ CPU – idle, reservedCapacity<br>■ Disk – All metrics, excluding numberRead and numberWrite.<br>■ Memory – All metrics, excluding memUsed and maximum and minimum rollup values.<br>■ Virtual Machine Operations – All metrics | Use for long-term performance monitoring when device statistics are not required but you want to monitor more than the basic statistics. |
| Level 3 | ■ Level 1 and Level 2 metrics<br>■ Metrics for all counters, excluding minimum and maximum rollup values.<br>■ Device metrics | Use for short-term performance monitoring after encountering problems or when device statistics are required.<br><br>Because of the large quantity of troubleshooting data retrieved and recorded, use level 3 for the shortest time period (Day or Week collection interval). |
| Level 4 | All metrics supported by the vCenter Server, including minimum and maximum rollup values. | Use for short-term performance monitoring after encountering problems or when device statistics are required.<br><br>Because of the large quantity of troubleshooting data retrieved and recorded, use level 4 for the shortest amount of time. |

**Note** When you increase the collection level the storage and system requirements might change. You might need to allocate more system resources to avoid decrease in the performance.

# View Performance Charts

The vCenter Server statistics settings, the type of object selected, and the features that are enabled on the selected object determine the amount of information displayed in charts. Charts are organized into views. You can select a view to see related data together on one screen. You

can also specify the time range, or data collection interval. The duration extends from the selected time range to the present time.

Overview charts display multiple data sets in one panel to easily evaluate different resource statistics, display thumbnail charts for child objects, and display charts for a parent and a child object. Advanced charts display more information than overview charts, are configurable, and can be printed or exported. You can export data in the PNG, JPEG, or CSV formats. See View Advanced Performance Charts.

**Procedure**

1   Select a valid inventory object in the vSphere Web Client.

    Overview and advanced performance charts are available for datacenter, cluster, host, resource pool, vApp, and virtual machine objects. Overview charts are also available for datastores and datastore clusters. Performance charts are not available for network objects.

2   Click the **Monitor** tab, and click **Performance**.

3   Select a view.

    Available views depend on the type of object. For views that might contain a large number of charts in a large environment, the vSphere Web Client displays the charts distributed on multiple pages. You can use the arrow buttons to navigate between pages.

4   Select a predefined or custom time range.

## Performance Charts Options Available Under the View Menu

The performance chart options that you can access under the **View** menu vary depending on the type of inventory object you select.

For example, the **Virtual Machines** view is available when you view host performance charts only if there are virtual machines on the selected host. Likewise, the **Fault Tolerance** view for virtual machine performance charts is available only when that feature is enabled for the selected virtual machine.

Table 1-6. Performance Chart Views by Inventory Object

| Object | View List Items |
|---|---|
| Data center | ■ **Storage** - space utilization charts for datastores in the data center, including space by file type and storage space used by each datastore in the data center.<br>■ **Clusters** - thumbnail CPU and memory charts for each cluster, and stacked charts for total CPU and memory usage in the data center. This view is the default. |
| Datastore and datastore cluster | ■ **Space** - space utilization charts for the datastore:<br>　■ space utilization by file type<br>　■ space utilization by virtual machine<br>　■ space usage<br>■ **Performance** - performance charts for the datastore or datastore cluster and for virtual machine disks on the resource.<br><br>**Note** The Performance view for datastores is only available when all hosts that are connected to the datastores are ESX/ESXi 4.1 or greater. The Performance view for datastore clusters is only available when the Storage DRS is enabled. |
| Cluster | ■ **Home** - CPU and memory charts for the cluster.<br>■ **Resource Pools & Virtual Machines** - thumbnail charts for resource pools and virtual machines, and stacked charts for total CPU and memory usage in the cluster.<br>■ **Hosts** - thumbnail charts for each host in the cluster, and stacked charts for total CPU, memory, disk usage, and network usage. |
| Host | ■ **Home** - CPU, memory, disk, and network charts for the host.<br>■ **Virtual Machines** - thumbnail charts for virtual machines, and stacked charts for total CPU usage and total memory usage on the host. |
| Resource Pool and vApps | ■ **Home** - CPU and memory charts for the resource pool.<br>■ **Resource Pools & Virtual Machines** - thumbnail charts for resource pools, and virtual machines and stacked charts for CPU and memory usage in the resource pool or vApp. |
| Virtual Machine | ■ **Storage** - space utilization charts for the virtual machine: space by file type, space by datastore, and total gigabytes.<br>■ **Fault Tolerance** - CPU and memory charts that display comparative metrics for the fault-tolerant primary and secondary virtual machines.<br>■ **Home** - CPU, memory, network, host (thumbnail charts), and disk usage charts for the virtual machine. |

# Overview Performance Charts

The overview performance charts display the most common metrics for an object in the inventory. Use these charts to monitor and troubleshoot performance problems.

The metrics provided in Overview performance charts are a subset of those collected for hosts and the vCenter Server. For a complete list of all metrics collected by hosts and the vCenter Server, see the *vSphere API Reference*.

# Clusters

The cluster charts contain information about CPU, disk, memory, and network usage for clusters. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (MHz)

The CPU (MHz) chart displays CPU usage for the cluster.

### Cluster Counters

This chart is located in the Home view of the Cluster **Performance** tab.

Table 1-7. Data Counters

| Chart Label | Description |
| --- | --- |
| Usage | Sum of the average CPU usage values, in Megahertz, of all virtual machines in the cluster.<br>■ Counter: usagemhz<br>■ Stats Type: Rate<br>■ Unit: Megahertz (MHz)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |
| Total | Total amount of CPU resources available in the cluster. The maximum value is equal to the number of cores multiplied by the frequency of the processors.<br>For example, a cluster has two hosts, each of which has four CPUs that are 3GHz each, and one virtual machine that has two virtual CPUs.<br>VM totalmhz = 2 vCPUs * 3000MHz = 6000MHz<br>Host totalmhz = 4 CPUs * 3000MHz = 12000MHz<br>Cluster totalmhz = 2 x 4 * 3000MHz = 24000MHz<br>■ Counter: totalmhz<br>■ Stats Type: Rate<br>■ Unit: Megahertz (MHz)<br>■ Rollup Type: Summation<br>■ Collection Level: 1 |

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

## Table 1-8. CPU Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If the cluster is not a DRS cluster, enable DRS. Select the cluster, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |
| 3 | If the cluster is a DRS cluster:<br>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.<br>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 4 | Migrate one or more virtual machines to a new host. |
| 5 | Upgrade the physical CPUs or cores on each host in the cluster if necessary. |
| 6 | Enable CPU-saving features, such as TCP Segmentation Offload. |
| 7 | Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs. |

## CPU Usage

The cluster CPU Usage charts monitors the CPU utilization of the hosts, resource pools, and virtual machines in the cluster. This chart displays the 10 child objects in the cluster with the most CPU usage.

This chart is located in the Resource Pools and Virtual Machines view of the Cluster **Performance** tab.

## Table 1-9. Data Counters

| Chart Label | Description |
|---|---|
| <host>, <resource pool>, or <virtual machine> | Amount of CPU actively used by the host, resource pool, or virtual machine in the cluster.<br>■ Counter: usagemhz<br>■ Stats Type: Rate<br>■ Unit: MegaHertz (MHz)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

## Table 1-10. CPU Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If the cluster is not a DRS cluster, enable DRS. Select the cluster, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |
| 3 | If the cluster is a DRS cluster:<br>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.<br>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 4 | Migrate one or more virtual machines to a new host. |
| 5 | Upgrade the physical CPUs or cores on each host in the cluster if necessary. |
| 6 | Enable CPU-saving features, such as TCP Segmentation Offload. |
| 7 | Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs. |

## Disk (KBps)

The Disk (KBps) chart displays the disk I/O of the 10 hosts in the cluster with the most disk usage.

This chart is located in the Hosts view of the cluster **Performance** tab.

## Table 1-11. Data Counters

| Chart Label | Description |
|---|---|
| *host_name* | Average data I/O rate across all hosts in the cluster.<br>■ Counter: usage<br>■ Stats Type: Rate<br>■ Unit: Kilobytes per second (KBps)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

■ The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-12. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Memory (MB)

The Memory (MB) chart displays consumed memory for the cluster. The chart appears only at collection level 1.

This chart is located in the Home view of the cluster **Performance** tab.

Table 1-13. Data Counters

| Chart Label | Description |
| --- | --- |
| Consumed | Amount of host machine memory used by all powered on virtual machines in the cluster. A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel.<br><br>■ Counter: consumed<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |
| Total | Total amount of machine memory of all hosts in the cluster that is available for virtual machine memory (physical memory for use by the Guest OS) and virtual machine overhead memory.<br><br>Memory Total = Aggregate host machine memory - (VMkernel memory + Service Console memory + other service memory)<br><br>**Note** The totalmb data counter is the same as the effectivemem data counter, which is supported only for backward compatibility.<br><br>■ Counter: totalmb<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

## Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

Table 1-14. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host. |
| 4 | If the cluster is not a DRS cluster, enable DRS. Select the cluster in the object navigator, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |

## Table 1-14. Memory Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 5 | If the cluster is a DRS cluster:<br>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.<br>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 6 | Add more physical memory to one or more hosts. |

## Memory (MB)

The Memory (MB) chart displays memory data counters for clusters. The chart appears at all collection levels except level 1.

### Description

This chart is located in the **Home** view of the cluster **Performance** tab.

**Note** These data counter definitions are for hosts. At the cluster level, the values are collected and totaled. The counter values in the chart represent the aggregate amounts of the host data. The counters that appear in the chart depend on the collection level set for your vCenter Server.

### Table 1-15. Data Counters

| Chart Label | Description |
|---|---|
| Active | Sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMkernel applications. Active memory is estimated by the VMkernel.<br>■ Counter: active<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 2 (4) |
| Balloon | Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host.<br>■ Counter: vmmemctl<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

Table 1-15. Data Counters (continued)

| Chart Label | Description |
| --- | --- |
| Consumed | Amount of machine memory used on the host.<br><br>Consumed memory includes virtual machine memory, service console memory, and VMkernel memory.<br><br>consumed memory = total host memory - free host memory<br><ul><li>Counter: consumed</li><li>Stats Type: Absolute</li><li>Unit: Megabytes (MB)</li><li>Rollup Type: Average (Minimum/Maximum)</li><li>Collection Level: 1 (4)</li></ul> |
| Granted | Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory.<br><br>Granted memory for a host includes the shared memory of each virtual machine on the host.<br><ul><li>Counter: granted</li><li>Stats Type: Absolute</li><li>Unit: Megabytes (MB)</li><li>Rollup Type: Average (Minimum/Maximum)</li><li>Collection Level: 2 (4)</li></ul> |
| Swap Used | Sum of the memory swapped by all powered on virtual machines on the host.<br><ul><li>Counter: swapused</li><li>Stats Type: Absolute</li><li>Unit: Megabytes (MB)</li><li>Rollup Type: Average (Minimum/Maximum)</li><li>Collection Level: 2 (4)</li></ul> |
| Total | Aggregate total memory available to the cluster.<br><ul><li>Counter: totalmb</li><li>Stats Type: Absolute</li><li>Unit: Megabytes (MB)</li><li>Rollup Type: Average (Minimum/Maximum)</li><li>Collection Level: 1 (4)</li></ul> |

Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. This allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. This leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

Table 1-16. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory Consumed

The Memory Consumed chart displays memory usage for the 10 child objects in the cluster with the most consumed memory.

For resource pools and virtual machines in a cluster, this chart is located in the **Resource Pools & Virtual Machines** view of the cluster **Performance** tab. For hosts in a cluster, this chart is located in the **Hosts** view of the cluster **Performance** tab.

## Table 1-17. Data Counters

| Chart Label | Description |
| --- | --- |
| *resource_pool*, *virtual_machine*, or *host* | Amount of machine memory used by all resource pools and virtual machines in the cluster or by all hosts in the cluster, depending on the cluster view.<br><br>Consumed memory includes virtual machine memory, service console memory, and VMkernel memory.<br><br>consumed memory = total host memory - free host memory<br><br>■ Counter: consumed<br>■ Stats Type: Absolute<br>■ Unit: MegaBytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

## Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

## Table 1-18. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host. |
| 4 | If the cluster is not a DRS cluster, enable DRS. Select the cluster in the object navigator, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |
| 5 | If the cluster is a DRS cluster:<br>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.<br>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 6 | Add more physical memory to one or more hosts. |

## Network (Mbps)

The Network (Mbps) chart displays network speed for the 10 hosts in the cluster with the most network usage.

This chart is located in the **Hosts** view of the Cluster **Performance** tab.

## Table 1-19. Data Counters

| Chart Label | Description |
| --- | --- |
| <host> | Average rate at which data is transmitted and received across all NIC instances on the host. |
| | ■ Counter: usage |
| | ■ Stats Type: Rate |
| | ■ Unit: Megabits per second (Mbps) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

### Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note** In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

## Table 1-20. Networking Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |

## Table 1-20. Networking Performance Enhancement Advice (continued)

| # | Resolution |
| --- | --- |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

# Datacenters

The data center charts contain information about CPU, disk, memory, and storage usage for data centers. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (MHz)

The CPU (MHz) chart displays CPU usage for the 10 clusters in the data center with the most CPU usage.

This chart is located in the Clusters view of the Datacenters **Performance** tab.

### Table 1-21. Data Counters

| Chart Label | Description |
| --- | --- |
| <cluster> | Amount of CPU currently in use by the cluster. The active CPU usage is approximately equal to the ratio of the used CPU cycles to the available CPU cycles. |
| | The maximum possible value is the frequency of the processors multiplied by the number of cores. For example, a two-way SMP virtual machine using 4000MHz on a host that has four 2GHz processors is using 50% of the CPU (4000 ÷ 4 × 2000) = 0.5). |
| | ■ Counter: usagemhz |
| | ■ Stats Type: Rate |
| | ■ Unit: MegaHertz (MHz) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

Table 1-22. CPU Performance Enhancement Advice

| # | Resolution |
|---|------------|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If the cluster is not a DRS cluster, enable DRS. Select the cluster, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |
| 3 | If the cluster is a DRS cluster:<br>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.<br>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 4 | Migrate one or more virtual machines to a new host. |
| 5 | Upgrade the physical CPUs or cores on each host in the cluster if necessary. |
| 6 | Enable CPU-saving features, such as TCP Segmentation Offload. |
| 7 | Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs. |

## Memory (MB)

The Memory (MB) chart displays the average amount of consumed memory for the 10 clusters in the data center with the most consumed memory.

This chart is located in the **Clusters** view of the Datacenters **Performance** tab.

Table 1-23. Data Counters

| Chart Label | Description |
|-------------|-------------|
| <cluster> | Amount of host machine memory used by all powered on virtual machines in the cluster.<br>■ Counter: consumed<br>■ Stats Type: Absolute<br>■ Unit: MegaBytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel.

If you experience problems with cluster memory usage, use the thumbnail cluster charts to examine memory usage for each cluster and increase memory resources if needed.

If the cluster is a DRS cluster, check the aggressiveness threshold. If the value is low, increase the threshold. Increasing the threshold might help avoid hot spots in the cluster.

## Space used in GB

The Space in GB chart displays the 10 datastores in the data center with the most used disk space.

This chart is located in the **Storage** view of the Datacenter **Performance** tab.

Table 1-24. Data Counters

| Chart Label | Description |
| --- | --- |
| <datastore> | Amount of used storage space on the 10 datastores with the most used space.<br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: GigaBytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization By File Type

The Space Utilization By File Type chart displays datastore space usage for virtual disks, swap files, snapshot files, and other virtual machine files.

**Note**  This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the data center **Performance** tab.

## Datastore Counters

### Table 1-25. Data Counters

| File Type | Description |
|---|---|
| Virtual Disks | Amount of disk space used by virtual disk files.<br><br>Virtual disk files store the contents of the virtual machine's hard disk drive, including information that you write to a virtual machine's hard disk, such as the operating system, program files, and data files. The files have the extension `.vmdk` and appear as a physical disk drive to a guest operating system.<br><br>**Note**  Delta disks, which also have an extension `.vmdk`, are not included in this file type.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Swap Files | Amount of disk space used by swap files.<br><br>Swap files back up the virtual machine's physical memory.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Snapshots | Amount of disk space used by virtual machine snapshot files.<br><br>Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension `.vmsn`. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Other VM Files | Amount of disk space used by all other virtual machine files, such as configuration files and log files.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Other | Amount of disk space used by all other non-virtual machine files, such as documentation files and backup files. |

Table 1-25. Data Counters (continued)

| File Type | Description |
| --- | --- |
| Free Space | Amount of disk space not currently in use. |
| Total Space | Amount of disk space available to the datastore. This is the datastore capacity. The chart displays the information for datastores but not for data centers. |
| | total space = virtual disk space + swap file space + snapshot space + other VM file space + other space + free space |

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Datastores

The datastore charts contain information about disk usage for datastores. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

### Space in GB

The Space in GB chart displays space usage data counters for datastores.

This chart is located in the **Space** view of the datastore **Performance** tab.

Table 1-26. Data Counters

| Chart Label | Description |
| --- | --- |
| Allocated | Amount of physical space provisioned by an administrator for the datastore. It is the storage size up to which files on the datastore can grow. Allocated space is not always in use.<br><br>■ Counter: provisioned<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |
| Used | Amount of physical datastore space in use.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |
| Capacity | Maximum capacity of the datastore.<br><br>capacity = virtual machine file space + non-virtual machine file space + free space<br><br>**Note** Storage data is collected and updated in the overview charts every 30 minutes. Therefore, if you refresh the datastore, the capacity value might only be updated in the datastore **Summary** tab, and not in the overview charts.<br><br>■ Counter: capacity<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

# Space Utilization By File Type

The Space Utilization by File Type chart displays space used by virtual disks, swap files, snapshot files, and other virtual machine files on the the datastore.

**Note**  This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Space** view of the datastore **Performance** tab.

Datastore Counters

Table 1-27. Data Counters

| File Type | Description |
| --- | --- |
| Virtual Disks | Amount of disk space used by virtual disk files.<br><br>Virtual disk files store the contents of the virtual machine's hard disk drive, including information that you write to a virtual machine's hard disk, such as the operating system, program files, and data files. The files have the extension `.vmdk` and appear as a physical disk drive to a guest operating system.<br><br>**Note**  Delta disks, which also have an extension `.vmdk`, are not included in this file type.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Swap Files | Amount of disk space used by swap files.<br><br>Swap files back up the virtual machine's physical memory.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Snapshots | Amount of disk space used by virtual machine snapshot files.<br><br>Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension `.vmsn`. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |

Table 1-27. Data Counters (continued)

| File Type | Description |
| --- | --- |
| Other VM Files | Amount of disk space used by all other virtual machine files, such as configuration files and log files.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 (4) |
| Other | Amount of disk space used by all other non-virtual machine files, such as documentation files and backup files. |
| Free Space | Amount of disk space not currently in use. |
| Total Space | Amount of disk space available to the datastore. This is the datastore capacity. The chart displays the information for datastores but not for data centers.<br><br>total space = virtual disk space + swap file space + snapshot space + other VM file space + other space + free space |

## Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization by Virtual Machine

The Space Utilization by Virtual Machine chart displays the amount of space used by the five virtual machines with the most space used on the datastore.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by Virtual Machine chart is located in the **Space** view of the datastore **Performance** tab.

Table 1-28. Data Counters

| Chart Label | Description |
|---|---|
| *virtual_machine* | Amount of datastore space used by the five virtual machines with the most used datastore space.<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Storage I/O Control Normalized Latency

This chart displays the normalized latency in microseconds on the datastore, which is the latency that is monitored by Storage I/O Control to detect congestion on the datastore. This metric computes a weighted response time across all hosts and VMs accessing the datastore. I/O count is used as the weight for the response time. It captures the device level latency and does not include any queuing inside the hypervisor storage stack or inside the VM. It is adjusted for I/O size, meaning that high latencies that are the result of very large I/Os are discounted so as not to make the datastore seem slower than it really is. Data for all virtual machines is combined. This chart will display zero values when Storage I/O Control is disabled.

This chart is located in the **Performance** view of the datastore **Performance** tab. The sizeNormalizedDatastoreLatency counter can also be displayed for datastore cluster charts.

Table 1-29. Data Counters

| Chart Label | Description |
|---|---|
| Storage I/O Control Normalized Latency | This is the latency that is monitored by Storage I/O Control to detect congestion on the datastore.<br><br>■ Counter: sizeNormalizedDatastoreLatency<br>■ Stats Type: Absolute<br>■ Unit: Microseconds<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Storage I/O Control Aggregate IOPs

This chart displays the number of I/O operations per second on the datastore, aggregated across all hosts and virtual machines accessing the datastore. The chart displays zero values when Storage I/O Control is disabled.

This chart is located in the **Performance** view of the datastore **Performance** tab. The datastoreIops counter can also be displayed for datastore cluster charts.

Table 1-30. Data Counters

| Chart Label | Description |
| --- | --- |
| Storage I/O Control Aggregate IOPs | Number of I/O operations per second on the datastore, aggregated across all hosts and virtual machines accessing the datastore. |
| | ■ Counter: datastoreIops |
| | ■ Stats Type: Absolute |
| | ■ Unit: Number |
| | ■ Rollup Type: Average |
| | ■ Collection Level: 3 |

## Storage I/O Control Activity

This chart displays the percentage of time during which Storage I/O Control actively controlled the datastore latency.

This chart is located in the **Performance** views of the datastore and datastore cluster **Performance** tabs.

Table 1-31. Data Counters

| Chart Label | Description |
| --- | --- |
| Storage I/O Control Activity | This is the percentage of time during which the Storage I/O Control actively controlled the I/O latency for the datastore. |
| | ■ Counter: siocActiveTimePercentage |
| | ■ Stats Type: Absolute |
| | ■ Unit: Percent |
| | ■ Rollup Type: Average |
| | ■ Collection Level: 3 |

## Average Device Latency per Host

This chart displays the average amount of latency on a host device. The chart displays the ten hosts with the highest device latency.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-32. Data Counters

| Chart Label | Description |
|---|---|
| Average Device Latency per Host | Measures the amount of time, in milliseconds, to complete a SCSI command issued from the physical device.<br>■ Counter: deviceLatency<br>■ Stats Type: Absolute<br>■ Unit: Milliseconds (ms)<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Maximum Queue Depth per Host

This chart displays maximum queue depth that hosts are currently maintaining for the datastore. When Storage I/O is enabled, queue depth can change over time when congestion is detected at the array.

This chart is located in the **Performance** view of the datastore **Performance** tab. The chart displays information about the ten hosts with the highest values.

Table 1-33. Data Counters

| Chart Label | Description |
|---|---|
| Max Queue Depth per Host | Maximum queue depth. Queue depth is the number of commands the SCSI driver queues to the HBA.<br>■ Counter: maxQueueDepth<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Read IOPs per Host

This chart displays the per-host disk read rates for a datastore. The chart displays information about the ten hosts with the highest values.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-34. Data Counters

| Chart Label | Description |
|---|---|
| Read IOPs per Host | Number of disk read commands completed on each disk on the host, per second.<br>Read rate = blocks read per second × block size<br>■ Counter: numberReadAveraged<br>■ Stats Type: Rate<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Write IOPs per Host

This chart displays the per-host disk write rates for a datastore. The chart displays information about the ten hosts with the highest values.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-35. Data Counters

| Chart Label | Description |
| --- | --- |
| Write IOPs per Host | Number of disk write commands completed on each disk on the host, per second. |
| | Write rate = blocks written per second × block size |
| | ■ Counter: numberWriteAveraged |
| | ■ Stats Type: Rate |
| | ■ Unit: Number |
| | ■ Rollup Type: Average |
| | ■ Collection Level: 3 |

## Average Read Latency per Virtual Machine Disk

This chart displays the top ten virtual machine disks with the highest average read latency in milliseconds. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-36. Data Counters

| Chart Label | Description |
| --- | --- |
| Average Read Latency per Virtual Machine Disk | Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request. |
| | Total latency = kernelLatency + deviceLatency. |
| | ■ Counter: totalReadLatency |
| | ■ Stats Type: Absolute |
| | ■ Unit: Milliseconds (ms) |
| | ■ Rollup Type: Average |
| | ■ Collection Level: 3 |

## Average Write Latency per Virtual Machine Disk

This chart displays the top ten virtual machine disks with the highest average write latency in milliseconds. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-37. Data Counters

| Chart Label | Description |
| --- | --- |
| Average Write Latency per Virtual Machine Disk | Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request.<br><br>Total latency = kernelLatency + deviceLatency.<br><br>■ Counter: totalWriteLatency<br>■ Stats Type: Absolute<br>■ Unit: Milliseconds (ms)<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Read IOPs per Virtual Machine Disk

This chart displays the top ten virtual machines with the highest number of read operations. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-38. Data Counters

| Chart Label | Description |
| --- | --- |
| Read IOPs per Virtual Machine Disk | Number of disk read commands completed on each virtual machine disk, per second.<br><br>Read rate = blocks read per second × block size<br><br>■ Counter: numberReadAveraged<br>■ Stats Type: Rate<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Write IOPs per Virtual Machine Disk

This chart displays the ten virtual machines with the highest number of write operations. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 1-39. Data Counters

| Chart Label | Description |
| --- | --- |
| Write IOPs per Virtual Machine Disk | Number of disk write commands completed on each virtual machine disk on the host.<br><br>Write rate = blocks read per second × block size<br><br>■ Counter: numberWriteAveraged<br>■ Stats Type: Rate<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Virtual Machine Observed Latency per Datastore

This chart displays the average datastore latency as observed by the virtual machines.

This chart is located in the **Performance** view of the datastore cluster **Performance** tab.

### Table 1-40. Data Counters

| Chart Label | Description |
| --- | --- |
| VM observed latency report per Datastore | This is the average datastore latency as observed by the virtual machines in the datastore cluster. <br>■ Counter: datastoreVMObservedLatency <br>■ Stats Type: Absolute <br>■ Unit: Microseconds <br>■ Rollup Type: Latest <br>■ Collection Level: 3 |

# Hosts

The hosts charts contain information about CPU, disk, memory, network, and storage usage for hosts. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (%)

The CPU (%) chart displays CPU usage for the host.

This chart is located in the Home view of the Host **Performance** tab.

### Table 1-41. Data Counters

| Chart Label | Description |
| --- | --- |
| Usage | Actively used CPU, as a percentage of the total available CPU, for each physical CPU on the host. <br><br>Active CPU is approximately equal to the ratio of the used CPU to the available CPU. <br><br>Available CPU = # of physical CPUs × clock rate. <br><br>100% represents all CPUs on the host. For example, if a four-CPU host is running a virtual machine with two CPUs, and the usage is 50%, the host is using two CPUs completely. <br>■ Counter: usage <br>■ Stats Type: Rate <br>■ Unit: Percentage (%) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 1 (4) |

## Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the host resources. However, if the value is constantly high, the host is probably lacking the CPU required to meet the demand. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

If performance is impacted, consider taking the actions listed below.

**Table 1-42. CPU Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 4 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |
| 5 | Upgrade the physical CPUs or cores on the host if necessary. |
| 6 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## CPU (MHz)

The CPU (MHz) chart displays CPU usage for the host.

This chart is located in the Home view of the Host **Performance** tab.

**Table 1-43. Data Counters**

| Chart Label | Description |
|---|---|
| Usage | The sum, in megahertz, of the actively used CPU of all powered on virtual machines on a host. |
| | The maximum possible value is the frequency of the processors multiplied by the number of processors. For example, if you have a host with four 2GHz CPUs running a virtual machine that is using 4000MHz, the host is using two CPUs completely. |
| | $4000 \div (4 \times 2000) = 0.50$ |
| | ■ Counter: usagemhz |
| | ■ Stats Type: Rate |
| | ■ Unit: MegaHertz (MHz) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

## Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the host resources. However, if the value is constantly high, the host is probably lacking the CPU required to meet the demand. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

If performance is impacted, consider taking the actions listed below.

**Table 1-44. CPU Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 4 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |
| 5 | Upgrade the physical CPUs or cores on the host if necessary. |
| 6 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## CPU Usage

The CPU Usage chart displays CPU usage of the 10 virtual machines on the host with the most CPU usage.

This chart is located in the Virtual Machines view of the host **Performance** tab.

**Table 1-45. Counters**

| Name | Description |
|---|---|
| *virtual_machine* | Amount of CPU actively being used by each virtual machine on the host. 100% represents all CPUs. |
| | For example, if a virtual machine has one virtual CPU that is running on a host with four CPUs and the CPU usage is 100%, the virtual machine is using one CPU resource. |
| | virtual CPU usage = usagemhz ÷ (number of virtual CPUs × core frequency) |
| | **Note**   This is the host's view of the CPU usage, not the guest operating system view. |
| | ■ Counter: usage |
| | ■ Stats Type: Rate |
| | ■ Unit: Percentage (%). Precision is to 1/100%. A value between 0 and 100. |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the actions listed below.

**Table 1-46. CPU Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's **Virtual Machine** view shows the CPU usage for virtual machines on the host. |
| 4 | Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine. |
| 5 | Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles. |
| 6 | Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time. |
| 7 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 8 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |
| 9 | Upgrade the physical CPUs or cores on the host if necessary. |
| 10 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## Disk (KBps)

The Disk (KBps) chart displays disk I/O of the host.

This chart is located in the Home view of the host **Performance** tab.

**Table 1-47. Data Counters**

| Chart Label | Description |
|---|---|
| Usage | Average data I/O rate across all LUNs on the host. <br> ■ Counter: usage <br> ■ Stats Type: Rate <br> ■ Unit: Kilobytes per second (KBps) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 1 (4) |

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

### Table 1-48. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access.<br><br>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |

Table 1-48. Disk I/O Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk Rate (KBps)

The Disk Rate chart displays disk read and write rates for LUNs on a host, including average rates.

This chart is located in the **Home** view of the host **Performance** tab.

Table 1-49. Data Counters

| Chart Label | Description |
|---|---|
| Read | Number of disk read commands completed on each disk on the host, per second. The aggregate number of all disk read commands is also displayed in the chart.<br><br>Read rate = blocksRead per second × blockSize<br><br>■ Counter: read<br>■ Stats Type: Rate<br>■ Unit: Kilobytes per second (KBps)<br>■ Rollup Type: Average<br>■ Collection Level: 3 |
| Write | Number of disk write commands completed on each disk on the host, per second. The aggregate number of all disk write commands is also displayed in the chart.<br><br>Write rate = blocksWritten per second × blockSize<br><br>■ Counter: write<br>■ Stats Type: Rate<br>■ Unit: Kilobytes per second (KBps)<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

### Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-50. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. |
| | To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |

## Table 1-50. Disk I/O Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk Requests (Number)

The Disk Requests chart displays disk usage for the host.

This chart is located in the **Home** view of the host **Performance** tab.

## Table 1-51. Data Counters

| Chart Label | Description |
|---|---|
| Read Requests | Number of disk read commands completed on each LUN on the host. The aggregate number of all disk read commands is also displayed in the chart.<br>■ Counter: numberRead<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Summation<br>■ Collection Level: 3 |
| Write Requests | Number of disk write commands completed on each LUN on the host. The aggregate number of all disk write commands is also displayed in the chart.<br>■ Counter: numberWrite<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Summation<br>■ Collection Level: 3 |

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

■ The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-52. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. <br><br> To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk (Number)

The Disk (Number) chart displays maximum queue depth for the top ten LUNs on a host.

This chart is located in the **Home** view of the host **Performance** tab.

Table 1-53. Data Counters

| Chart Label | Description |
| --- | --- |
| Maximum Queue Depth | Maximum queue depth. Queue depth is the number of commands the SCSI driver queues to the HBA.<br>■ Counter: maxQueueDepth<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 1 |

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

■ The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

■ The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

■ The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-54. Disk I/O Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access.<br>To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |

## Table 1-54. Disk I/O Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk (ms)

The Disk (ms) chart displays the amount of time taken to process commands on a host.

This chart is located in the **Home** view of the host **Performance** tab.

## Table 1-55. Data Counters

| Chart Label | Description |
|---|---|
| Highest Disk Latency | Highest latency value of all disks used by the host. |
| | Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request. |
| | Total latency = kernelLatency + deviceLatency. |
| | ■ Counter: maxTotalLatency |
| | ■ Stats Type: Absolute |
| | ■ Unit: Milliseconds (ms) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

### Table 1-56. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|------------|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. |
|  | To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |

Table 1-56. Disk I/O Performance Enhancement Advice (continued)

| # | Resolution |
|---|-----------|
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk (KBps)

The Disk (KBps) chart displays disk usage for the 10 virtual machines on the host with the most disk usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

Table 1-57. Data Counters

| Chart Label | Description |
|-------------|-------------|
| *virtual_machine* | Sum of the data read from the virtual machine.<br>■ Counter: usage<br>■ Stats Type: Rate<br>■ Unit: KiloBytes per second (KBps)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

■ The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-58. Disk I/O Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. <br><br> To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Memory (%)

The Memory (%) chart displays host memory usage.

This chart is located in the **Home** view of the host **Performance** tab.

## Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. This allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. This leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

### Table 1-59. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (Balloon)

The Memory (Balloon) chart displays balloon memory on a host.

This chart is located in the **Home** view of the host **Performance** tab.

Table 1-60. Data Counters

| Chart Label | Description |
| --- | --- |
| Balloon | Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host. <br> ■ Counter: vmmemctl <br> ■ Stats Type: Absolute <br> ■ Unit: Megabytes (MB) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 1 (4) |

## Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. This allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. This leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

Table 1-61. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |

### Table 1-61. Memory Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (MBps)

The Memory (MBps) chart displays the swap in and swap out rates for a host.

This chart is located on the **Home** view of the Host **Performance** tab.

### Table 1-62. Data Counters

| Chart Label | Description |
|---|---|
| swapinRate | Average rate at which memory is swapped in from the host swap file. <br>■ Counter: swapinRate <br>■ Stats Type: Rate <br>■ Unit: MegaBytes per second (MBps) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 1 (4) |
| swapoutRate | Average rate at which memory is swapped out to the host swap file. <br>■ Counter: swapoutRate <br>■ Stats Type: Rate <br>■ Unit: MegaBytes per second (MBps) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 1 (4) |

### Chart Analysis

Host memory must be large enough to accommodate virtual machine workload. Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when there is a spike in virtual machine workload.

However, a consistently high memory usage value (94% or greater) indicates the host does not have the memory resources required to meet the demand. If the memory balloon and swap values are not high, performance is probably not affected. If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host requires more memory resources.

If the host is not lacking memory resources, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host is lacking memory resources or you notice a degredation in performance, consider taking the following actions.

Table 1-63. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (MB)

The Memory (MB) chart displays memory data counters for hosts.

This chart is located in the **Home** view of the host **Performance** tab.

**Note**  Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system. Machine memory is the actual physical RAM in the host.

Not all counters are collected at collection level 1.

Table 1-64. Data Counters

| Chart Label | Description |
|---|---|
| Active | Sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMKernel applications. Active memory is estimated by the VMkernel and is based on the current workload of the host. <br>■ Counter: active <br>■ Stats Type: Absolute <br>■ Unit: Megabytes (MB) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 2 (4) |
| Balloon | Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host. <br>■ Counter: vmmemctl <br>■ Stats Type: Absolute <br>■ Unit: Megabytes (MB) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 1 (4) |

## Table 1-64. Data Counters (continued)

| Chart Label | Description |
| --- | --- |
| Balloon Target | Sum of the balloon target memory of all powered on virtual machines on the host. |
| | If the balloon target value is greater than the balloon value, the VMkernel inflates the balloon, causing more virtual machine memory to be reclaimed. If the balloon target value is less than the balloon value, the VMkernel deflates the balloon, which allows the virtual machine to consume additional memory if needed. |
| | Virtual machines initiate memory reallocation. Therefore, it is possible to have a balloon target value of 0 and a balloon value greater than 0. |
| | ■ Counter: vmmemctltarget |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 2 (4) |
| Consumed | Amount of machine memory used on the host. |
| | Consumed memory includes virtual machine memory, service console memory, and VMkernel memory. |
| | consumed memory = total host memory - free host memory |
| | ■ Counter: consumed |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |
| Granted | Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory. |
| | Granted memory for a host includes the shared memory of each virtual machine on the host. |
| | ■ Counter: granted |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 2 (4) |

Table 1-64. Data Counters (continued)

| Chart Label | Description |
|---|---|
| Shared Common | Amount of machine memory shared by all powered on virtual machines. |
| | Shared common memory consists of the entire pool of memory from which sharing is possible, including the amount of physical RAM required by the guest memory. |
| | memory shared - memory shared common = amount of memory saved on the host from sharing |
| | ■ Counter: sharedcommon |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 2 (4) |
| Swap Used | Sum of the memory swapped by all powered on virtual machines on the host. |
| | ■ Counter: swapused |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 2 (4) |

## Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. This allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. This leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

## Table 1-65. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory Usage

The Memory Usage chart displays memory usage for the 10 virtual machines on the host with the most memory usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

### Virtual Machine Counters

**Note**  Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system.

## Table 1-66. Data Counters

| Chart Label | Description |
|---|---|
| Usage | Amount of guest physical memory currently in use on the virtual machine.<br>■ Counter: usage<br>■ Stats Type: Absolute<br>■ Unit: Percentage (%)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 1-67. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Network (Mbps)

The Network (Mbps) chart displays network usage for the host.

This chart is located in the **Home** view of the Host **Performance** tab.

Table 1-68. Host Counters

| Chart Label | Description |
| --- | --- |
| Usage | Average rate at which data is transmitted and received across all NIC instances connected to the host.<br>■ Counter: usage<br>■ Stats Type: Rate<br>■ Unit: Megabits per second (Mbps)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

## Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note**  In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

**Table 1-69. Networking Performance Enhancement Advice**

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |

Table 1-69. Networking Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

## Network Rate (Mbps)

The Network Rate chart displays network bandwidth on a host.

The Network Data Transmitted/Received chart for hosts is located in the **Home** view of the Host **Performance** tab.

Table 1-70. Data Counters

| Chart Label | Description |
|---|---|
| Data Receive Rate | Rate at which data is received across the top ten physical NIC instances on the host. This represents the bandwidth of the network. The chart also displays the aggregated data receive rate of all physical NICs. <br>■ Counter: received <br>■ Stats Type: Rate <br>■ Unit: Megabits per second (Mbps) <br>■ Rollup Type: Average <br>■ Collection Level: 3 (4) |
| Data Transmit Rate | Rate at which data is transmitted across the top ten physical NIC instances on the host. This represents the bandwidth of the network. The chart also displays the aggregated data transmit rate of all physical NICs. <br>■ Counter: transmitted <br>■ Stats Type: Rate <br>■ Unit: Megabits per second (Mbps) <br>■ Rollup Type: Average <br>■ Collection Level: 3 (4) |

## Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note**  In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

Table 1-71. Networking Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

## Network Packets (Number)

The Network Packets chart displays the network bandwith on a host.

This chart is located in the **Home** view of the Host **Performance** tab.

Table 1-72. Data Counters

| Chart Label | Description |
| --- | --- |
| Packets Received | Number of network packets received across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs.<br>■ Counter: packetRx<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Summation<br>■ Collection Level: 3 |
| Packets Transmitted | Number of network packets transmitted across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs.<br>■ Counter: packetTx<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Summation<br>■ Collection Level: 3 |

Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note**  In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

## Table 1-73. Networking Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

## Network (Mbps)

The Network (Mbps) chart displays network usage for the 10 virtual machines on the host with the most network usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

## Table 1-74. Data Counters

| Chart Label | Description |
|---|---|
| <virtual machine> | Sum of the data transmitted and received across all virtual NIC instances connected to the virtual machine.<br>■ Counter: usage<br>■ Stats Type: Rate<br>■ Unit: Megabits per second (Mbps)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note**  In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

### Table 1-75. Networking Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

# Resource Pools

The resource pool charts contain information about CPU and memory usage for resource pools. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (MHz)

The CPU (MHz) chart displays CPU usage in the resource pool or vApp.

This chart is located in the Home view of the Resource Pool or vApp **Performance** tab.

### Counters

Table 1-76. Data Counters

| Chart Label | Description |
| --- | --- |
| Usage | CPU usage is the sum of the average CPU usage values of the virtual machines in the resource pool or vApp. <br> CPU usage = number of cores * CPU frequency <br> ■ Counter: usagemhz <br> ■ Stats Type: Rate <br> ■ Unit: Megahertz (MHz) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 1 (4) |

### Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the resources available. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines in the resource pool. Generally, if the CPU usage value for a virtual machine is above 90% and the CPU ready value for a virtual machine is above 20%, performance is impacted.

If performance is impacted, consider taking the following actions.

Table 1-77. CPU Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | Deploy single-threaded applications on uniprocessor virtual machines instead of SMP virtual machines. |
| 3 | Migrate one or more virtual machines to a new host. |
| 4 | Upgrade the physical CPUs or cores on each host if necessary. |
| 5 | Enable CPU-saving features, such as TCP Segmentation Offload. |
| 6 | Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs. |

## CPU Usage

The CPU Usage chart displays CPU usage of virtual machines in the resource pool or vApp. The chart displays the top ten virtual machines with the highest CPU usage.

This chart is located in the Resource Pools & Virtual Machines view of the Resource Pool or vApp **Performance** tab.

### Table 1-78. Data Counters

| Chart Label | Description |
| --- | --- |
| *virtual_machine* | Amount of CPU actively used by virtual machines. <br> ■ Counter: usagemhz <br> ■ Stats Type: Rate <br> ■ Unit: Megahertz (MHz) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 1 (4) |

### Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the actions listed below.

### Table 1-79. CPU Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's **Virtual Machine** view shows the CPU usage for virtual machines on the host. |
| 4 | Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine. |
| 5 | Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles. |
| 6 | Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time. |
| 7 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 8 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |

### Table 1-79. CPU Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 9 | Upgrade the physical CPUs or cores on the host if necessary. |
| 10 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## Memory (MB)

The Memory (MB) chart displays memory usage in the resource pool or vApp.

This chart is located in the **Home** view of the resource pool or vApp **Performance** tab.

### Table 1-80. Data Counters

| Chart Label | Description |
|---|---|
| *resource_pool* or *vApp* | Sum of the active memory used by all virtual machines in the resource pool or vApp. Active memory is determined by the VMkernel and includes overhead memory.<br><br>memory usage = active memory / configured virtual machine memory size<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average<br>■ Collection Level: 1 |

### Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

### Table 1-81. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host. |
| 4 | If the cluster is not a DRS cluster, enable DRS. Select the cluster in the object navigator, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |

## Table 1-81. Memory Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 5 | If the cluster is a DRS cluster: |
| | ■ Increase the number of hosts, and migrate one or more virtual machines to the new host. |
| | ■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 6 | Add more physical memory to one or more hosts. |

## Memory Consumed

The Memory Consumed chart displays the memory performance of all virtual machines in the resource pool or vApp.

This chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

For resource pools and virtual machines in a resource pool or vApp, this chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

### Table 1-82. Data Counters

| Chart Label | Description |
|---|---|
| *virtual_machine* | Amount of host memory used by the virtual machine for its guest operating system's physical memory. Memory overhead is not included in consumed memory. |
| | consumed memory = memory granted - memory saved from page sharing |
| | For example, if a virtual machine has 100MB of memory that is shared equally with three other virtual machines, its portion of the shared memory is 25MB (100MB ÷ 4 VMs). This amount is counted in the memory consumed data counter. |
| | ■ Counter: consumed |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 1-83. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (MB)

The Memory (MB) chart displays memory data counters for resource pools or vApps.

### Description

This chart is located in the **Home** view of the resource pool or vApp **Performance** tab.

**Note**   These data counter definitions are for virtual machines. At the resource pool level, the values are collected and totaled. The counter values in the chart represent the aggregate amounts of the virtual machine data.The counters that appear in the chart depend on the collection level set for your vCenter Server.

## Table 1-84. Data Counters

| Chart Label | Description |
| --- | --- |
| Active | Sum of the active guest physical memory of all powered on virtual machines in the resource pool. <br> ■ Counter: active <br> ■ Stats Type: Absolute <br> ■ Unit: Megabytes (MB) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 2 (4) |
| Balloon | Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines in the resource pool. <br> ■ Counter: vmmemctl <br> ■ Stats Type: Absolute <br> ■ Unit: Megabytes (MB) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 1 (4) |
| Consumed | Amount of physical memory consumed by the virtual machine for guest memory. Consumed memory does not include overhead memory. It includes shared memory and memory that might be reserved, but not actually used. <br> consumed memory = memory granted – memory saved due to memory sharing <br> ■ Counter: consumed <br> ■ Stats Type: Absolute <br> ■ Unit: Megabytes (MB) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 1 (4) |
| Granted | Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory. <br> ■ Counter: granted <br> ■ Stats Type: Absolute <br> ■ Unit: Megabytes (MB) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 2 (4) |
| Shared | Amount of guest physical memory shared with other virtual machines in the resource pool. |
| Swapped | Sum of the memory swapped by all powered on virtual machines in the resource pool. <br> ■ Counter: swapused <br> ■ Stats Type: Absolute <br> ■ Unit: Megabytes (MB) <br> ■ Rollup Type: Average (Minimum/Maximum) <br> ■ Collection Level: 2 (4) |

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-85. Memory Performance Enhancement Advice**

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

# vApps

The vApp charts contain information about CPU and memory usage for vApps. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (MHz)

The CPU (MHz) chart displays CPU usage in the vApp or resource pool.

This chart is located in the **Home** view of the vApp or resource pool **Performance** tab.

Counters

Table 1-86. Data Counters

| Chart Label | Description |
|---|---|
| Usage | CPU usage is the sum of the average CPU usage values of the virtual machines in the resource pool or vApp.<br><br>CPU usage = number of cores * CPU frequency<br><br>■ Counter: usagemhz<br>■ Stats Type: Rate<br>■ Unit: Megahertz (MHz)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the resources available. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines in the resource pool. Generally, if the CPU usage value for a virtual machine is above 90% and the CPU ready value for a virtual machine is above 20%, performance is impacted.

If performance is impacted, consider taking the following actions.

Table 1-87. CPU Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | Deploy single-threaded applications on uniprocessor virtual machines instead of SMP virtual machines. |
| 3 | Migrate one or more virtual machines to a new host. |
| 4 | Upgrade the physical CPUs or cores on each host if necessary. |
| 5 | Enable CPU-saving features, such as TCP Segmentation Offload. |
| 6 | Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs. |

## CPU Usage

The CPU Usage chart displays CPU usage of each virtual machine in the vApp or resource pool.

This chart is located in the **Virtual Machines** view of the vApp or resource pool **Performance** tab.

Table 1-88. Data Counters

| Chart Label | Description |
| --- | --- |
| *virtual_machine* | Amount of CPU actively used by virtual machines.<br>■ Counter: usagemhz<br>■ Stats Type: Rate<br>■ Unit: Megahertz (MHz)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the actions listed below.

Table 1-89. CPU Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's **Virtual Machine** view shows the CPU usage for virtual machines on the host. |
| 4 | Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine. |
| 5 | Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles. |
| 6 | Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time. |
| 7 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 8 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |
| 9 | Upgrade the physical CPUs or cores on the host if necessary. |
| 10 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## Memory (MB)

The Memory (MB) chart displays memory usage in the vApp or resource pool.

This chart is located in the **Home** view of the vApp or resource pool **Performance** tab.

Table 1-90. Data Counters

| Chart Label | Description |
| --- | --- |
| *resource_pool* or *vApp* | Sum of the active memory used by all virtual machines in the resource pool or vApp. Active memory is determined by the VMkernel and includes overhead memory.<br><br>memory usage = active memory / configured virtual machine memory size<br><br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average<br>■ Collection Level: 1 |

Chart Analysis

Memory usage is not an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

Table 1-91. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host. |
| 4 | If the cluster is not a DRS cluster, enable DRS. Select the cluster in the object navigator, and click **Manage > Settings > vSphere DRS > Edit > Turn ON vSphere DRS**. |
| 5 | If the cluster is a DRS cluster:<br>■ Increase the number of hosts, and migrate one or more virtual machines to the new host.<br>■ Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster. |
| 6 | Add more physical memory to one or more hosts. |

## Memory Consumed

The Memory Consumed chart displays the memory performance of the top ten virtual machines in the vApp or resource pool.

This chart is located in the **Virtual Machines** view of the vApp or resource pool **Performance** tab.

For resource pools and virtual machines in a resource pool or vApp, this chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

Table 1-92. Data Counters

| Chart Label | Description |
| --- | --- |
| *virtual_machine* | Amount of host memory used by the virtual machine for its guest operating system's physical memory. Memory overhead is not included in consumed memory. |
| | consumed memory = memory granted - memory saved from page sharing |
| | For example, if a virtual machine has 100MB of memory that is shared equally with three other virtual machines, its portion of the shared memory is 25MB (100MB ÷ 4 VMs). This amount is counted in the memory consumed data counter. |
| | ■ Counter: consumed |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 1-93. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |

Table 1-93. Memory Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

# Virtual Machines

The virtual machine charts contain information about CPU, disk, memory, network, storage, and fault tolerance for virtual machines. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

## CPU (%)

The CPU (%) chart displays virtual machine CPU usage and ready values.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

Table 1-94. Data Counters

| Chart Label | Description |
|---|---|
| Usage | Amount of actively used virtual CPU as a percentage of total available CPU. |
| | CPU usage is the average CPU utilization over all available virtual CPUs in the virtual machine. |
| | For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. |
| | virtual CPU usage = usagemhz ÷ (number of virtual CPUs × core frequency) |
| | **Note** This is the host's view of the CPU usage, not the guest operating system view. |
| | ■ Counter: usage |
| | ■ Stats Type: Rate |
| | ■ Unit: Percentage (%). Precision is to 1/100%. A value between 0 and 100. |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |
| Ready | Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU. |
| | CPU ready time is dependent on the number of virtual machines on the host and their CPU loads. At collection level 1, the average CPU ready time of all virtual CPUs on the virtual machine is displayed. At collection level 3, the average CPU ready time of each virtual CPU is also displayed. |
| | ■ Counter: ready |
| | ■ Stats Type: Rate |
| | ■ Unit: Percentage (%) |
| | ■ Rollup Type: Summation |
| | ■ Collection Level: 1 |

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the actions listed below.

**Table 1-95. CPU Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's **Virtual Machine** view shows the CPU usage for virtual machines on the host. |
| 4 | Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine. |
| 5 | Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles. |
| 6 | Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time. |
| 7 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 8 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |
| 9 | Upgrade the physical CPUs or cores on the host if necessary. |
| 10 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## CPU Usage (MHz)

The CPU Usage (MHz) chart displays virtual machine CPU usage.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

## Table 1-96. Data Counters

| Chart Label | Description |
| --- | --- |
| Usage | Amount of actively used virtual CPU. |
| | **Note** This is the host's view of the CPU usage, not the guest operating system view. |
| | ■ Counter: usagemhz |
| | ■ Stats Type: rate |
| | ■ Unit: MegaHertz (MHz) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

## Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the actions listed below.

## Table 1-97. CPU Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on every virtual machine on the host. |
| 2 | Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required. |
| 3 | Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's **Virtual Machine** view shows the CPU usage for virtual machines on the host. |
| 4 | Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine. |
| 5 | Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles. |
| 6 | Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time. |
| 7 | Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work. |
| 8 | If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host. |
| 9 | Upgrade the physical CPUs or cores on the host if necessary. |
| 10 | Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames. |

## Disk (KBps)

The Disk (KBps) chart displays disk usage for the virtual machine.

It is located in the **Home** view of the virtual machine **Performance** tab.

Table 1-98. Data Counters

| Chart Label | Description |
| --- | --- |
| Usage | Average data I/O rate across all virtual disks on the virtual machine. |
| | ■ Counter: usage |
| | ■ Stats Type: Rate |
| | ■ Unit: Kilobytes per second (KBps) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

## Table 1-99. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|------------|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. |
|   | To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk Rate (KBps)

The Disk Rate chart displays disk usage for the virtual machine.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

Table 1-100. Data Counters

| Chart Label | Description |
|---|---|
| Read | Number of disk read commands completed on each virtual disk on the virtual machine, per second. The aggregate number of all disk read commands per second is also displayed in the chart.<br><br>Read rate = blocksRead per second × blockSize<br><br>■ Counter: read<br>■ Stats Type: Rate<br>■ Unit: Kilobytes per second (KBps)<br>■ Rollup Type: Average<br>■ Collection Level: 2 |
| Write | Number of disk write commands completed on each virtual disk on the virtual machine, per second. The aggregate number of all disk write commands per second is also displayed in the chart.<br><br>Write rate = blocksWritten per second × blockSize<br><br>■ Counter: write<br>■ Stats Type: Rate<br>■ Unit: Kilobytes per second (KBps)<br>■ Rollup Type: Average<br>■ Collection Level: 2 |

## Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

■ The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

■ The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

■ The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-101. Disk I/O Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. <br><br> To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Disk Requests (Number)

The Disk Requests chart displays disk usage for the virtual machine.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

Table 1-102. Data Counters

| Chart Label | Description |
| --- | --- |
| Read Requests | Number of disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all disk read commands is also displayed in the chart.<br><br>■ Counter: numberRead<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Summation<br>■ Collection Level: 3 |
| Write Requests | Number of disk write commands completed on each virtual disk on the virtual machine. The aggregate number of all disk write commands is also displayed in the chart.<br><br>■ Counter: numberWrite<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Summation<br>■ Collection Level: 3 |

**Chart Analysis**

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running at that time.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

■ The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value should be 0-1 milliseconds. If the value is greater than 4ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

■ The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.

■ The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the actions listed below.

Table 1-103. Disk I/O Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this might require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. |
| | To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur. |
| 2 | Defragment the file systems on all guests. |
| 3 | Disable antivirus on-demand scans on the VMDK and VMEM files. |
| 4 | Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput. |
| 5 | Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts. |
| 6 | Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency. |
| 7 | Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see *vSphere Storage*. |
| 8 | For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use. |
| 9 | On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file. |
| 10 | If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links. |
| 11 | For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented. |
| 12 | Use the most current host hardware. |

## Virtual Disk Requests (Number)

The Virtual Disk Requests chart displays virtual disk usage for the virtual machine.

After you click **Overview** on the **Performance** tab of the virtual machine, you can view this chart by selecting **Home** from the **View** drop-down menu. It is available at collection (display) levels 3 and 4.

Table 1-104. Data Counters

| Chart Label | Description |
| --- | --- |
| Read Requests | Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands is also displayed in the chart.<br>■ Counter: numberRead<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 2 |
| Write Requests | Number of virtual disk write commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk write commands is also displayed in the chart.<br>■ Counter: numberWrite<br>■ Stats Type: Absolute<br>■ Unit: Number<br>■ Rollup Type: Average<br>■ Collection Level: 2 |

## Virtual Disk Rate (KBps)

The Virtual Disk Rate chart displays virtual disk usage rate for the virtual machine.

After you click **Overview** on the **Performance** tab of the virtual machine, you can view this chart by selecting **Home** from the **View** drop-down menu. It is available only at collection levels 3 and 4.

Table 1-105. Data Counters

| Chart Label | Description |
| --- | --- |
| Read Requests | Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands per second is also displayed in the chart. Read rate = blocksRead per second × blockSize<br>■ Counter: read<br>■ Stats Type: Rate<br>■ Unit: KiloBytes per second (KBps)<br>■ Rollup Type: Average<br>■ Collection Level: 3 |
| Write Requests | Number of virtual disk write commands completed on each virtual disk on the virtual machine per second. The aggregate number of all virtual disk write commands per second is also displayed in the chart. Write rate = blocksWritten per second × blockSize<br>■ Counter: write<br>■ Stats Type: Rate<br>■ Unit: KiloBytes per second (KBps)<br>■ Rollup Type: Average<br>■ Collection Level: 3 |

## Memory (%)

The Memory (%) chart monitors virtual machine memory usage.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

### Virtual Machine Counters

**Note**   Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system.

Table 1-106. Data Counters

| Chart Label | Description |
| --- | --- |
| Usage | Amount of guest physical memory currently in use on the virtual machine. |
| | ■ Counter: usage |
| | ■ Stats Type: Absolute |
| | ■ Unit: Percentage (%) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 1-107. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (MB)

The Memory (MB) chart displays virtual machine balloon memory.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

Table 1-108. Data Counters

| Chart Label | Description |
|---|---|
| Balloon | Amount of guest physical memory reclaimed from the virtual machine by the balloon driver. |
| | ■ Counter: vmmemctl |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

**Table 1-109. Memory Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (MBps)

The Memory (MBps) chart displays virtual machine memory swap rates.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

**Table 1-110. Data Counters**

| Chart Label | Description |
|---|---|
| swapinRate | Average rate at which memory is swapped into the virtual machine.<br>■ Counter: swapinRate<br>■ Stats Type: Rate<br>■ Unit: MegaBytes per second (MBps)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |
| swapoutRate | Average rate at which memory is swapped out of the virtual machine.<br>■ Counter: swapoutRate<br>■ Stats Type: Rate<br>■ Unit: MegaBytes per second (MBps)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 1 (4) |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If there is sufficient swap space, a high balloon value is not a performance issue. However, if the swapin and swapout values for the host are large, the host is probably lacking the memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. The host might require more memory resources. If it does not, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If memory usage is high or you notice degredation in performance, consider taking the following actions.

Table 1-111. Memory Performance Enhancement Advice

| # | Resolution |
|---|------------|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Memory (MB)

The Memory (MB) chart displays memory data counters for virtual machines.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It appears only at collection levels 2, 3, and 4.

In the descriptions below, guest physical memory refers to the virtual-hardware memory presented to a virtual machine for its guest operating system. Machine memory is actual physical RAM in the host. Note that not all counters are collected at collection level 1.

**Table 1-112. Data Counters**

| Chart Label | Description |
| --- | --- |
| Active | Amount of guest physical memory in use by the virtual machine. |
| | Active memory is estimated by VMkernel statistical sampling and represents the actual amount of memory the virtual machine needs. The value is based on the current workload of the virtual machine. |
| | ■ Counter: active |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 2 (4) |
| Balloon | Amount of guest physical memory reclaimed from the virtual machine by the balloon driver. |
| | ■ Counter: vmmemctl |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |
| Balloon Target | Desired amount of virtual machine balloon memory. |
| | Balloon target memory is estimated by the VMkernel. |
| | If the balloon target amount is greater than the balloon amount, the VMkernel inflates the balloon amount, which reclaims more virtual machine memory. If the balloon target amount is less than the balloon amount, the VMkernel deflates the balloon, which allows the virtual machine to reallocate memory when needed. |
| | ■ Counter: vmmemctltarget |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 2 (4) |
| Consumed | Amount of guest physical memory consumed by the virtual machine for guest memory. |
| | Consumed memory does not include overhead memory. It includes shared memory and memory that might be reserved, but not actually used. |
| | consumed memory = memory granted - memory saved due to memory sharing |
| | ■ Counter: consumed |
| | ■ Stats Type: Absolute |
| | ■ Unit: Megabytes (MB) |
| | ■ Rollup Type: Average (Minimum/Maximum) |
| | ■ Collection Level: 1 (4) |

**Table 1-112. Data Counters (continued)**

| Chart Label | Description |
|---|---|
| Shared | Amount of guest physical memory available for sharing. Memory sharing occurs through transparent page sharing. <br>■ Counter: shared <br>■ Stats Type: Absolute <br>■ Unit: Megabytes (MB) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 2 (4) |
| Swapped | The amount of guest physical memory swapped out to the disk by the VMkernel. This data counter measures VMkernel swapping and not to guest OS swapping. <br>swapped = swapout – swapin <br><br>**Note**  In some cases, vMotion can skew these values and cause a virtual machine to arrive on a host with some memory already swapped out. As a result, the swapped value can be greater than the swapout – swapin value. <br><br>■ Counter: swapped <br>■ Stats Type: Absolute <br>■ Unit: Megabytes (MB) <br>■ Rollup Type: Average (Minimum/Maximum) <br>■ Collection Level: 2 (4) |

## Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

## Table 1-113. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

## Network (Mbps)

The Network (Mbps) chart displays network bandwidth for the virtual machine.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

## Table 1-114. Virtual Machine Counters

| Chart Label | Description |
|---|---|
| Usage | Average rate at which data is transmitted and received across all virtual NIC instances connected to the virtual machine. <ul><li>Counter: usage</li><li>Stats Type: Rate</li><li>Unit: Megabits per second (Mbps)</li><li>Rollup Type: Average (Minimum/Maximum)</li><li>Collection Level: 1 (4)</li></ul> |

### Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note** In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

**Table 1-115. Networking Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

## Network Rate (Mbps)

The Network Rate chart displays network usage for virtual machines.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab. It appears only at collection levels 3 and 4.

Table 1-116. Data Counters

| Chart Label | Description |
| --- | --- |
| Data Receive Rate | Rate at which data is received across each virtual NIC instance on the virtual machine.<br>■ Counter: received<br>■ Stats Type: Rate<br>■ Unit: Megabits per second (Mbps)<br>■ Rollup Type: Average<br>■ Collection Level: 2 (4) |
| Data Transmit Rate | Rate at which data is transmitted across each virtual NIC instance on the virtual machine.<br>■ Counter: transmitted<br>■ Stats Type: Rate<br>■ Unit: Megabits per second (Mbps)<br>■ Rollup Type: Average<br>■ Collection Level: 2 (4) |

## Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note** In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

## Table 1-117. Networking Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

## Network Packets (Number)

The Network Packets monitors network bandwidth for virtual machines.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab. It appears only at collection levels 3 and 4.

## Table 1-118. Data Counters

| Chart Label | Description |
|---|---|
| Packets Transmitted | Number of network packets transmitted across the top ten virtual NIC instances on the virtual machine. The chart also displays the aggregated value for each NIC. <ul><li>Counter: packetTx</li><li>Stats Type: Absolute</li><li>Unit: Number</li><li>Rollup Type: Summation</li><li>Collection Level: 3</li></ul> |
| Packets Received | Number of network packets received across the top ten virtual NIC instances on the virtual machine. The chart also displays the aggregated value for each NIC. <ul><li>Counter: packetRx</li><li>Stats Type: Absolute</li><li>Unit: Number</li><li>Rollup Type: Summation</li><li>Collection Level: 3</li></ul> |

## Chart Analysis

Network performance depends on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

**Note** In some instances, large packets might result in high network latency. To check network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the actions listed below.

**Table 1-119. Networking Performance Enhancement Advice**

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. |
| 2 | If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance. |
| 3 | If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network. |
| 4 | Assign each physical NIC to a port group and a vSwitch. |
| 5 | Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks. |
| 6 | Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch. |
| 7 | If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable. |
| 8 | Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch. |

Table 1-119. Networking Performance Enhancement Advice (continued)

| # | Resolution |
| --- | --- |
| 9 | Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode. |
| 10 | Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible. |

## Space in GB

The Space in GB chart displays space utilization data counters for virtual machines.

This chart is located in the **Storage** view of the virtual machine **Performance** tab.

Table 1-120. Data Counters

| Chart Label | Description |
| --- | --- |
| Allocated | Total amount of logical datastore space provisioned by an administrator for the virtual machine. It is the storage size up to which the virtual machine files on datastores can grow. This includes log files, VMX files, and other miscellaneous files. Allocated space is not always in use.<br>■ Counter: provisioned<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |
| Used | Amount of physical datastore space in use by the virtual machine files.<br>■ Counter: used<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |
| Not Shared | Amount of datastore space that belongs only to this virtual machine and is not shared with other virtual machines. Only unshared space is guaranteed to be reclaimed for the virtual machine if, for example, it is moved to a different datastore and then back again.The value is an aggregate of all unshared space for the virtual machine, across all datastores.<br>■ Counter: unshared<br>■ Stats Type: Absolute<br>■ Unit: Gigabytes (GB)<br>■ Rollup Type: Latest<br>■ Collection Level: 1 |

## Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization by Datastores

The Space Utilization by Datastores chart displays the amount of space used by a virtual machine on different datastores in the data center.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by Datastores chart is located in the **Storage** view of the virtual machine **Performance** tab.

Table 1-121. Data Counters

| Chart Label | Description |
| --- | --- |
| *datastore_name* | Amount of disk space in the datastore currently in use by the virtual machine. |
| | ■ Counter: used |
| | ■ Stats Type: Absolute |
| | ■ Unit: Gigabytes (GB) |
| | ■ Rollup Type: Latest |
| | ■ Collection Level: 1 |

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Space Utilization By File Type

The Space Utilization by File Type chart displays the datastore usage by virtual machine files.

**Note** This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the virtual machine **Performance** tab.

### Datastore counters

Table 1-122. Data Counters

| File Type | Description |
| --- | --- |
| Virtual Disks | Amount of disk space used by virtual disk files. |
| | Virtual disk files store the contents of the virtual machine's hard disk drive, including information that you write to a virtual machine's hard disk - the operating system, program files, and data files. The files have the extension `.vmdk` and appear as a physical disk drive to a guest operating system. |
| | **Note**  Delta disks, which also have an extension `.vmdk`, are not included in this file type. |
| | ■ Counter: used |
| | ■ Stats Type: Absolute |
| | ■ Unit: Gigabytes (GB) |
| | ■ Rollup Type: Latest |
| | ■ Collection Level: 1 (4) |
| Swap Files | Amount of disk space used by swap files. |
| | Swap files back up the virtual machine's physical memory. |
| | ■ Counter: used |
| | ■ Stats Type: Absolute |
| | ■ Unit: Gigabytes (GB) |
| | ■ Rollup Type: Latest |
| | ■ Collection Level: 1 (4) |
| Snapshots | Amount of disk space used by virtual machine snapshot files. |
| | Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension `.vmsn`. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken. |
| | ■ Counter: used |
| | ■ Stats Type: Absolute |
| | ■ Unit: GigaBytes (GB) |
| | ■ Rollup Type: Latest |
| | ■ Collection Level: 1 (4) |
| Other VM Files | Amount of disk space used by all other virtual machine files, such as configuration files and log files. |
| | ■ Counter: used |
| | ■ Stats Type: Absolute |
| | ■ Unit: Gigabytes (GB) |
| | ■ Rollup Type: Latest |
| | ■ Collection Level: 1 (4) |
| Total Space | Amount of disk space used by the virtual machine. |
| | total space = virtual disk space + swap file space + snapshot space + other VM file space |

### Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming a lot of datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface. For information about consolidating the data center, see the vSphere documentation.

## Fault Tolerance Performance Counters

The fault tolerance charts contain information about CPU and memory for fault tolerant virtual machines.

**Note** The performance charts and help topics for fault tolerance are available only when you have vSphere Fault Tolerance enabled. If you select a link for a secondary virtual machine in the thumbnail section of the Resource Pools and Virtual Machines view of the cluster Performance tab, the navigation in the inventory updates to the primary virtual machine. This occurs because secondary machines are not displayed in the inventory.

### CPU (MHz)

The CPU (MHz) chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the virtual machine **Peformance** tab. It is available only at collection levels 3 and 4.

Table 1-123. Data Counters

| Name | Description |
|---|---|
| Usage | The average amount of virtual CPU, per CPU instance, in use on the primary and secondary fault tolerant virtual machines. |
|  | ■ Counter: usagemhz |
|  | ■ Stats Type: Rate |
|  | ■ Unit: Megahertz (MHz) |
|  | ■ Rollup Type: Average (Minimum/Maximum) |
|  | ■ Collection Level: 3 (4) |

### Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

### Table 1-124. CPU Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level. |
| 2 | Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations. |
| 3 | Verify that the network and datastore connections for both virtual machines are similar. |
| 4 | Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion. |
| 5 | Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion. |

### CPU System Time for vCPU (%)

The CPU System Time chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Peformance** tab. It is available only at collection levels 3 and 4.

### Table 1-125. Data Counters

| Chart Label | Description |
|---|---|
| System | Amount of time spent on system processes on each virtual CPU in the virtual machine. |
| | **Note**  This is the host view of the CPU usage, not the guest operating system view. |
| | ■ Counter: system |
| | ■ Stats Type: Delta |
| | ■ Unit: Percentage (%) |
| | ■ Rollup Type: Summation |
| | ■ Collection Level: 3 |

### Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

### Table 1-126. CPU Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level. |
| 2 | Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations. |

## Table 1-126. CPU Performance Enhancement Advice (continued)

| # | Resolution |
|---|---|
| 3 | Verify that the network and datastore connections for both virtual machines are similar. |
| 4 | Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion. |
| 5 | Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion. |

### CPU Used Time for vCPU (%)

The CPU Used Time chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Peformance** tab. The chart is available only for collection levels 3 and 4.

## Table 1-127. Data Counters

| Chart Label | Description |
|---|---|
| used | Amount of used virtual CPU as a percentage of total available CPU on the primary and secondary virtual machines. <br> A high value indicates excessive use of CPU resources. <br> ■ Counter: used <br> ■ Stats Type: Delta <br> ■ Unit: Percentage (%) <br> ■ Rollup Type: Summation <br> ■ Collection Level: 1 |

### Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

## Table 1-128. CPU Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level. |
| 2 | Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations. |
| 3 | Verify that the network and datastore connections for both virtual machines are similar. |

## Table 1-128. CPU Performance Enhancement Advice (continued)

| # | Resolution |
|---|------------|
| 4 | Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion. |
| 5 | Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion. |

### Memory Active (MB)

The Memory Active chart displays active memory usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Peformance** tab. It is not available at collection level 1.

## Table 1-129. Data Counters

| Chart Label | Description |
|-------------|-------------|
| Active | Amount of guest physical memory in use by the fault tolerant virtual machine. Active memory is estimated by VMkernel statistical sampling and represents the actual amount of memory the virtual machine needs. Additional, unused memory may be swapped out or ballooned with no performance impact.<br><br>■ Counter: active<br>■ Stats Type: Absolute<br>■ Unit: Megabytes (MB)<br>■ Rollup Type: Average (Minimum/Maximum)<br>■ Collection Level: 2 (4)<br><br>Make sure that the primary and secondary virtual machines have enough memory. If the secondary system is not provisioned well, it might slow down performance of the primary virtual machine or fail. |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 1-130. Memory Performance Enhancement Advice

| # | Resolution |
| --- | --- |
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

### Memory Swapout (MB)

The Memory Swapout chart displays the swapout memory usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Peformance** tab. It is not available at collection level 1.

Table 1-131. Data Counters

| Chart Label | Description |
| --- | --- |
| Swapout | Amount of machine memory written to the VMkernel swap file. |
| | ▪ Counter: swapout |
| | ▪ Stats Type: Absolute |
| | ▪ Unit: MegaBytes |
| | ▪ Rollup Type: Average (Minimum/Maximum) |
| | ▪ Collection Level: 2 (4) |
| | Make sure that the primary and secondary virtual machines have enough memory and that the swapout value is not high. If the secondary system is not provisioned well, it might slow down performance of the primary virtual machine or fail. |

### Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If there is sufficient swap space, a high balloon value is not a performance issue. However, if the swapin and swapout values for the host are large, the host is probably lacking the memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. The host might require more memory resources. If it does not, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If memory usage is high or you notice degredation in performance, consider taking the following actions.

Table 1-132. Memory Performance Enhancement Advice

| # | Resolution |
|---|---|
| 1 | Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance. |
| 2 | Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance. |
| 3 | Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines. |
| 4 | If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host. |
| 5 | Migrate one or more virtual machines to a host in a DRS cluster. |
| 6 | Add physical memory to the host. |

# Working with Advanced and Custom Charts

Use advanced charts, or create your own custom charts, to see more performance data. Advanced charts can be useful when you are aware of a problem but need more statistical data to pinpoint the source of the trouble.

Advanced charts include the following features:

- More information. Hover over a data point in a chart and details about that specific data point are displayed.

- Customizable charts. Change chart settings. Save custom settings to create your own charts.

- Export to spreadsheet.

- Save to image file or spreadsheet.

## View Advanced Performance Charts

Advanced charts support data counters that are not supported in other performance charts.

**Procedure**

**1** Navigate to an inventory object in the vSphere Web Client.

**2** Click the **Monitor** tab, and click **Performance**.

**3** Click **Advanced**.

**4** (Optional) To view a different chart, select an option from the **View** list.

The amount of historical data displayed in a chart depends on the collection interval and statistics level set for vCenter Server.

## Change Advanced Chart Settings

You can customize a performance chart by specifying the objects to monitor, the counters to include, the time range, and chart type. You can customize preconfigured chart views and create new chart views.

**Procedure**

**1** Select an inventory object in the vSphere Web Client.

**2** Click the **Monitor** tab, and click **Performance**.

**3** Click **Advanced**.

**4** Click **Chart Options**.

**5** In Chart Metrics, select a metric group for the chart.

**6** Select a time range for the metric group.

Time range options are not active unless you select **Custom interval** in the **Timespan** menu.

If you choose **Custom interval**, do one of the following.

- Select **Last** and set the number of hours, days, weeks, or months for the amount of time to monitor the object.

- Select **From** and select the beginning and end dates.

You can also customize the time range options by customizing the statistics collection interval setting.

**7** In Target Objects, select the inventory objects to display in the chart.

You can also specify the objects using the **All** or **None** buttons.

**8** Select the chart type.

When selecting the stacked graph option, consider the following.

- You can select only one item from the list of measurements.

- Per-virtual-machine stacked graphs are available only for hosts.

- Click a counter description name to display information about the counter's function and whether the selected metric can be stacked for per-virtual-machine graphs.

**9**    In Counters, select the data counters to display in the chart.

You can also specify counters using the **All** or **None** buttons. The **All** button is inactive when there are more than two different counter units for the corresponding metric group.

**10**    Click **OK**.

## Create a Custom Advanced Chart

You can create your own charts by saving customized chart settings. New charts are added to the **View** menu and will appear there only when charts for the selected object are being displayed.

**Procedure**

**1**    Select an inventory object in the vSphere Web Client.

**2**    Click the **Monitor** tab, click **Performance**, and navigate to the Chart Options dialog of a chart.

**3**    Customize chart settings.

**4**    Click **Save Options As...**.

**5**    Enter a name for your settings.

**6**    Click **OK**.

**Results**

The chart settings are saved and an entry for your chart is added to the **View** menu.

## Delete a Custom Advanced Chart View

You can delete custom chart views from the vSphere Web Client.

**Procedure**

**1**    Select an inventory object in the vSphere Web Client.

**2**    Click the **Monitor** tab, and click **Performance**.

**3**    Click **Advanced**.

**4**    Click **Chart Options**.

**5**    Select a chart and click **Delete Options**.

**6**    Click **OK** to confirm deletion.

The chart is deleted, and it is removed from the **View** menu.

## Save Chart Data to a File

You can save data from the advanced performance charts to a file in various graphic formats or in comma-separated values (CSV) format.

Procedure

**1**   In the vSphere Web Client select an inventory object.

**2**   Click the **Monitor** tab, and click **Performance**.

**3**   Click **Advanced**.

**4**   Click the **Export** icon (⬇).

**5**   Select a file type.

| Option | Description |
| --- | --- |
| **To PNG** | Exports a bitmap image in the PNG format. |
| **To JPEG** | Exports a bitmap image in the JPEG format. |
| **To CSV** | Exports plain-text data in the CSV format. |

**6**   Enter a name and location for the file.

**7**   Click **Save**.

Results

The file is saved to the location and format you specified.

# Troubleshoot and Enhance Performance

This section presents tips for identifying and solving performance problems.

The suggestions in this section are not meant to be a comprehensive guide to diagnosing and troubleshooting problems in the virtual environment. It is meant to provide information about some common problems that can be solved without contacting VMware Technical Support.

## Solutions for Consistently High CPU Usage

Temporary spikes in CPU usage indicate that you are making the best use of CPU resources. Consistently high CPU usage might indicate a problem. You can use the CPU performance charts to monitor CPU usage for hosts, clusters, resource pools, virtual machines, and vApps.

Problem

- Host CPU usage constantly is high. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

- Virtual machine CPU usage is above 90% and the CPU ready value is above 20%. Application performance is impacted.

Cause

- The host probably is lacking the CPU resources required to meet the demand.

- There might be too many virtual CPUs relative to the number of regular CPUs.

- There might be an IO storage or networking operation that places the CPU in a wait state.

- The Guest OS generates too much load for the CPU.

**Solution**

- Verify that VMware Tools is installed on every virtual machine on the host.

- Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked bar chart on the host's **Virtual Machine** view shows the CPU usage for all virtual machines on the host.

- Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.

- Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.

- Increase the amount of memory allocated to the virtual machine. This action decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.

- Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.

- If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.

- Upgrade the physical CPUs or cores on the host if necessary.

- Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

## Solutions for Memory Performance Problems

Host machine memory is the hardware backing for guest virtual memory and guest physical memory. Host machine memory must be at least slightly larger than the combined active memory of the virtual machines on the host. A virtual machine's memory size must be slightly larger than the average guest memory usage. Increasing the virtual machine memory size results in more overhead memory usage.

**Problem**

- Memory usage is constantly high (94% or greater) or constantly low (24% or less).

- Free memory consistently is 6% or less and swapping frequently occurs.

Cause

- The host probably is lacking the memory required to meet the demand. The active memory size is the same as the granted memory size, which results in memory resources that are not sufficient for the workload. Granted memory is too much if the active memory is constantly low.

- Host machine memory resources are not enough to meet the demand, which leads to memory reclamation and degraded performance.

- The active memory size is the same as the granted memory size, which results in memory resources that are not sufficient for the workload.

Solution

- Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.

- Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.

- Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.

- If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.

- Migrate one or more virtual machines to a host in a DRS cluster.

- Add physical memory to the host.

## Solutions for Storage Performance Problems

Datastores represent storage locations for virtual machine files. A storage location can be a VMFS volume, a directory on Network Attached Storage, or a local file system path. Datastores are platform-independent and host-independent.

Problem

- Snapshot files are consuming a lot of datastore space.

- The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks.

Solution

- Consider consolidating snapshots to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Web Client user interface.

- You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

## Solutions for Disk Performance Problems

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read/write requests, check if any such applications were running at that time.

Problem

- The value for the kernelLatency data counter is greater than 4ms.

- The value for the deviceLatency data counter is greater than 15ms indicates there are probably problems with the storage array.

- The queueLatency data counter measures above zero.

- Spikes in latency.

- Unusual increases in read/write requests.

Cause

- The virtual machines on the host are trying to send more throughput to the storage system than the configuration supports.

- The storage array probably is experiencing internal problems.

- The workload is too high and the array cannot process the data fast enough.

Solution

- The virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.

- Move the active VMDK to a volume with more spindles or add disks to the LUN.

- Increase the virtual machine memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this may require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access.

- Check swap statistics in the guest operating system to verify that virtual machines have adequate memory. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.

- Defragment the file systems on all guests.

- Disable antivirus on-demand scans on the VMDK and VMEM files.

- Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput.

◆ Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.

◆ Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.

◆ Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter.

◆ For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use.

◆ On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the virtual machine's .VMX file.

◆ If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.

◆ For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.

◆ Use the most current hypervisor software.

## Solutions for Poor Network Performance

Network performance is dependent on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. Slow network performance can be a sign of load-balancing problems.

**Problem**

Network problems can manifest in many ways:

- Packets are being dropped.

- Network latency is high.

- Data receive rate is low.

**Cause**

Network problems can have several causes:

- Virtual machine network resource shares are too few.

- Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.

- Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough to handle the load.

**Solution**

- Determine whether packets are being dropped by using `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values. Verify that VMware Tools is installed on each virtual machine.

- Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different virtual switches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

- If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.

- If virtual machines running on the same host communicate with each other, connect them to the same virtual switch to avoid the cost of transferring packets over the physical network.

- Assign each physical NIC to a port group and a virtual switch.

- Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.

- Ensure that the physical NIC capacity is large enough to handle the network traffic on that virtual switch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a virtual switch with a lighter load or to a new virtual switch.

- If packets are being dropped at the virtual switch port, increase the virtual network driver ring buffers where applicable.

- Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch.

- Verify that all NICs are running in full duplex mode. Hardware connectivity issues might result in a NIC resetting itself to a lower speed or half duplex mode.

- Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TCP Segmentation Offload-Jumbo Frames are enabled where possible.

## Empty Performance Charts

No graphics or data are displayed in performance charts.

Problem

When data is missing for a performance chart, the chart appears empty and you see the message `No data available`.

Cause

The causes for missing data in performance charts that are described herein are based on the assumption that the default rollup configuration for the vCenter Server system has not changed. The causes include but are not limited to the following scenarios:

- Metrics introduced in ESXi 5.0 are not available for hosts running earlier versions.

- Data is deleted when you remove or add objects to vCenter Server.

- Performance charts data for inventory objects that were moved to a new site by VMware vCenter Site Recovery Manager is deleted from the old site and not copied to the new site.

- Performance charts data is deleted when you use VMware vMotion across vCenter Server instances.

- Real-time statistics are not available for disconnected hosts or powered off virtual machines.

- Real-time statistics are collected on hosts and aggregated every 5 minutes. After six data points are collected for approximately 30 minutes, they are rolled up to the vCenter Server database to create the 1 Day statistics. 1 Day statistics might not be available for 30 minutes after the current time, depending on when the sample period began.

- The 1 Day statistics are rolled up to create one data point every 30 minutes. If a delay occurs in the rollup operation, the 1 Week statistics might not be available for 1 hour after the current time. It takes 30 minutes for the 1 Week collection interval, plus 30 minutes for the 1 Day collection interval.

- The 1 Week statistics are rolled up to create one data point every two hours. If a delay occurs in the rollup operations, the 1 Month statistics might not be available for 3 hours. It takes 2 hours for the 1 Month collection interval, plus 1 hour for the 1 Week collection interval.

- The 1 Month statistics are rolled up to create one data point every day. If a delay occurs in the rollup operations, the statistics might not be available for 1 day and 3 hours. It takes 1 day for the past year collection interval, plus 3 hours for the past month collection interval. During this time, the charts are empty.

Solution

◆ There is no solution available.

# Using VMware vCenter Operations Manager in the vSphere Web Client

2

The VMware vCenter Operations Manager vApp integrates with the vSphere Web Client to provide a high-level overview of the health and performance of vSphere objects.

A default installation of the vSphere Web Client includes a shortcut to deploying and configuring vCenter Operations Manager. On the **Home** tab of the vSphere Web Client, you can click the **vCenter Operations Manager** button to start the deployment and configuration process. The deployment wizard is pre-configured with a URL address for the OVF template. You need a My VMware account to download the template and deploy the vApp.

You can also install vCenter Operations Manager by downloading and deploying the vApp manually.

Problems that occur during deployment of the vApp by using the shortcut in the vSphere Web Client return error messages, which appear in the OVF deployment wizard.

Configuration steps and installation requirements are identical for both approaches of installing vCenter Operations Manager. See *vApp Deployment and Configuration Guide*.

After deployment and configuration,vCenter Operations Manager can be monitored in the vSphere Web Client interface. See Using Badges to Monitor Objects in the Virtual Environment.

This chapter includes the following topics:

- Deploy vRealize Operations Manager from the Home Tab of the vSphere Web Client

- Attempts to Download the OVF Template of vCenter Operations Manager Fail

- Using Badges to Monitor Objects in the Virtual Environment

## Deploy vRealize Operations Manager from the Home Tab of the vSphere Web Client

To use vRealize Operations Manager, you must first deploy the OVF template and then configure the application.

For more information about configuring vRealize Operations Manager, see the *vApp Deployment and Configuration Guide*.

Prerequisites

- Required privileges:

  - **Extension.Register extension**

  - **vApp.Power On**

  - **Virtual machine.Provisioning.Deploy template**

- Verify that you have a My VMware account.

- Verify that you have a vRealize Operations Manager license key.

- A host running ESX/ESXi 4.0 or later.

- Verify that the time of the ESX/ESXi system that hosts the vRealize Operations Manager vApp and the time of the vCenter Server is synchronized.

- Verify that the virtual machine network that you connect this vApp to has an IP pool and select that network during the OVF deployment.

- Verify that you have installed the Client Integration Plug-in.

Procedure

1. Log in to the vSphere Web Client.

2. Click the **Home** tab, and click **vCenter Operations Manager**.

3. Click **Deploy vRealize Operations Manager**.

   If your environment has multiple vCenter Server instances, select the instance that you want to monitor with vRealize Operations Manager. You can only select vCenter Server instances of version 5.1 or later.

4. In the **Deploy OVF Template** wizard, type your My VMware credentials, and click **Next**.

5. Complete the **Deploy OVF Template** wizard by accepting the default values where applicable.

   During deployment, the wizard contains a pre-populated URL address for the OVF template. You can monitor the deployment progress in the Work in Progress panel.

6. Click **Configure vRealize Operations Manager**.

   The login screen of the configuration dashboard appears in a new browser window.

7. Log in to the vRealize Operations Manager configuration dashboard and perform the required configuration steps.

   a   Change the user name and password.

   b   Register vRealize Operations Manager with a vCenter Server instance.

8. Navigate to **Home > Administration > Licenses > Solutions** and assign a license key for vRealize Operations Manager.

**9** On the **Getting Started** tab of vRealize Operations Manager, click **Open vRealize Operations Manager**.

If more than one vRealize Operations Manager vApps are available, you can select the instance that you want to open.

**Results**

You have deployed and configured the vApp. You can start using vRealize Operations Manager. Health, risk, and efficiency information appears in a panel under the **Summary** tab. The **Health** tab is available on the **Monitor** tab.

**What to do next**

Click **Learn more about vRealize Operation Manager** to navigate to the official documentation page.

## Attempts to Download the OVF Template of vCenter Operations Manager Fail

You might be unable to download the OVF template of vCenter Operations Manager because of loss of connection with the download service.

**Problem**

When you attempt to download the OVF template of vCenter Operations Manager, the download does not start, and the following error message appears:

```
Unable to communicate with the download service. Contact VMware support or visit My VMware to
download the software manually.
```

**Cause**

An internal error has occurred in My VMware or the OVF deployment functionality. You cannot use the automatic download service to get the OVF template of vCenter Operations Manager.

**Solution**

- Go to My VMware, download the OVF template of vCenter Operations Manager, and install it by using the **Deploy OVF Template** wizard in the vSphere Web Client.

- Troubleshoot the loss of connection with the download service.

## Using Badges to Monitor Objects in the Virtual Environment

Badges appear in the vSphere Web Client interface when you register an instance of vCenter Operations Manager with a vCenter Server that you want to monitor.

In the vSphere Web Client interface, you can view the vCenter Operations Manager badges and widgets on the **Summary** tab, and on the **Monitor** tab, under **Health**, for the following objects in the inventory tree.

Table 2-1. Supported Object Types

| Icon | Description |
| --- | --- |
|  | vCenter Server system |
|  | Datacenter |
|  | Cluster |
|  | ESX/ESXi host |
|  | Datastore |
|  | Powered on virtual machine |
|  | Powered off virtual machine |

Your vCenter Operations Manager license determines which badges and widgets you can see in the vSphere Web Client.

Table 2-2. Badges that You See Depending on the Licenses

| vSphere Web Client Component | Foundation License | Standard or Advanced License |
| --- | --- | --- |
| **Summary** tab | Health badge | ■ Health badge<br>■ Risk badge<br>■ Efficiency badge |
| **Health** tab under the **Monitor** tab | ■ Health Trend widget<br>■ Workload badge<br>■ Faults badge | ■ Health Trend widget<br>■ Workload badge<br>■ Anomalies badge<br>■ Faults badge |

If you want to view more details about the causes for a badge score, you can click the **View details** icon  at the upper right of the **Health** tab to open the vCenter Operations Manager user interface.

## The Health Badge

The health rating gives you an overview of the current operational state of an inventory object.

vRealize Operations Manager calculates the Health score by using the scores of the sub-badges that the Health badge contains. Faults are given precedence in the Health score because they describe existing problems, while Workload and Anomalies are combined to identify performance problems. This approach ensures that the score of the Health badge reflects the actual state of the object, without exaggerating or underestimating problems.

The Health score ranges between 0 (bad) and 100 (good). The badge changes its color based on the badge score thresholds that are set by the vRealize Operations Manager administrator.

Table 2-3. Object Health States

| Badge Icon | Description | User Action |
|---|---|---|
| | The health of the object is normal. | No attention required. |
| | The object is experiencing some level of problems. | |
| | The object might have serious problems. | |
| | The object is either not functioning properly or will stop functioning soon. | |
| | No data is available or the object is offline. | |

A vRealize Operations Manager administrator can change the badge score thresholds. For example, a green Health badge can indicate a score above 80 instead of 75, as set by default.

## Health Trend Graph

For virtual machine objects, vCenter Operations Manager displays a health trend graph under the **Health** badge.

The Health Trend graph appears under the **Health** badge when virtual machine objects are selected in the inventory pane. You can use the Health Trend Graph for a quick overview of the current condition of the virtual machine, and to see how it has changed during the past six hours. You cannot change or manipulate this graph.

## The Workload Badge

The analytics algorithms of vRealize Operations Manager combine the metrics that show the demand for resources on a virtual machine or another object in a single value, called the workload.

These metrics include CPU use, memory use, and so on.

vRealize Operations Manager indicates the workload by a colored icon that is based on the defined badge score thresholds.

The Workload score ranges from 0 (good) to over 100 (bad). The badge changes its color based on the badge score thresholds that the vRealize Operations Manager administrator sets.

An object with a workload score of 0 is not being used. An object with a workload score above 100 is trying to access more resources than are available to it. In this case, you might have to allocate more resources to the object or move some tasks to other objects.

Table 2-4. Object Workload States

| Badge Icon | Description | User Action |
|---|---|---|
| | Workload on the object is not excessive. | No attention required. |
| | Object is experiencing some high-resource workloads. | Check and take appropriate action. |
| | Workload on the object is approaching its capacity in at least one area. | Check and take appropriate action as soon as possible. |
| | Workload on the object is at or over its capacity in one or more areas. | Act immediately to avoid or correct problems. |
| | No data is available or the object is offline. | |

A vRealize Operations Manager administrator can change the badge score thresholds. For example, an administrator might change a threshold so that a green Workload badge indicates a score below 80 instead of 85, as set by default.

## The Anomalies Badge

The Anomalies score represents how abnormal the behavior of the object is, based on its historical metrics data.

vRealize Operations Manager calculates dynamic thresholds for each metric that is collected for an object. vRealize Operations Manager also analyses the number of metrics that are violating their dynamic thresholds to determine trends and normal levels of threshold violations. Based on these trends, the Anomalies score is calculated using the total number of threshold violations for all metrics for the selected object and its child objects.

A low Anomalies score indicates that an object is behaving in accordance with its established historical parameters. Most or all of the object metrics, especially its KPIs, are within their thresholds. Because changes in behavior often indicate developing problems, if the metrics of an object go outside the calculated thresholds, the anomalies score for the object grows. As more metrics breach the thresholds, anomalies continue to increase. Violations by KPI metrics increase the Anomalies score more than violations by non-KPI metrics. A high number of anomalies usually indicates a problem or at least a situation that requires your attention.

Anomalies and Workload differ in that workload is computing an absolute measurement of how hard an object is working while anomalies is computing how different from normal the object is behaving. Both Workload and Anomalies are useful when searching for and troubleshooting performance problems.

The Anomalies score ranges between 0 (good) and 100 (bad). The badge changes its color based on the badge score thresholds that are set by the vRealize Operations Manager administrator.

Table 2-5. Object Anomalies States

| Badge Icon | Description | User Action |
|---|---|---|
| 🟢 | The Anomalies score is normal. | No attention required. |
| 🟡 | The Anomalies score exceeds the normal range. | |
| 🟠 | The Anomalies score is very high. | |
| 🔴 | Most of the metrics are beyond their thresholds. This object might not be working properly or might stop working soon. | |
| ❓ | No data is available or the object is offline. | |

A vRealize Operations Manager administrator can change the badge score thresholds. For example, a green Anomalies badge can indicate a score below 60 instead of 50, as set be default.

## The Anomalies Graph

The Anomalies graph provides a visual comparison between the actual number of metrics that have abnormal values and the calculated allowable number of metrics with abnormal values.

### The Abnormal Metric Count

Any object in vCenter Operations Manager can have hundreds or thousands of collected metrics associated with it. At any time, a subset of these metrics can be abnormal or different from the estimated normal values for the metrics. The current number of metrics that have abnormal values is presented with a blue dot on the right of the Anomalies graph. A blue line represents the number of abnormal metrics for the last 6 hours.

### The Noise Line

vCenter Operations Manager uses one of its algorithms to determine how many metrics on an object are usually acting abnormally and create a noise line for the object. The noise line helps to remove the miscellaneous abnormalities, or day to day noise, for an object. Therefore, the more dynamic an object is, the higher its noise level.

In the Anomalies graph, the noise line is presented with a gray line.

If the blue line of the abnormal metric count is far below the noise line, the level of anomalies is normal. If the blue line of abnormal metric count approaches or surpasses the noise line, the object might be experiencing health degradation.

## The Faults Badge

The Faults badge quantifies the severity of problems the selected object is experiencing.

The Faults score is calculated based on events published by the vCenter Server. This score includes events like loss of redundancy in NICs or HBAs, memory checksum errors, HA failover problems, CIM events, and so on. Faults are included in the health score because they require immediate resolution, while items that contribute to the risk score might not be immediate, but still require your attention.

Each resource in vRealize Operations Manager has a faults score ranging from 0 (no faults) to 100 (critical faults). The scores are computed based on the severity of the underlying problems. When more than one fault-related problems exists on the resource, the faults score is based on the most severe problem.

The higher the Faults score, the lower the resulting health for that resource. Resolution of the problem indicated by the Fault will restore the resource's health score.

While the Faults score ranges between 0 to 100, the badge changes color based on the badge score thresholds that are set by the vRealize Operations Manager administrator. For example, a green Faults badge can indicate a score below 40 instead of a score below 25 (the system default).

Unlike other badges in vRealize Operations Manager, the Faults badge does not have an alert generated from its threshold score. Instead, each problem generates its own fault alert, and resolution of the problem both clears or cancels the alert and lowers the badge score.

Table 2-6. Object Faults States

| Badge Icon | Description | User Action |
| --- | --- | --- |
|  | No faults are registered on the selected object. | No attention required. |
|  | Faults of low importance are registered on the selected object. | |
|  | Faults of high importance are registered on the selected object. | |
|  | Faults of critical importance are registered on the selected object. | |
|  | No data is available or the object is offline. | |

## The Risk Badge

The Risk badge indicates potential problems that might eventually degrade the performance of the system. Risk does not necessarily imply a current problem. Risk indicates problems that might require your attention in the near future, but not immediately.

vRealize Operations Manager calculates the risk score using the scores of the sub-badges that the Risk badge contains. The formula that is applied to calculate the risk score is inverse geometric weighted mean.

The overall risk score for an object ranges between 0 (no risk) to 100 (serious risk). The badge changes its color based on the badge score thresholds that are set by the vRealize Operations Manager administrator.

Table 2-7. Object Risk States

| Badge Icon | Description | User Action |
|---|---|---|
|  | The selected object has no current problems. No problems are expected in the future. | No attention required. |
|  | A low chance of future problems exists or a potential problem might occur in the far future. | |
|  | A chance of a more serious problem exists or a problem might occur in the medium-term future. | |
|  | The chances of a serious future problem are high or a problem might occur in the near future. | |
|  | No data is available or the object is offline. | |

A vRealize Operations Manager administrator can change the default badge score thresholds. For example, green can indicate a score below 30 instead of 25.

## The Efficiency Badge

The Efficiency badge helps you identify optimization opportunities in your systems. Efficiency does not tell you about current or future performance problems, but tells you how to run a more efficient datacenter.

Efficiency badge helps you identify optimizations and opportunities to improve your system performance.

The Efficiency score ranges between 0 (bad) and 100 (good). The badge changes its color based on the badge score thresholds that are set by the vRealize Operations Manager administrator.

Table 2-8. Object Efficiency States

| Badge Icon | Description | User Action |
|---|---|---|
|  | The resource use on the selected object is optimal. | No attention required. |
|  | The efficiency is good, but can be improved. Some resources are not fully used. | |
|  | The resources on the selected object are not used in the most optimal way. | |

### Table 2-8. Object Efficiency States (continued)

| Badge Icon | Description | User Action |
| --- | --- | --- |
| | The efficiency is bad. Many resources are wasted. | |
| | No data is available or the object is offline. | |

A vRealize Operations Manager administrator can change the badge score thresholds. For example, a green Efficiency badge can indicate a score above 40 instead of 25, as set by default.

# Monitoring Guest Operating System Performance

3

This section describes how to install and view VMware-specific performance data for virtual machines that run Microsoft Windows operating systems. VMware provides performance counters that enable you to view data on many aspects of guest operating system performance for the Microsoft Windows Perfmon utility.

Some virtualization processes dynamically allocate available resources depending on the status, or utilization rates, of virtual machines in the environment. This can make obtaining accurate information about the resource utilization (CPU utilization, in particular) of individual virtual machines, or applications running within virtual machines, difficult. VMware now provides virtual machine-specific performance counter libraries for the Windows Perfmon utility that enable application administrators to access accurate virtual machine resource utilization statistics from within the Windows Perfmon utility.

You can take advantage of virtualized CPU performance counters to use performance tuning tools inside the guest operating system. See the *vSphere Virtual Machine Administration* documentation.

This chapter includes the following topics:

- Enable Statistics Collection for Guest Operating System Performance Analysis
- View Performance Statistics for Windows Guest Operating Systems

## Enable Statistics Collection for Guest Operating System Performance Analysis

VMware-specific performance objects are loaded into Microsoft Windows Perfmon and enabled when VMware Tools is installed.

To display a performance chart for any performance object, you must add counters. See View Performance Statistics for Windows Guest Operating Systems

## View Performance Statistics for Windows Guest Operating Systems

You can display VMware specific statistics in the Microsoft Windows Perfmon utility.

Prerequisites

Verify that a virtual machine with a Microsoft Windows operating system and VMware Tools is installed.

Procedure

1 Open a console to the virtual machine and log in.

2 Select **Start > Run**.

3 Enter **Perfmon** and press **Enter**.

4 In the Performance dialog box, click **Add** .

5 In the Add Counters dialog box, select **Use local computer counters**.

6 Select a virtual machine performance object.

   Virtual machine performance object names begin with **VM**.

7 Select the counters that you want to display for that object.

8 If the performance object has multiple instances, select the instances you want to display.

9 Click **Add**.

   The Performance dialog box displays data for the selected performance object.

10 Click **Close** to close the Add Counter dialog box and return to the Performance dialog box.

# Monitoring Host Health Status

4

You can use the vSphere Web Client to monitor the state of host hardware components, such as CPU processors, memory, fans, and other components.

The host health monitoring tool allows you to monitor the health of a variety of host hardware components including:

- CPU processors

- Memory

- Fans

- Temperature

- Voltage

- Power

- Network

- Battery

- Storage

- Cable/Interconnect

- Software components

- Watchdog

- PCI devices

- Other

The host health monitoring tool presents data gathered using Systems Management Architecture for Server Hardware (SMASH) profiles. The information displayed depends on the sensors available on your server hardware. SMASH is an industry standard specification providing protocols for managing a variety of systems in the data center. For more information, see http://www.dmtf.org/standards/smash.

You can monitor host health status either by connecting the vSphere Client directly to a host, or by connecting the vSphere Web Client to a vCenter Server system. You can also set alarms to trigger when the host health status changes.

**Note**  The interpretation of hardware monitoring information is specific for each hardware vendor. Your hardware vendor can help you understand the results of the host hardware components monitoring.

This chapter includes the following topics:

- Monitor Hardware Health Status

- Reset Health Status Sensors

# Monitor Hardware Health Status

You can monitor the health status of host hardware in the vSphere Web Client.

**Procedure**

**1** Select a host in the vSphere Web Client navigator.

**2** Click the **Monitor** tab, and click **Hardware Status**.

**3** Select the type of information to view.

| Option | Description |
|---|---|
| **Sensors** | Displays all sensors arranged in a tree view. If the status is blank, the health monitoring service cannot determine the status of the component. |
| | ■ Click the **Expand All** icon to expand the tree view to show all sensors under each group. |
| | ■ Click **Collapse All** icon to expand the tree view to show descriptive details for every sensor. |
| **Alerts and warnings** | Displays alerts and warnings. |
| **System event log** | Displays the system event log. |

# Reset Health Status Sensors

Some host hardware sensors display data that is cumulative over time. You can reset these sensors to clear the data in them and begin collecting new data.

If you need to preserve sensor data for troubleshooting or other purposes, take a screenshot, export the data, or download a support bundle before resetting sensors.

**Prerequisites**

Verify that the vCenter Hardware Status plug-in is enabled.

**Procedure**

**1**   Select a host in the vSphere Web Client navigator.

**2**   Click the **Monitor** tab, and click **Hardware Status**.

**3**   Click **Reset sensors**.

# Monitoring Events, Alarms, and Automated Actions

# 5

vSphere includes a user-configurable events and alarms subsystem. This subsystem tracks events happening throughout vSphere and stores the data in log files and the vCenter Server database. This subsystem also enables you to specify the conditions under which alarms are triggered. Alarms can change state from mild warnings to more serious alerts as system conditions change, and can trigger automated alarm actions. This functionality is useful when you want to be informed, or take immediate action, when certain events or conditions occur for a specific inventory object, or group of objects.

## Events

Events are records of user actions or system actions that occur on objects in vCenter Server or on a host. Actions that might be recordered as events include, but are not limited to, the following examples:

- A license key expires

- A virtual machine is powered on

- A user logs in to a virtual machine

- A host connection is lost

Event data includes details about the event such as who generated it, when it occured, and what type of event it is. There are three types of events:

- Information

- Warning

- Error

In the vSphere Web Client, event data is displayed in the **Monitor** tab. See View Events.

## Alarms

Alarms are notifications that are activated in response to an event, a set of conditions, or the state of an inventory object. An alarm definition consists of the following elements:

- Name and description - Provides an identifying label and description.

- Alarm type - Defines the type of object that will be monitored.

- Triggers - Defines the event, condition, or state that will trigger the alarm and defines the notification severity.

- Tolerance thresholds (Reporting) - Provides additional restrictions on condition and state triggers thresholds that must be exceeded before the alarm is triggered. Thresholds are not available in the vSphere Web Client.

- Actions - Defines operations that occur in response to triggered alarms. VMware provides sets of predefined actions that are specific to inventory object types.

Alarms have the following severity levels:

- Normal – green

- Warning – yellow

- Alert – red

Alarm definitions are associated with the object selected in the inventory. An alarm monitors the type of inventory objects specified in its definition.

For example, you might want to monitor the CPU usage of all virtual machines in a specific host cluster. You can select the cluster in the inventory, and add a virtual machine alarm to it. When enabled, that alarm will monitor all virtual machines running in the cluster and will trigger when any one of them meets the criteria defined in the alarm. If you want to monitor a specific virtual machine in the cluster, but not others, you would select that virtual machine in the inventory and add an alarm to it. One easy way to apply the same alarms to a group of objects is to place those objects in a folder and define the alarm on the folder.

**Note** You can enable, disable, and modify alarms only from the object in which the alarm is defined. For example, if you defined an alarm in a cluster to monitor virtual machines, you can only enable, disable, or modify that alarm through the cluster; you can not make changes to the alarm at the individual virtual machine level.

## Alarm Actions

Alarm actions are operations that occur in response to the trigger. For example, you can have an email notification sent to one or more administrators when an alarm is triggered.

**Note** Default alarms are not preconfigured with actions. You must manually set what action occurs when the triggering event, condition, or state occurs.

This chapter includes the following topics:

- View Events

- View System Logs

- Export Events Data

- View Triggered Alarms and Alarm Definitions

- Live Refresh of Recent Tasks and Alarms

- Set an Alarm

- Acknowledge Triggered Alarms

- Reset Triggered Event Alarms

- Preconfigured vSphere Alarms

# View Events

You can view events associated with a single object or view all vSphere events. The events list for a selected inventory object includes events associated with child objects. vSphere keeps information about tasks and events for 30 days.

**Procedure**

1 Select an inventory object in the vSphere Web Client.

2 Click the **Monitor** tab, and click **Events**.

3 Select an event to see details.

4 (Optional) Use the filter controls above the list to filter the list.

5 (Optional) Click a column heading to sort the list.

# View System Logs

vSphere records events in the vCenter Server database. System log entries include such information as who generated the event, when the event was created, and the type of event.

**Prerequisites**

- Required privilege: **Global.Diagnostics**

**Procedure**

1 Select a vCenter Server instance in the vSphere Web Client navigator.

2 Click **Monitor**, and click **System Logs**.

3 From the drop-down menu, select the log.

4 (Optional) Click **Show All Lines** or **Show Next 2000 Lines** to see additional log entries.

# Export Events Data

You can export all or part of the events data stored in the vCenter Server database.

**Prerequisites**

Required Role: **Read-only**

**Procedure**

**1** Select an inventory object in the vSphere Web Client.

**2** Click the **Monitor** tab, and click **Events**.

**3** Click the **Export** icon (⬚).

**4** In the **Export Events** window, specify what types of event information you want to export.

**5** Click **Generate CSV Report**, and click **Save**.

**6** Specify a file name and location and save the file.

# View Triggered Alarms and Alarm Definitions

Triggered alarms are visible in several locations throughout the vSphere Web Client.

**Procedure**

◆ To view all triggered alarms, click **All** in the Alarms sidebar panel.

> **Note** The list of alarms in the sidebar refreshes every 120 seconds. For information about changing the default refresh period, see the VMware knowledge base article at http://kb.vmware.com/kb/2020290.

◆ To view only newly triggered alarms, click **New** in the Alarms sidebar panel.

The sidebar panel displays the latest 30 most critical alarms.

◆ To view acknowledged alarms, click **Acknowledged** in the Alarms sidebar panel.

◆ To view alarms triggered on a selected inventory object, click the **Monitor** tab, click **Issues**, and click **Triggered Alarms**.

◆ To view a list of available alarm definitions for a selected inventory object, click the **Manage** tab, and click **Alarm Definitions**.

# Live Refresh of Recent Tasks and Alarms

You can configure the vSphere Web Client to live refresh the recent tasks and the alarms that result from operations that other users perform in your environment.

By design the vSphere Web Client displays tasks initiated by other users and the resulting alarms from these tasks only when you manually refresh the vSphere Web Client. If you want to see the tasks from other users, or monitor alarms resulting from other users actions, perform the following procedure.

**Procedure**

**1**  On the computer where the vSphere Web Client is installed, locate the
`webclient.properties` file.

The location of this file depends on the operating system on which the vSphere Web Client is installed.

| Operating System | File path |
|---|---|
| **Windows** | `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\webclient.properties` |
| **vCenter Server Appliance** | `/etc/vmware/vsphere-client/webclient.properties` |

**2**  Open the `webclient.properties` file, add the following configuration line, and save it.

```
live.updates.enabled=true
```

Live refresh of recent tasks and alarms is enabled for the vSphere Web Client.

**3**  Log out from the vSphere Web Client.

**4**  Use `https://hostname:9443/vsphere-client/` to log in to the vSphere Web Client.

*hostname* stands for the name or the IP address of the host where vCenter Server system runs.

If you log in to the vSphere Web Client by using the `https://hostname/vsphere-client/` you will see no recent tasks or alarms under the respective Recent Tasks or Alarms portlets in the vSphere Web Client.

**Results**

In an environment with multiple vCenter Server systems that are connected to the same vCenter Server Single-Sign On domain, the vSphere Web Client that you configured for live refresh displays recent tasks and alarms for all the vCenter Server instances in the domain. However, if you log in to a different vSphere Web Client, you will not see live refresh for recent tasks or alarms for any of the vCenter Server systems in the vCenter Server Single-Sign On domain.

**Example**

In this example, you have two vCenter Server instances (A and B) connected to the same vCenter Server Single-Sign On domain. With each of the vCenter Server instances you installed a vSphere Web Client instance.

You log in to vSphere Web Client A by using `https://hostnameA/vsphere-client/`.

You log in to vSphere Web Client B by using `https://hostnameB/vsphere-client/`.

You enable live refresh of recent tasks and alarms on vSphere Web Client A, and log out from it.

You can observe the following results:

- You log in to vSphere Web Client A from `https://hostnameA/vsphere-client/`. You do not see any recent tasks or alarms in the respective Recent Tasks or Alarms portlets.

- You log in to vSphere Web Client A from `https://hostnameA:9443/vsphere-client/`. You can see live refresh of recent tasks and alarms for all the users currently performing operations on both vCenter Server systems in the vCenter Server Single-Sign On domain.

- You log in to vSphere Web Client B from `https://hostnameB/vsphere-client/`. You can see the recent tasks and alarms of only operations that you perform on vCenter Server system A or vCenter Server system B. Only after you manually refresh the vSphere Web Client B you see the latest recent tasks and alarms that result from operations performed by other users on vCenter Server system A and vCenter Server system B.

# Set an Alarm

You can monitor inventory objects by setting alarms on them. Setting an alarm involves selecting the type of inventory object to monitor, defining when and for how long the alarm will trigger, and defining actions that will be performed as a result of the alarm being triggered. You define alarms in the alarm definition wizard under the **Manage** tab.

**Prerequisites**

Verify that you have logged in to the vSphere Web Client.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

- View and Edit Alarm Settings

    To monitor your environment, you can create and modify alarm definitions in the vSphere Web Client. You can view alarm settings from any object, but you can modify settings only through the object on which the alarm is defined.

- Specify Alarm Name, Description, and Type

    General settings of an alarm definition include alarm name, description, and type. You can also enable and disable the alarm from the general settings page. When you select the alarm type, you also select the type of inventory object and the type of activity (events, or conditions and states) that you want to monitor. The options on the Triggers page change depending on the type of activity that you choose to monitor.

- Specify How an Event-Based Alarm is Triggered

    You can specify the events, states, or conditions that trigger the alarm on the Triggers page of the alarm definition wizard. The alarm type options that you selected determine the options available on the Triggers page. An alarm definition must contain at least one trigger before you can save it.

- Specify How a Condition-Based or State-Based Alarm is Triggered

   You can select and configure the events, states, or conditions that trigger the alarm on the Triggers page of the alarm definition wizard. The options that you choose on the General page of the alarm definition wizard determine the options available on the Triggers page. An alarm definition must contain at least one trigger before you can save it.

- Specify Alarm Actions

   You can define actions that the system performs when the alarm is triggered or changes status. You can enable and disable alarms and alarm actions independently of each other.

## View and Edit Alarm Settings

To monitor your environment, you can create and modify alarm definitions in the vSphere Web Client. You can view alarm settings from any object, but you can modify settings only through the object on which the alarm is defined.

You can access alarm definitions in the **Manage** tab or in the pop-up menu.

### Prerequisites

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

### Procedure

- Create or edit alarms in the **Manage** tab.

   a   Select an inventory object and click the **Manage** tab.

   b   Click **Alarm Definitions**.

   c   Right-click the list of alarms and select one of the options to add or edit an alarm.

- Add an alarm to an object in the object navigator.

   a   Right-click an inventory object and select **Alarms > New Alarm Definition**.

## Specify Alarm Name, Description, and Type

General settings of an alarm definition include alarm name, description, and type. You can also enable and disable the alarm from the general settings page. When you select the alarm type, you also select the type of inventory object and the type of activity (events, or conditions and states) that you want to monitor. The options on the Triggers page change depending on the type of activity that you choose to monitor.

### Prerequisites

Verify that you have navigated to the General page of the alarm definition wizard. See View and Edit Alarm Settings.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

**Procedure**

**1** Type a name and description.

**2** Select the type of inventory object that this alarm will monitor.

**3** Select the type of activity that this alarm will monitor.

The options on the **Triggers** page change depending on the type of activity you select.

**4** Click **Next**.

**What to do next**

Specify alarm triggers.

## Specify How an Event-Based Alarm is Triggered

You can specify the events, states, or conditions that trigger the alarm on the Triggers page of the alarm definition wizard. The alarm type options that you selected determine the options available on the Triggers page. An alarm definition must contain at least one trigger before you can save it.

**Prerequisites**

Verify that you have navigated to the Triggers page of the alarm definition wizard. See View and Edit Alarm Settings.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

**Procedure**

**1** Select the trigger that you want to change, or click the **Add** icon to add a new trigger.

**2** Click in the **Event** column, and select an option from the drop-down menu.

**3** Click in the **Status** column, and select an option from the drop-down menu.

**4** (Optional) Configure additional conditions to be met before the alarm triggers.

a Click the **Add** icon to add an argument.

b Click in the **Argument** column, and select an option from the drop-down menu.

c Click in the **Operator** column, and select an option from the drop-down menu.

d Click in the **Value** column, and enter a value into the text field.

You can add more than one argument.

**5** Click **Next**.

**Results**

You selected and configured alarm triggers.

**What to do next**

Configure alarm actions.

## Specify How a Condition-Based or State-Based Alarm is Triggered

You can select and configure the events, states, or conditions that trigger the alarm on the Triggers page of the alarm definition wizard. The options that you choose on the General page of the alarm definition wizard determine the options available on the Triggers page. An alarm definition must contain at least one trigger before you can save it.

You can add multiple triggers and choose whether to trigger the alarm when one or all of them become active.

For information about defining triggers for an event-based alarm, see Specify How an Event-Based Alarm is Triggered.

**Prerequisites**

Verify that you have navigated to the Triggers page of the alarm definition wizard. See View and Edit Alarm Settings.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

**Procedure**

**1** Select the trigger that you want to change, or click the **Add** icon to add a new trigger.

**2** Click in the **Trigger** column, and select an option from the drop-down menu.

**3** Click in the **Operator** column, and select an option from the drop-down menu.

**4** Click in the **Warning Condition** column, and select an option from the drop-down menu to set the threshold for triggering a warning.

**5** Click in the **Critical Condition** column, and select an option from the drop-down menu.

**6** Click **Next**.

**Results**

You selected and configured alarm triggers.

**What to do next**

Configure alarm actions.

## Specify Alarm Actions

You can define actions that the system performs when the alarm is triggered or changes status. You can enable and disable alarms and alarm actions independently of each other.

vCenter Server can perform alarm actions in response to triggered alarms.

Prerequisites

Verify that you have navigated to the Actions page of the alarm definition wizard. See View and Edit Alarm Settings.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

■ Send Email as an Alarm Action

You can use the SMTP agent included with vCenter Server to send email notifications when alarms are triggered.

■ Send SNMP Traps as an Alarm

The SNMP agent included with vCenter Server can be used to send traps when alarms are triggered on a vCenter Server instance. The default hardware health alarms send SNMP traps by default.

■ Run a Script or a Command as an Alarm Action

You can configure an alarm to run a script or a command in the vSphere Web Client when the alarm is triggered.

Procedure

1   Select the action that you want to change, or click the **Add** icon to add one.

2   Click in the **Action** column, and select an option from the drop-down menu.

3   Click in the **Configuration** column, and enter configuration information for actions that require additional information:

| Option | Action |
| --- | --- |
| **Send a notification email** | Type email addresses, separated by a comma. |
| **Migrate VM** | Complete the virtual machine migration wizard. |
| **Run a command** | Take one of the following actions and press **Enter**:<br><br>■ If the command is a `.exe` file, enter the full path name of the command and include any parameters. For example, to run the `cmd.exe` command in the `C:\tools` directory, with the `alarmName` and `targetName` parameters, type: `c:\tools\cmd.exe alarmName targetName`<br><br>■ If the command is a `.bat` file, enter the full path name of the command as an argument to the `c:\windows\system32\cmd.exe` command. Include any parameters. For example, to run the `cmd.bat` command in the `C:\tools` directory, with the `alarmName` and `targetName` parameters, type:<br>`c:\windows\system32\cmd.exe /c c:\tools\cmd.bat alarmName targetName`<br><br>For `.bat` files, the command and its parameters must be formatted into one string. |

**4** (Optional) For each alarm status change column, select whether the alarm should be triggered when the alarm status changes.

Some actions do not support re-triggering when alarm status change.

**5** For repeat actions, select the time interval for the repetition.

**6** Click **Finish**.

Results

You configured the alarm general settings, triggers, and actions. The alarm monitors the object on which it is defined, as well as child objects.

## Send Email as an Alarm Action

You can use the SMTP agent included with vCenter Server to send email notifications when alarms are triggered.

Prerequisites

Ensure that the vCenter Server SMTP agent is properly configured to send email notifications.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

Procedure

**1** On the Actions page of the alarm definition wizard, click **Add** to add an action.

**2** In the **Actions** column, select **Send a notification email** from the drop-down menu.

**3** In the **Configuration** column, enter recipient addresses. Use commas to separate multiple addresses.

**4** (Optional) Configure alarm transitions and frequency.

### vCenter Server Email Agent Notifications

The following tables describe the information that is included in Alarm-based and Event-based email notifications. The first table described the information included in all email notifications; the second table describes additional information that is included in Event-based notifications.

Table 5-1. Basic SNMP Email Notification Details

| Email Entry | Description |
| --- | --- |
| Target | Object for which the alarm was triggered. |
| Old Status | Previous alarm status. Applies only to state triggers. |
| New Status | Current alarm status. Applies only to state triggers. |
| Metric Value | Threshold value that triggered the alarm. Applies only to metric condition triggers. |
| Alarm Definition | Alarm definition in vCenter Server, including the alarm name and status. |
| Description | Localized string containing a summary of the alarm. For example: |
| | `Alarm New_Alarm on host1.vmware.com changed from Gray to Red.` |

## Table 5-2. Additional Notification Details for Alarms Triggered by Events

| Detail | Description |
|--------|-------------|
| Event Details | VMODL event type name. |
| Summary | Alarm summary, including the event type, alarm name, and target object. |
| Date | Time and date the alarm was triggered. |
| UserName | Person who initiated the action that caused the event to be created. Events caused by an internal system activity do not have a UserName value. |
| Host | Host on which the alarm was triggered. |
| Resource Pool | Resource pool on which the alarm was triggered. |
| Datacenter | Data center on which the alarm was triggered. |
| Arguments | Arguments passed with the alarm and their values. |

## Send SNMP Traps as an Alarm

The SNMP agent included with vCenter Server can be used to send traps when alarms are triggered on a vCenter Server instance. The default hardware health alarms send SNMP traps by default.

### Prerequisites

Ensure that vCenter Server SNMP agents and ESXi SNMP agents are properly configured.

Ensure that SNMP trap receiver agents are properly configured.

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

### Procedure

1  On the Actions page of the alarm definition wizard, click **Add**.

2  In the **Actions** column, select **Send a notification trap** from the drop-down menu.

3  (Optional) Configure alarm transitions and frequency.

4  Click **Finish** to save the alarm settings.

### SNMP Trap Notifications

The following table describes the information that is included in vCenter Server and ESXi trap notifications.

## Table 5-3. SNMP Trap Notification Details

| Trap Entry | Description |
|------------|-------------|
| Type | The state vCenter Server is monitoring for the alarm. Options include Host Processor (or CPU) usage, Host Memory usage, Host State, Virtual Machine Processor (or CPU) usage, Virtual Machine Memory usage, Virtual Machine State, Virtual Machine Heartbeat. |
| Name | The name of the host or virtual machine that triggers the alarm. |
| Old Status | The alarm status before the alarm was triggered. |

Table 5-3. SNMP Trap Notification Details (continued)

| Trap Entry | Description |
| --- | --- |
| New Status | The alarm status when the alarm is triggered. |
| Object Value | The object value when the alarm is triggered. |

## Run a Script or a Command as an Alarm Action

You can configure an alarm to run a script or a command in the vSphere Web Client when the alarm is triggered.

Use the alarm environment variables to define complex scripts and attach them to multiple alarms or inventory objects. For example, you can write a script that enters the following trouble ticket information into an external system when an alarm is triggered:

- Alarm name

- Object on which the alarm was triggered

- Event that triggered the alarm

- Alarm trigger values

When you write the script, include the following environment variables in the script:

- VMWARE_ALARM_NAME

- VMWARE_ALARM_TARGET_NAME

- VMWARE_ALARM_EVENTDESCRIPTION

- VMWARE_ALARM_ALARMVALUE

You can attach the script to any alarm on any object without changing the script.

The script runs on the vCenter Server machine, and it runs even if you close the vSphere Web Client.

**Prerequisites**

Required Privilege: **Alarms.Create alarm** or **Alarms.Modify alarm**

**Procedure**

1    On the Actions tab of the alarm definitions wizard, click **Add** to add an action.

2    In the **Actions** column, select **Run a command** from the drop-down menu.

**3**   In the **Configuration** column, type script or command information:

| For this type of command... | Enter this... |
|---|---|
| **EXE executable files** | Full pathname of the command. For example, to run the `cmd.exe` command in the `C:\tools` directory, type:<br><br>`c:\tools\cmd.exe`. |
| **BAT batch file** | Full pathname of the command as an argument to the `c:\windows\system32\cmd.exe` command. For example, to run the `cmd.bat` command in the `C:\tools` directory, type:<br><br>`c:\windows\system32\cmd.exe /c c:\tools\cmd.bat`.<br><br>**Note**   The command and its parameters must be formatted into one string. |

If your script does not make use of the alarm environment variables, include any necessary parameters in the configuration field. Enclose parameters in curly brackets. For example:

```
c:\tools\cmd.exe {alarmName} {targetName}
c:\windows\system32\cmd.exe /c c:\tools\cmd.bat {alarmName} {targetName}
```

The script can run on any platform. You must provide the path to the script and argument keys. For example:

```
/var/myscripts/myAlarmActionScript {alarmName} {targetName}
```

**4**   (Optional) Configure alarm transitions and frequency.

**5**   Click **Finish** to save the alarm settings.

## Alarm Environment Variables for Scripts

To simplify script configuration for alarm actions, VMware provides environment variables for VMware alarms. Use the variables to define more complex scripts and attach them to multiple alarms or inventory objects so that the alarm action occurs when the alarm triggers.

Table 5-4. Alarm Environment Variables

| Variable Name | Variable Description | Supported Alarm Type |
|---|---|---|
| VMWARE_ALARM_NAME | The name of the triggered alarm. | Condition, State, Event |
| VMWARE_ALARM_ID | The MOID of the triggered alarm. | Condition, State, Event |
| VMWARE_ALARM_TARGET_NAME | The name of the entity on which the alarm triggered. | Condition, State, Event |
| VMWARE_ALARM_TARGET_ID | The MOID of the entity on which the alarm triggered. | Condition, State, Event |
| VMWARE_ALARM_OLDSTATUS | The old status of the alarm. | Condition, State, Event |
| VMWARE_ALARM_NEWSTATUS | The new status of the alarm. | Condition, State, Event |
| VMWARE_ALARM_TRIGGERINGSUMMARY | A multiline summary of the alarm. | Condition, State, Event |

## Table 5-4. Alarm Environment Variables (continued)

| Variable Name | Variable Description | Supported Alarm Type |
|---|---|---|
| VMWARE_ALARM_DECLARINGSUMMARY | A single-line declaration of the alarm expression. | Condition, State, Event |
| VMWARE_ALARM_ALARMVALUE | The value that triggered the alarm. | Condition, State |
| VMWARE_ALARM_EVENTDESCRIPTION | A description of the alarm status change event. | Condition, State |
| VMWARE_ALARM_EVENTDESCRIPTION | A description of the event that triggered the alarm. | Event |
| VMWARE_ALARM_EVENT_USERNAME | The user name associated with the event. | Event |
| VMWARE_ALARM_EVENT_DATACENTER | The name of the data center in which the event occurred. | Event |
| VMWARE_ALARM_EVENT_COMPUTERESOURCE | The name of the cluster or resource pool in which the event occurred. | Event |
| VMWARE_ALARM_EVENT_HOST | The name of the host on which the event occurred. | Event |
| VMWARE_ALARM_EVENT_VM | The name of the virtual machine on which the event occurred. | Event |
| VMWARE_ALARM_EVENT_NETWORK | The name of the network on which the event occurred. | Event |
| VMWARE_ALARM_EVENT_DATASTORE | The name of the datastore on which the event occurred. | Event |
| VMWARE_ALARM_EVENT_DVS | The name of the vSphere Distributed Switch on which the event occurred. | Event |

## Alarm Command-Line Parameters

VMware provides command-line parameters that function as a substitute for the default alarm environment variables. You can use these parameters when running a script as an alarm action for a condition, state, or event alarm.

The command-line parameters enable you to pass alarm information without having to change an alarm script. For example, you can use these parameters when you have an external program for which you do not have the source. You can pass in the necessary data by using the substitution parameters, which take precedence over the environment variables. You pass the parameters through the **Configuration** dialog box in the alarm definition wizard or on a command line.

Table 5-5. Command-Line Parameters for Alarm Action Scripts

| Variable | Description |
| --- | --- |
| {eventDescription} | The text of the alarmStatusChange event. The {eventDescription} variable is supported only for Condition and State alarms. |
| {targetName} | The name of the entity on which the alarm is triggered. |
| {alarmName} | The name of the alarm that is triggered. |
| {triggeringSummary} | A summary of the alarm trigger values. |
| {declaringSummary} | A summary of the alarm declaration values. |
| {oldStatus} | The alarm status before the alarm is triggered. |
| {newStatus} | The alarm status after the alarm is triggered. |
| {target} | The inventory object on which the alarm is set. |

# Acknowledge Triggered Alarms

After you acknowledge an alarm in the vSphere Web Client, its alarm actions are discontinued. Alarms are neither cleared, nor reset when acknowledged.

Acknowledging an alarm lets other users know that you are taking ownership of the issue. For example, a host has an alarm set on it that monitors CPU usage and that sends an email to an administrator when the alarm is triggered. The host CPU usage spikes, triggering the alarm which sends an email to the host's administrator. The administrator acknowledges the triggered alarm to let other administrators know he is working on the problem, and to prevent the alarm from sending more email messages. The alarm, however, is still visible in the system.

**Prerequisites**

Required privilege: **Alarm.Alarm Acknowledge**

**Procedure**

◆ Right-click the alarm in the Alarms sidebar panel and select **Acknowledge**.

◆ Acknowledge the alarm in the **Monitor** tab.

    a    Select an inventory object in the object navigator.

    b    Click the **Monitor** tab.

    c    Click **Issues**, and click **Triggered Alarms**.

    d    Right-click an alarm and select **Acknowledge**.

# Reset Triggered Event Alarms

An alarm triggered by an event might not reset to a normal state if vCenter Server does not retrieve the event that identifies the normal condition. In such cases, reset the alarm manually in the vSphere Web Client to return it to a normal state.

Prerequisites

Required privilege: **Alarm.Set Alarm Status**

Procedure

◆ Right-click an alarm in the Alarms sidebar pane and select **Reset to green**.

◆ Reset triggered alarms in the **Monitor** tab.

    a   Select an inventory object.

    b   Click the **Monitor** tab.

    c   Click **Issues**, and click **Triggered Alarms**.

    d   Select the alarms you want to reset.

        Use Shift+left-click or Ctrl+left-click to select multiple alarms.

    e   Right-click an alarm and select **Reset to Green**.

# Preconfigured vSphere Alarms

vCenter Server provides a list of default alarms, which monitor the operations of vSphere inventory objects. You only need to set up actions for these alarms.

Some alarms are stateless. vCenter Server does not keep data on stateless alarms and neither computes nor displays their status. Stateless alarms cannot be acknowledged or reset. Stateless alarms are indicated by an asterisk next to their name.

Table 5-6. Default vSphere Alarms

| Alarm Name | Description |
| --- | --- |
| Host connection and power state | Monitors the power state of the host and whether the host is reachable. |
| Host CPU usage | Monitors host CPU usage. |
| Host memory usage | Monitors host memory usage. |
| Virtual machine CPU usage | Monitors virtual machine CPU usage. |
| Virtual machine memory usage | Monitors virtual machine memory usage. |
| Datastore usage on disk | Monitors datastore disk usage. |
| | **Note** This alarm controls the Status value for datastores in vSphere Web Client. If you disable this alarm, the datastore status is displayed as Unknown. |
| Virtual machine CPU ready | Monitors virtual machine CPU ready time. |
| Virtual machine total disk latency | Monitors virtual machine total disk latency. |
| Virtual machine disk commands canceled | Monitors the number of virtual machine disk commands that are canceled. |
| Virtual machine disk reset | Monitors the number of virtual machine bus resets. |

## Table 5-6. Default vSphere Alarms (continued)

| Alarm Name | Description |
| --- | --- |
| License inventory monitoring | Monitors the license inventory for compliance. |
| License user threshold monitoring | Monitors whether a user-defined license threshold is exceeded. |
| License capacity monitoring | Monitors whether a license capacity is exceeded. |
| The host license edition is not compatible with the vCenter Server license edition | Monitors the compatibility of the vCenter Server and host licence editions. |
| Host flash capacity exceeds the licensed limit for Virtual SAN | Monitors whether the flash disk capacity on the host exceeds the limit of the Virtual SAN license. |
| Expired Virtual SAN license | Monitors the expiry of the Virtual SAN license and the end of the evaluation period. |
| Errors occurred on the disk(s) of a Virtual SAN host | Default alarm that monitors whether the host disks in the Virtual SAN cluster have errors. |
| Timed out starting Secondary VM * | Monitors whether starting a secondary virtual machine has timed out. . |
| No compatible host for Secondary VM | Monitors the availability of compatible hosts on which a secondary virtual machine can be created and run. |
| Virtual machine Fault Tolerance state changed | Monitors changes in the Fault Tolerance state of a virtual machine. |
| Virtual Machine Fault Tolerance vLockStep interval Status Changed | Monitors changes in the Fault Tolerance Secondary vLockStep interval. |
| Host processor status | Monitors the host processors. |
| Host memory status | Monitors host memory usage. |
| Host hardware fan status | Monitors host fans. |
| Host hardware voltage | Monitors host hardware voltage. |
| Host hardware temperature status | Monitors the temperature status of the host hardware. |
| Host hardware power status | Monitors the host power status. |
| Host hardware system board status | Monitors the status of host system boards. |
| Host battery status | Monitors the battery status of hosts. |
| Status of other host hardware objects | Monitors other host hardware objects. |
| Host storage status | Monitors host connectivity to storage devices. |
| Host IPMI System Event Log status | Monitors the capacity of the IPMI system event log. |
| Host Baseboard Management Controller status | Monitors the status of the Baseboard Management Controller. |
| Host error * | Monitors host error and warning events. |
| Virtual machine error * | Monitors virtual machine error and warning events. |
| Host connection failure * | Monitors host connection failures. |
| Unmanaged workload detected on SIOC-enabled datastore | Monitors the unmanaged I/O workload on a SIOC-enabled datastore. |

## Table 5-6. Default vSphere Alarms (continued)

| Alarm Name | Description |
|---|---|
| Thin-provisioned volume capacity threshold exceeded | Monitors whether the thin provisioning threshold on the storage array exceeds for volumes backing the datastore. |
| Datastore capability alarm | Monitors the change in the capability status for volumes backing the datastore. |
| VASA provider disconnected | Monitors the changes in the connection state of VASA providers. |
| VASA Provider certificate expiration alarm | Monitors whether VASA provider certificates are getting close to their expiry date. |
| VM storage compliance alarm | Monitors the virtual disk compliance with the object-based storage. |
| Datastore compliance alarm | Monitors whether the virtual disk on the datastore complies with the object-based storage. |
| Refreshing CA certificates and CRLs for a VASA provider failed | Monitors whether the refreshing of CA certificates and CRLs for some of the VASA providers has failed. |
| Insufficient vSphere HA failover resources | Monitors the sufficiency of failover cluster resources required for vSphere High Availability. |
| vSphere HA failover in progress | Monitors the failover progress of vSphere High Availability. |
| Cannot find vSphere HA primary agent | Monitors whether vCenter Server is able to connect to a vSphere High Availability primary agent. |
| vSphere HA host status | Monitors the host health status reported by vSphere High Availability. |
| vSphere HA virtual machine failover failed | Monitors whether a failover operation that uses vSphere High Availability failed. |
| vSphere HA virtual machine monitoring action | Monitors whether vSphere High Availability has restarted a virtual machine. |
| vSphere HA virtual machine monitoring error | Monitors whether vSphere High Availability failed to reset a virtual machine. |
| vSphere HA VM Component Protection could not power off a virtual machine | Monitors whether vSphere High Availability VM Component Protection cannot power off a virtual machine with an inaccessible datastore. |
| License error * | Monitors license errors. |
| Health status changed * | Monitors changes to service and extension health status. |
| Virtual machine component protection restart error | Monitors whether the vSphere HA VM Component Protection fails to restart a virtual machine. |
| Storage DRS recommendation | Monitors Storage DRS recommendations. |
| Storage DRS is not supported on a host | Monitors and alerts when Storage DRS is not supported on a host. |
| Datastore cluster is out of space | Monitors whether a datastore cluster runs out of disk space. |
| Datastore is in multiple datacenters | Monitors whether a datastore in a datastore cluster is visible in more than one data center. |

## Table 5-6. Default vSphere Alarms (continued)

| Alarm Name | Description |
| --- | --- |
| vSphere Distributed Switch VLAN trunked status | Monitors changes in vSphere Distributed Switch VLAN trunked status. |
| vSphere Distributed Switch MTU matched status | Monitors changes in vSphere Distributed Switch MTU matched status. |
| vSphere Distributed Switch MTU supported status | Monitors changes in vSphere Distributed Switch MTU supported status. |
| vSphere Distributed Switch teaming matched status | Monitors changes in vSphere Distributed Switch teaming matched status. |
| Virtual Machine network adapter reservation status | Monitors changes in the reservation status of a virtual machine network adapter. |
| Virtual machine Consolidation Needed status | Monitors changes in the virtual machine Consolidation Needed status. |
| Host virtual flash resource status | Monitors the Flash Read Cache resource status on the host. |
| Host virtual flash resource usage | Monitors the Flash Read Cache resource usage on the host. |
| Registration/unregistration of a VASA vendor provider on a Virtual SAN host fails | Default alarm that monitors whether the registration or unregistration of a VASA vendor provider on a Virtual SAN host fails. |
| Registration/unregistration of third-party IO filter storage providers fails on a host | Default alarm that monitors whether vCenter Server fails to register or unregister third-party IO filter storage providers on a host. |
| Service Control Agent Health Alarm | Monitors the health status of the VMware Service Control Agent. |
| Identity Health Alarm | Monitors the health status of the Identity Management Service. |
| vSphere Client Health Alarm | Monitors the health status of the vSphere Web Client. |
| ESX Agent Manager Health Alarm | Monitors the health status of the ESX Agent Manager. |
| Message Bus Config Health Alarm | Monitors the health status of the Message Bus Configuration Service. |
| Cis License Health Alarm | Monitors the health status of the License Service. |
| Appliance Management Health Alarm | Monitors the health status of the Appliance Management Service. |
| Inventory Health Alarm | Monitors the health status of the Inventory Service. |
| vCenter Server Health Alarm | Monitors the health status of vCenter Server. |
| Data Service Health Alarm | Monitors the health status of the Data Service. |
| RBD Health Alarm | Monitors the health status of the vSphere Auto Deploy Waiter. |
| vService Manager Health Alarm | Monitors the health status of the vService Manager. |
| Performance Charts Service Health Alarm | Monitors the health status of the Performance Charts Service. |

## Table 5-6. Default vSphere Alarms (continued)

| Alarm Name | Description |
|---|---|
| Content Library Service Health Alarm | Monitors the health status of the VMware Content Library Service. |
| Transfer Service Health Alarm | Monitors the health status of the VMware Transfer Service. |
| VMware vSphere ESXi Dump Collector Health Alarm | Monitors the health status of the VMware vSphere ESXi Dump Collector Service. |
| VMware vAPI Endpoint Service Health Alarm | Monitors the health status of the VMware vAPI Endpoint Service. |
| VMware System and Hardware Health Manager Service Health Alarm | Monitors the health status of the VMware System and Hardware Health Manager Service. |
| VMware vSphere Profile-Driven Storage Service Health Alarm | Monitors the health status of the VMware vSphere Profile-Driven Storage Service. |
| VMware Common Logging Service Health Alarm | Monitors the health status of the VMware Common Logging Service. |
| VMware vFabric Postgres Service Health Alarm | Monitors the health status of the VMware vFabric Postgres Service. |
| ESXi Host Certificates Update Failure Status | Monitors whether the update of the ESXi host certificates failed. |
| ESXi Host Certificate Status | Monitors the certificate status of an ESXi host. |
| ESXi Host Certificate Verification Failure Status | Monitors whether the verification of an ESXi host certificate failed. |
| vSphere vCenter Host Certificate Management Mode | Monitors changes in the certificate management mode of vCenter Server. |
| Root Certificate Status | Monitors whether a root certificate is getting close to its expiration date. |
| GPU ECC Uncorrected Memory Alarm | Monitors the GPU ECC uncorrected memory status. |
| GPU ECC Corrected Memory Alarm | Monitors the GPU ECC corrected memory status. |
| GPU Thermal Condition Alarm | Monitors the GPU Thermal condition status. |
| Network connectivity lost | Monitors the network connectivity on a virtual switch. |
| Network uplink redundancy lost | Monitors network uplink redundancy on a virtual switch. |
| Network uplink redundancy degraded * | Monitors network uplink redundancy degradation on a virtual switch. |
| VMKernel NIC not configured correctly * | Monitors incorrectly configured VMkernel NICs. |
| Cannot connect to storage * | Monitors host connectivity to a storage device. |
| Migration error * | Monitors whether a virtual machine cannot be migrated or relocated, or is orphaned. |
| Exit standby error | Monitors whether a host cannot exit standby mode. |

**Table 5-7. Deprecated vSphere Alarms**

| Alarm name | Description |
| --- | --- |
| Cannot connect to network | Monitors the network connectivity on a virtual switch. |
| IPv6 TSO not supported | Monitors whether the IPv6 TSO packets sent by the guest operating system of a virtual machine are dropped. |
| SRM Consistency Group Violation | Datastore cluster has datastores that belong to different SRM consistency groups. |
| Virtual machine high availability error | Monitors High Availability errors on a virtual machine. |
| Cluster high availability error * | Monitors High Availability errors on a cluster. |
| Health status monitoring | Monitors changes in the overall health status of vCenter Server components. |
| Pre-4.1 host connected to SIOC-enabled datastore | Monitors whether a host running ESX/ESXi 4.1 or earlier is connected to a SIOC-enabled datastore. |
| Host service console swap rates | Monitors host service console memory swap rates. |

# Monitoring Solutions with the vCenter Solutions Manager

6

In the vSphere Web Client, you can view an inventory of installed solutions, view detailed information about the solutions, and monitor the solution health status. A solution is an extension of vCenter Server that adds new functions to a vCenter Server instance.

VMware products that integrate with vCenter Server are also considered solutions. For example, vSphere ESX Agent Manager is a solution provided by VMware to let you manage host agents that add new capabilities to ESX/ESXi hosts.

You can install a solution to add functionality from third-party technologies to the standard functions of vCenter Server. Solutions typically are delivered as OVF packages. You can install and deploy solutions from the vSphere Web Client. You can integrate solutions into the vCenter Solutions Manager, which provides a view in the vSphere Web Client that lists all solutions.

If a virtual machine or vApp is running a solution, a custom icon represents it in the inventory of the vSphere Web Client. Each solution registers a unique icon to identify that the virtual machine or vApp is being managed by that solution. The icons show the power states (powered on, paused, or powered off). The solutions display more than one type of icon if they manage more than one type of virtual machine or vApp.

When you power on or power off a virtual machine or vApp, you are notified that you are performing this operation on an entity that is managed by the Solutions Manager. When you attempt an operation on a virtual machine or a vApp that is managed by a solution, an informational warning message appears.

For more information, see the *Developing and Deploying vSphere Solutions, vServices, and ESX Agents* documentation.

This chapter includes the following topics:

- View Solutions and vServices
- Monitoring Agents
- Monitoring vServices

# View Solutions and vServices

In the vSphere Web Client, you can view information about solutions and vService providers. A vService is a service that a solution provides to specific applications that run inside virtual machines and vApps.

**Procedure**

1 Navigate to the vCenter Server system in the object navigator.

2 Double-click the vCenter Server object.

3 Click **Extensions**.

4 Select a solution.

   The **Summary** tab displays more information about the solution.

5 To view vService provider information, click **Monitor**, and click **vServices**.

# Monitoring Agents

The vCenter Solutions Manager displays the vSphere ESX Agent Manager agents that you use to deploy and manage related agents on ESX/ESXi hosts.

You can use the Solutions Manager to keep track of whether the agents of a solution are working as expected. Outstanding issues are reflected by the solution's ESX Agent Manager status and a list of issues.

When the status of a solution changes, the Solutions Manager updates the ESX Agent Manager summary status and state. Administrators use this status to track whether the goal state is reached.

The agent health status is indicated by a specific color.

Table 6-1. ESX Agent Manager health status

| Status | Description |
| --- | --- |
| Red | The solution must intervene for the ESX Agent Manager to proceed. For example, if a virtual machine agent is powered off manually on a compute resource and the ESX Agent Manager does not attempt to power on the agent. The ESX Agent Manager reports this action to the solution, and the solution alerts the administrator to power on the agent. |
| Yellow | The ESX Agent Manager is actively working to reach a goal state. The goal state can be enabled, disabled, or uninstalled. For example, when a solution is registered, its status is yellow until the ESX Agent Manager deploys the solutions agents to all the specified compute resources. A solution does not need to intervene when the ESX Agent Manager reports its ESX Agent Manager health status as yellow. |
| Green | A solution and all its agents have reached the goal state. |

# Monitoring vServices

A vService is a service or function that a solution provides to virtual machines and vApps. A solution can provide one or more vServices. These vServices integrate with the platform and are able to change the environment in which the vApp or virtual machine runs.

A vService is a type of service for a virtual machine and a vApp provided by a vCenter extension. Virtual machines and vApps can have dependencies on vServices. Each dependency is associated with a vService type. The vService type must be bound to a particular vCenter extension that implements that vService type. This vService type is similar to a virtual hardware device. For example, a virtual machine can have a networking device that at deployment must be connected to a particular network.

The vService Manager allows a solution to connect to operations related to OVF templates:

- Importing OVF templates. Receive a callback when OVF templates with a vService dependancy of a certain type is imported.

- Exporting OVF templates. Inserts OVF sections when a virtual machine is exported.

- OVF environment generation. Inserts OVF sections into the OVF environment at the power-on instance.

The **vServices** tab in the Solution Manager provides details for each vCenter extension. This information allows you to monitor vService providers and list the virtual machines or vApps to which they are bound.

# Monitoring the Health of Services and Nodes

7

You can monitor the health status of services and nodes to determine whether problems exist in your environment.

The vSphere Web Client provides an overview of all services and nodes across the management stack of the vCenter Server system. A list of default services is available for each vCenter Server instance.

This chapter includes the following topics:

- View the Health Status of Services and Nodes

## View the Health Status of Services and Nodes

In the vSphere Web Client, you can view the health status of vCenter Server services and nodes.

vCenter Server instances and machines that run vCenter Server services are considered nodes. Graphical badges represent the health status of services and nodes.

### Prerequisites

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

### Procedure

1   Log in as administrator@*your_domain_name* to the vCenter Server instance by using the vSphere Web Client.

2   On the vSphere Web Client Home page, click **System Configuration**.

    You can view the health status badges for the services and nodes.

    Table 7-1. Health States

| Badge Icon | Description |
| --- | --- |
|  | Good. The health of the object is normal. |
|  | Warning. The object is experiencing some problems. |

Table 7-1. Health States (continued)

| Badge Icon | Description |
|---|---|
|  | Critical. The object is either not functioning properly or will stop functioning soon. |
|  | Unknown. No data is available for this object. |

**3**  (Optional) In the Services Health and Nodes Health panes, click the hyperlink next to the health badge to view all services and nodes in this health state.

For example, in the Services Health pane, click the hyperlink of the Warning health status, and in the dialog box that pops up, select a service to view more information about the service and attempt to resolve the health issues of the service.

# Performance Monitoring Utilities: resxtop and esxtop

# 8

The `resxtop` and `esxtop` command-line utilities provide a detailed look at how ESXi uses resources in real time. You can start either utility in one of three modes: interactive (default), batch, or replay.

The fundamental difference between `resxtop` and `esxtop` is that you can use `resxtop` remotely, whereas you can start `esxtop` only through the ESXi Shell of a local ESXi host.

This chapter includes the following topics:

- Using the esxtop Utility
- Using the resxtop Utility
- Using esxtop or resxtop in Interactive Mode
- Using Batch Mode
- Using Replay Mode

## Using the esxtop Utility

You can run the `esxtop` utility using the ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

Type the command, using the options you want:

```
esxtop [-h] [-v] [-b] [-s] [-a] [-c config file] [-R vm-support_dir_path]  [-d delay] [-n iterations]
```

The `esxtop` utility reads its default configuration from `.esxtop50rc` on the ESXi system. This configuration file consists of nine lines.

The first eight lines contain lowercase and uppercase letters to specify which fields appear in which order on the CPU, memory, storage adapter, storage device, virtual machine storage, network, interrupt, and CPU power panels. The letters correspond to the letters in the Fields or Order panels for the respective `esxtop` panel.

The ninth line contains information on the other options. Most important, if you saved a configuration in secure mode, you do not get an insecure esxtop without removing the s from the seventh line of your .esxtop50rc file. A number specifies the delay time between updates. As in interactive mode, typing c, m, d, u, v, n, I, or p determines the panel with which esxtop starts.

**Note** Do not edit the .esxtop50rc file. Instead, select the fields and the order in a running esxtop process, make changes, and save this file using the W interactive command.

## Using the resxtop Utility

The resxtop utility is a vSphere CLI command.

Before you can use any vSphere CLI commands, you must either download and install a vSphere CLI package or deploy the vSphere Management Assistant (vMA) to your ESXi host or vCenter Server system.

After it is set up, start resxtop from the command line. For remote connections, you can connect to a host either directly or through vCenter Server.

The command-line options listed in the following table are the same as for esxtop (except for the R option) with additional connection options.

**Note** resxtop does not use all the options shared by other vSphere CLI commands.

Table 8-1. resxtop Command-Line Options

| Option | Description |
| --- | --- |
| [server] | Name of the remote host to connect to (required). If connecting directly to the ESXi host, use the name of that host. If your connection to the ESXi host is indirect (that is, through vCenter Server), use the name of the vCenter Server system for this option. |
| [vihost] | If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the host, this option is not used. Note that the host name needs to be the same as what appears in the vSphere Web Client. |
| [portnumber] | Port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. |
| [username] | User name to be authenticated when connecting to the remote host. The remote server prompts you for a password. |

You can also use resxtop on a local ESXi host by omitting the server option on the command line. The command defaults to localhost.

## Using esxtop or resxtop in Interactive Mode

By default, resxtop and esxtop run in interactive mode. Interactive mode displays statistics in different panels.

A help menu is available for each panel.

## Interactive Mode Command-Line Options

You can use various command-line options with `esxtop` and `resxtop` in interactive mode.

**Table 8-2. Interactive Mode Command-Line Options**

| Option | Description |
| --- | --- |
| h | Prints help for `resxtop` (or `esxtop`) command-line options. |
| v | Prints `resxtop` (or `esxtop`) version number. |
| s | Calls `resxtop` (or `esxtop`) in secure mode. In secure mode, the –d command, which specifies delay between updates, is disabled. |
| d | Specifies the delay between updates. The default is five seconds. The minimum is two seconds. Change this with the interactive command s. If you specify a delay of less than two seconds, the delay is set to two seconds. |
| n | Number of iterations. Updates the display n times and exits. Default value is 10000. |
| server | The name of the remote server host to connect to (required for `resxtop` only). |
| vihost | If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the ESXi host, this option is not used. Note that the host name needs to be the same as what is displayed in the vSphere Web Client. |
| portnumber | The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. (`resxtop` only) |
| username | The user name to be authenticated when connecting to the remote host. The remote server prompts you for a password, as well (`resxtop` only). |
| a | Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default ~/.esxtop50rc configuration file or a user-defined configuration file. |
| c filename | Load a user-defined configuration file. If the -c option is not used, the default configuration filename is ~/.esxtop50rc. Create your own configuration file, specifying a different filename, using the W single-key interactive command. |

## Common Statistics Description

Several statistics appear on the different panels while `resxtop` (or `esxtop`) is running in interactive mode. These statistics are common across all four panels.

The Uptime line, found at the top of each of the four `resxtop` (or `esxtop`) panels, displays the current time, time since last reboot, number of currently running worlds and load averages. A world is an ESXi VMkernel schedulable entity, similar to a process or thread in other operating systems.

Below that the load averages over the past one, five, and fifteen minutes appear. Load averages take into account both running and ready-to-run worlds. A load average of 1.00 means that there is full utilization of all physical CPUs. A load average of 2.00 means that the ESXi system might need twice as many physical CPUs as are currently available. Similarly, a load average of 0.50 means that the physical CPUs on the ESXi system are half utilized.

## Statistics Columns and Order Pages

You can define the order of fields displayed in interactive mode.

If you press f, F, o, or 0, the system displays a page that specifies the field order on the top line and short descriptions of the field contents. If the letter in the field string corresponding to a field is uppercase, the field is displayed. An asterisk in front of the field description indicates whether a field is displayed.

The order of the fields corresponds to the order of the letters in the string.

From the Field Select panel, you can:

- Toggle the display of a field by pressing the corresponding letter.

- Move a field to the left by pressing the corresponding uppercase letter.

- Move a field to the right by pressing the corresponding lowercase letter.

## Interactive Mode Single-Key Commands

When running in interactive mode, resxtop (or esxtop) recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table. The command to specify the delay between updates is disabled if the s option is given on the command line. All sorting interactive commands sort in descending order.

Table 8-3. Interactive Mode Single-Key Commands

| Key | Description |
| --- | --- |
| h or ? | Displays a help menu for the current panel, giving a brief summary of commands, and the status of secure mode. |
| space | Immediately updates the current panel. |
| ^L | Erases and redraws the current panel. |
| f or F | Displays a panel for adding or removing statistics columns (fields) to or from the current panel. |
| o or O | Displays a panel for changing the order of statistics columns on the current panel. |
| # | Prompts you for the number of statistics rows to display. Any value greater than 0 overrides automatic determination of the number of rows to show, which is based on window size measurement. If you change this number in one resxtop (or esxtop) panel, the change affects all four panels. |
| s | Prompts you for the delay between updates, in seconds. Fractional values are recognized down to microseconds. The default value is five seconds. The minimum value is two seconds. This command is not available in secure mode. |
| W | Write the current setup to an esxtop (or resxtop) configuration file. This is the recommended way to write a configuration file. The default filename is the one specified by -c option, or ~/.esxtop50rc if the -c option is not used. You can also specify a different filename on the prompt generated by this W command. |

Table 8-3. Interactive Mode Single-Key Commands (continued)

| Key | Description |
| --- | --- |
| q | Quit interactive mode. |
| c | Switch to the CPU resource utilization panel. |
| p | Switch to the CPU Power utilization panel. |
| m | Switch to the memory resource utilization panel. |
| d | Switch to the storage (disk) adapter resource utilization panel. |
| u | Switch to storage (disk) device resource utilization screen. |
| v | Switch to storage (disk) virtual machine resource utilization screen. |
| n | Switch to the network resource utilization panel. |
| i | Switch to the interrupt panel. |

## CPU Panel

The CPU panel displays server-wide statistics as well as statistics for individual world, resource pool, and virtual machine CPU utilization.

Resource pools, virtual machines that are running, or other worlds are at times called groups. For worlds belonging to a virtual machine, statistics for the virtual machine that is running are displayed. All other worlds are logically aggregated into the resource pools that contain them.

Table 8-4. CPU Panel Statistics

| Line | Description |
| --- | --- |
| PCPU USED(%) | A PCPU refers to a physical hardware execution context. This can be a physical CPU core if hyperthreading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) if hyperthreading is enabled.<br><br>PCPU USED(%) displays the following percentages:<br><br>■ percentage of CPU usage per PCPU<br><br>■ percentage of CPU usage averaged over all PCPUs<br><br>CPU Usage (%USED) is the percentage of PCPU nominal frequency that was used since the last screen update. It equals the total sum of %USED for Worlds that ran on this PCPU.<br><br>**Note** If a PCPU is running at frequency that is higher than its nominal (rated) frequency, then PCPU USED(%) can be greater than 100%.<br><br>If a PCPU and its partner are busy when hyperthreading is enabled, each PCPU accounts for half of the CPU usage. |
| PCPU UTIL(%) | A PCPU refers to a physical hardware execution context. This can be a physical CPU core if hyperthreading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) if hyperthreading is enabled.<br><br>PCPU UTIL(%) represents the percentage of real time that the PCPU was not idle (raw PCPU utilization) and it displays the percentage CPU utilization per PCPU, and the percentage CPU utilization averaged over all PCPUs.<br><br>**Note** PCPU UTIL(%) might differ from PCPU USED(%) due to power management technologies or hyperthreading. |

## Table 8-4. CPU Panel Statistics (continued)

| Line | Description |
|---|---|
| ID | Resource pool ID or virtual machine ID of the resource pool or virtual machine of the world that is running, or world ID of the world that is running. |
| GID | Resource pool ID of the resource pool or virtual machine of the world that is running. |
| NAME | Name of the resource pool or virtual machine of the world that is running, or name of the world that is running. |
| NWLD | Number of members in the resource pool or virtual machine of the world that is running. If a Group is expanded using the interactive command e, then NWLD for all the resulting worlds is 1. |
| %STATE TIMES | Set of CPU statistics made up of the following percentages. For a world, the percentages are a percentage of one physical CPU core. |
| %USED | Percentage of physical CPU core cycles used by the resource pool, virtual machine, or world. %USED might depend on the frequency with which the CPU core is running. When running with lower CPU core frequency, %USED can be smaller than %RUN. On CPUs which support turbo mode, CPU frequency can also be higher than the nominal (rated) frequency, and %USED can be larger than %RUN. <br> %USED = %RUN + %SYS - %OVRLP |
| %SYS | Percentage of time spent in the ESXi VMkernel on behalf of the resource pool, virtual machine, or world to process interrupts and to perform other system activities. This time is part of the time used to calculate %USED. <br> %USED = %RUN + %SYS - %OVRLP |
| %WAIT | Percentage of time the resource pool, virtual machine, or world spent in the blocked or busy wait state. This percentage includes the percentage of time the resource pool, virtual machine, or world was idle. <br> 100% = %RUN + %RDY + %CSTP + %WAIT |
| %VMWAIT | The total percentage of time the Resource Pool/World spent in a blocked state waiting for events. |
| %IDLE | Percentage of time the resource pool, virtual machine, or world was idle. Subtract this percentage from %WAIT to see the percentage of time the resource pool, virtual machine, or world was waiting for some event. The difference, %WAIT- %IDLE, of the VCPU worlds can be used to estimate guest I/O wait time. To find the VCPU worlds, use the single-key command e to expand a virtual machine and search for the world NAME starting with "vcpu". (The VCPU worlds might wait for other events in addition to I/O events, so this measurement is only an estimate.) |
| %RDY | Percentage of time the resource pool, virtual machine, or world was ready to run, but was not provided CPU resources on which to execute. <br> 100% = %RUN + %RDY + %CSTP + %WAIT |
| %MLMTD (max limited) | Percentage of time the ESXi VMkernel deliberately did not run the resource pool, virtual machine, or world because doing so would violate the resource pool, virtual machine, or world's limit setting. Because the resource pool, virtual machine, or world is ready to run when it is prevented from running in this way, the %MLMTD (max limited) time is included in %RDY time. |
| %SWPWT | Percentage of time a resource pool or world spends waiting for the ESXi VMkernel to swap memory. The %SWPWT (swap wait) time is included in the %WAIT time. |
| EVENT COUNTS/s | Set of CPU statistics made up of per second event rates. These statistics are for VMware internal use only. |
| CPU ALLOC | Set of CPU statistics made up of the following CPU allocation configuration parameters. |
| AMIN | Resource pool, virtual machine, or world attribute Reservation. |

## Table 8-4. CPU Panel Statistics (continued)

| Line | Description |
| --- | --- |
| AMAX | Resource pool, virtual machine, or world attribute Limit. A value of -1 means unlimited. |
| ASHRS | Resource pool, virtual machine, or world attribute Shares. |
| SUMMARY STATS | Set of CPU statistics made up of the following CPU configuration parameters and statistics. These statistics apply only to worlds and not to virtual machines or resource pools. |
| AFFINITY BIT MASK | Bit mask showing the current scheduling affinity for the world. |
| HTSHARING | Current hyperthreading configuration. |
| CPU | The physical or logical processor on which the world was running when resxtop (or esxtop) obtained this information. |
| HTQ | Indicates whether the world is currently quarantined or not. N means no and Y means yes. |
| TIMER/s | Timer rate for this world. |
| %OVRLP | Percentage of system time spent during scheduling of a resource pool, virtual machine, or world on behalf of a different resource pool, virtual machine, or world while the resource pool, virtual machine, or world was scheduled. This time is not included in %SYS. For example, if virtual machine A is currently being scheduled and a network packet for virtual machine B is processed by the ESXi VMkernel, the time spent appears as %OVRLP for virtual machine A and %SYS for virtual machine B.<br><br>%USED = %RUN + %SYS - %OVRLP |
| %RUN | Percentage of total time scheduled. This time does not account for hyperthreading and system time. On a hyperthreading enabled server, the %RUN can be twice as large as %USED.<br><br>%USED = %RUN + %SYS - %OVRLP<br><br>100% = %RUN + %RDY + %CSTP + %WAIT |
| %CSTP | Percentage of time a resource pool spends in a ready, co-deschedule state.<br><br>**Note** You might see this statistic displayed, but it is intended for VMware use only.<br><br>100% = %RUN + %RDY + %CSTP + %WAIT |
| POWER | Current CPU power consumption for a resource pool (in Watts). |
| %LAT_C | Percentage of time the resource pool or world was ready to run but was not scheduled to run because of CPU resource contention. |
| %LAT_M | Percentage of time the resource pool or world was ready to run but was not scheduled to run because of memory resource contention. |
| %DMD | CPU demand in percentage. It represents the average active CPU load in the past minute. |
| CORE UTIL(%) | Percentage of CPU cycles per core when at least one of the PCPUs in this core is unhalted, and its average over all cores.<br><br>This statistic only appears when hyperthreading is enabled.<br><br>In batch mode, the correspondening CORE UTIL(%) statistic is displayed for each PCPU. For example, PCPU 0 and PCPU 1 have the same the CORE UTIL(%) number, and that is the number for core 0. |

You can change the display using single-key commands.

Table 8-5. CPU Panel Single-Key Commands

| Command | Description |
|---------|-------------|
| e | Toggles whether CPU statistics are displayed expanded or unexpanded. |
| | The expanded display includes CPU resource utilization statistics broken down by individual worlds belonging to a resource pool or virtual machine. All percentages for the individual worlds are percentage of a single physical CPU. |
| | Consider these examples: |
| | ■ If the %Used by a resource pool is 30% on a two-way server, the resource pool is utilizing 30 percent of one physical core. |
| | ■ If the %Used by a world belonging to a resource pool is 30 percent on a two-way server, that world is utilizing 30% of one physical core. |
| U | Sorts resource pools, virtual machines, and worlds by the resource pool's or virtual machine's %Used column. This is the default sort order. |
| R | Sorts resource pools, virtual machines, and worlds by the resource pool's or virtual machine's %RDY column. |
| N | Sorts resource pools, virtual machines, and worlds by the GID column. |
| V | Displays virtual machine instances only. |
| L | Changes the displayed length of the NAME column. |

# CPU Power Panel

The CPU Power panel displays CPU Power utilization statistics.

On the CPU Power panel, statistics are arranged per PCPU. A PCPU is a physical hardware execution context -- a physical CPU core if hyper-threading is unavailable or disabled, or a logical CPU (LCPU or SMT thread) if hyper-threading is enabled.

Table 8-6. CPU Power Panel Statistics

| Line | Description |
|------|-------------|
| Power Usage | Current total power usage (in Watts). |
| Power Cap | Total power cap (in Watts). |
| PSTATE MHZ | Clock frequency per state. |
| %USED | Percentage of PCPU nominal frequency used since the last screen update. It is the same as PCPU USED(%) shown in the CPU Screen. |
| %UTIL | Raw PCPU utilization, that is the percentage of real time that PCPU was not idle. It is the same as PCPU UTIL(%) shown in the CPU Screen. |
| %Cx | Percentage of time the PCPU spent in C-State 'x'. |
| %Px | Percentage of time the PCPU spent in P-State 'x'. On systems with Processor Clocking Control, P-states are not directly visible to ESXi, so `esxtop` instead shows the percentage of time spent at full speed under the heading 'P0' and the percentage of time spent at any lower speed under 'P1'. |
| %Tx | Percentage of time the PCPU spent in T-State 'x'. |

# Memory Panel

The Memory panel displays server-wide and group memory utilization statistics. As on the CPU panel, groups correspond to resource pools, running virtual machines, or other worlds that are consuming memory.

The first line, found at the top of the Memory panel displays the current time, time since last reboot, number of currently running worlds, and memory overcommitment averages. The memory overcommitment averages over the past one, five, and fifteen minutes appear. Memory overcommitment of 1.00 means a memory overcommitment of 100 percent.

**Table 8-7. Memory Panel Statistics**

| Field | Description |
| --- | --- |
| PMEM (MB) | Displays the machine memory statistics for the server. All numbers are in megabytes. |
| | **total** |
| | Total amount of machine memory in the server. |
| | **vmk** |
| | Amount of machine memory being used by the ESXi VMkernel. |
| | **other** |
| | Amount of machine memory being used by everything other than the ESXi VMkernel. |
| | **free** |
| | Amount of machine memory that is free. |
| VMKMEM (MB) | Displays the machine memory statistics for the ESXi VMkernel. All numbers are in megabytes. |
| | **managed** |
| | Total amount of machine memory managed by the ESXi VMkernel. |
| | **min free** |
| | Minimum amount of machine memory that the ESXi VMkernel aims to keep free. |
| | **rsvd** |
| | Total amount of machine memory currently reserved by resource pools. |
| | **ursvd** |
| | Total amount of machine memory currently unreserved. |
| | **state** |
| | Current machine memory availability state. Possible values are high, soft, hard and low. High means that the machine memory is not under any pressure and low means that it is. |

## Table 8-7. Memory Panel Statistics (continued)

| Field | Description |
|---|---|
| NUMA (MB) | Displays the ESXi NUMA statistics. This line appears only if the ESXi host is running on a NUMA server. All numbers are in megabytes.<br><br>For each NUMA node in the server, two statistics are displayed:<br>■ The total amount of machine memory in the NUMA node that is managed by ESXi.<br>■ The amount of machine memory in the node that is currently free (in parentheses).<br><br>Shared memory for the ESXi host might be larger than the total amount of memory if memory is over-commited. |
| PSHARE (MB) | Displays the ESXi page-sharing statistics. All numbers are in megabytes.<br><br>**shared**<br>Amount of physical memory that is being shared.<br><br>**common**<br>Amount of machine memory that is common across worlds.<br><br>**saving**<br>Amount of machine memory that is saved because of page sharing.<br><br>shared = common + saving |
| SWAP (MB) | Displays the ESXi swap usage statistics. All numbers are in megabytes.<br><br>**curr**<br>Current swap usage.<br><br>**rclmtgt**<br>Where the ESXi system expects the reclaimed memory to be. Memory can be reclaimed by swapping or compression.<br><br>**r/s**<br>Rate at which memory is swapped in by the ESXi system from disk.<br><br>**w/s**<br>Rate at which memory is swapped to disk by the ESXi system. |
| ZIP (MB) | Displays the ESXi memory compression statistics. All numbers are in megabytes.<br><br>**zipped**<br>Total compressed physical memory.<br><br>**saved**<br>Saved memory by compression. |

## Table 8-7. Memory Panel Statistics (continued)

| Field | Description |
|---|---|
| MEMCTL (MB) | Displays the memory balloon statistics. All numbers are in megabytes. |
| | **curr** |
| | Total amount of physical memory reclaimed using the `vmmemctl` module. |
| | **target** |
| | Total amount of physical memory the ESXi host attempts to reclaim using the `vmmemctl` module. |
| | **max** |
| | Maximum amount of physical memory the ESXi host can reclaim using the `vmmemctl` module. |
| AMIN | Memory reservation for this resource pool or virtual machine. |
| AMAX | Memory limit for this resource pool or virtual machine. A value of -1 means Unlimited. |
| ASHRS | Memory shares for this resource pool or virtual machine. |
| NHN | Current home node for the resource pool or virtual machine. This statistic is applicable only on NUMA systems. If the virtual machine has no home node, a dash (-) appears. |
| NRMEM (MB) | Current amount of remote memory allocated to the virtual machine or resource pool. This statistic is applicable only on NUMA systems. |
| N% L | Current percentage of memory allocated to the virtual machine or resource pool that is local. |
| MEMSZ (MB) | Amount of physical memory allocated to a resource pool or virtual machine. The values are the same for the VMM and VMX groups.<br>MEMSZ = GRANT + MCTLSZ + SWCUR + "never touched" |
| GRANT (MB) | Amount of guest physical memory mapped to a resource pool or virtual machine. The consumed host machine memory is equal to GRANT - SHRDSVD. The values are the same for the VMM and VMX groups. |
| CNSM | Amount of the memory currently consumed by the virtual machine. The memory currently consumed by the virtual machine is equal to the amount of memory that the VM guest operating system currently uses, excluding the amount of memory saved for sharing if memory sharing is enabled on the VM, excluding the amount of memory saved if some of the VM memory is compressed. For more information on memory sharing and memory compression, see the *vSphere Resource Management* documentation. |
| SZTGT (MB) | Amount of machine memory the ESXi VMkernel wants to allocate to a resource pool or virtual machine. The values are the same for the VMM and VMX groups. |
| TCHD (MB) | Working set estimate for the resource pool or virtual machine. The values are the same for the VMM and VMX groups. |
| %ACTV | Percentage of guest physical memory that is being referenced by the guest. This is an instantaneous value. |
| %ACTVS | Percentage of guest physical memory that is being referenced by the guest. This is a slow moving average. |
| %ACTVF | Percentage of guest physical memory that is being referenced by the guest. This is a fast moving average. |

## Table 8-7. Memory Panel Statistics (continued)

| Field | Description |
|---|---|
| %ACTVN | Percentage of guest physical memory that is being referenced by the guest. This is an estimation. (You might see this statistic displayed, but it is intended for VMware use only.) |
| MCTL? | Memory balloon driver is installed or not. **N** means no, **Y** means yes. |
| MCTLSZ (MB) | Amount of physical memory reclaimed from the resource pool by way of ballooning. |
| MCTLTGT (MB) | Amount of physical memory the ESXi system attempts to reclaim from the resource pool or virtual machine by way of ballooning. |
| MCTLMAX (MB) | Maximum amount of physical memory the ESXi system can reclaim from the resource pool or virtual machine by way of ballooning. This maximum depends on the guest operating system type. |
| SWCUR (MB) | Current swap usage by this resource pool or virtual machine. |
| SWTGT (MB) | Target where the ESXi host expects the swap usage by the resource pool or virtual machine to be. |
| SWR/s (MB) | Rate at which the ESXi host swaps in memory from disk for the resource pool or virtual machine. |
| SWW/s (MB) | Rate at which the ESXi host swaps resource pool or virtual machine memory to disk. |
| LLSWR/s (MB) | Rate at which memory is read from the host cache. The reads and writes are attributed to the VMM group only, LLSWAP statistics are not displayed for the VM group. |
| LLSWW/s (MB) | Rate at which memory is written to the host cache from various sources. The reads and writes are attributed to the VMM group only, LLSWAP statistics are not displayed for the VM group. |
| CPTRD (MB) | Amount of data read from checkpoint file. |
| CPTTGT (MB) | Size of checkpoint file. |
| ZERO (MB) | Resource pool or virtual machine physical pages that are zeroed. |
| SHRD (MB) | Resource pool or virtual machine physical pages that are shared. |
| SHRDSVD (MB) | Machine pages that are saved because of resource pool or virtual machine shared pages. |
| OVHD (MB) | Current space overhead for resource pool. |
| OVHDMAX (MB) | Maximum space overhead that might be incurred by resource pool or virtual machine. |
| OVHDUW (MB) | Current space overhead for a user world. (You might see this statistic displayed, but it is intended for VMware use only.) |
| GST_NDx (MB) | Guest memory allocated for a resource pool on NUMA node x. This statistic is applicable on NUMA systems only. |
| OVD_NDx (MB) | VMM overhead memory allocated for a resource pool on NUMA node x. This statistic is applicable on NUMA systems only. |
| TCHD_W (MB) | Write working set estimate for resource pool. |
| CACHESZ (MB) | Compression memory cache size. |
| CACHEUSD (MB) | Used compression memory cache. |
| ZIP/s (MB/s) | Compressed memory per second. |
| UNZIP/s (MB/s) | Decompressed memory per second. |

Table 8-8. Memory Panel Interactive Commands

| Command | Description |
| --- | --- |
| M | Sort resource pools or virtual machines by MEMSZ column. This is the default sort order. |
| B | Sort resource pools or virtual machines by Group Memctl column. |
| N | Sort resource pools or virtual machines by GID column. |
| V | Display virtual machine instances only. |
| L | Changes the displayed length of the NAME column. |

## Storage Adapter Panel

Statistics in the Storage Adapter panel are aggregated per storage adapter by default. Statistics can also be viewed per storage path.

Table 8-9. Storage Adapter Panel Statistics

| Column | Description |
| --- | --- |
| ADAPTR | Name of the storage adapter. |
| PATH | Storage path name. This name is only visible if the corresponding adapter is expanded. See interactive command e in Table 8-10. Storage Adapter Panel Interactive Commands. |
| NPTH | Number of paths. |
| AQLEN | Current queue depth of the storage adapter. |
| CMDS/s | Number of commands issued per second. |
| READS/s | Number of read commands issued per second. |
| WRITES/s | Number of write commands issued per second. |
| MBREAD/s | Megabytes read per second. |
| MBWRTN/s | Megabytes written per second. |
| RESV/s | Number of SCSI reservations per second. |
| CONS/s | Number of SCSI reservation conflicts per second. |
| DAVG/cmd | Average device latency per command, in milliseconds. |
| KAVG/cmd | Average ESXi VMkernel latency per command, in milliseconds. |
| GAVG/cmd | Average virtual machine operating system latency per command, in milliseconds. |
| QAVG/cmd | Average queue latency per command, in milliseconds. |
| DAVG/rd | Average device read latency per read operation, in milliseconds. |
| KAVG/rd | Average ESXi VMkernel read latency per read operation, in milliseconds. |
| GAVG/rd | Average guest operating system read latency per read operation, in milliseconds. |
| QAVG/rd | Average queue latency per read operation, in milliseconds. |
| DAVG/wr | Average device write latency per write operation, in milliseconds. |
| KAVG/wr | Average ESXi VMkernel write latency per write operation, in milliseconds. |
| GAVG/wr | Average guest operating system write latency per write operation, in milliseconds. |

Table 8-9. Storage Adapter Panel Statistics (continued)

| Column | Description |
|---|---|
| QAVG/wr | Average queue latency per write operation, in milliseconds. |
| FCMDS/s | Number of failed commands issued per second. |
| FREAD/s | Number of failed read commands issued per second. |
| FWRITE/s | Number of failed write commands issued per second. |
| FMBRD/s | Megabytes of failed read operations per second. |
| FMBWR/s | Megabytes of failed write operations per second. |
| FRESV/s | Number of failed SCSI reservations per second. |
| ABRTS/s | Number of commands cancelled per second. |
| RESETS/s | Number of commands reset per second. |
| PAECMD/s | The number of PAE (Physical Address Extension) commands per second. |
| PAECP/s | The number of PAE copies per second. |
| SPLTCMD/s | The number of split commands per second. |
| SPLTCP/s | The number of split copies per second. |

The following table displays the interactive commands you can use with the storage adapter panel.

Table 8-10. Storage Adapter Panel Interactive Commands

| Command | Description |
|---|---|
| e | Toggles whether storage adapter statistics appear expanded or unexpanded. Allows you to view storage resource utilization statistics broken down by individual paths belonging to an expanded storage adapter. You are prompted for the adapter name. |
| r | Sorts by READS/s column. |
| w | Sorts by WRITES/s column. |
| R | Sorts by MBREAD/s read column. |
| T | Sorts by MBWRTN/s written column. |
| N | Sorts first by ADAPTR column, then by PATH column. This is the default sort order. |

# Storage Device Panel

The storage device panel displays server-wide storage utilization statistics.

By default, the information is grouped per storage device. You can also group the statistics per path, per world, or per partition.

## Table 8-11. Storage Device Panel Statistics

| Column | Description |
| --- | --- |
| DEVICE | Name of the storage device. |
| PATH | Path name. This name is visible only if the corresponding device is expanded to paths. See the interactive command p in Table 8-12. Storage Device Panel Interactive Commands. |
| WORLD | World ID. This ID is visible only if the corresponding device is expanded to worlds. See the interactive command e in Table 8-12. Storage Device Panel Interactive Commands. The world statistics are per world per device. |
| PARTITION | Partition ID. This ID is visible only if the corresponding device is expanded to partitions. See interactive command t in Table 8-12. Storage Device Panel Interactive Commands. |
| NPH | Number of paths. |
| NWD | Number of worlds. |
| NPN | Number of partitions. |
| SHARES | Number of shares. This statistic is applicable only to worlds. |
| BLKSZ | Block size in bytes. |
| NUMBLKS | Number of blocks of the device. |
| DQLEN | Current device queue depth of the storage device. |
| WQLEN | World queue depth. This is the maximum number of ESXi VMkernel active commands that the world is allowed to have. This is a per device maximum for the world. It is valid only if the corresponding device is expanded to worlds. |
| ACTV | Number of commands in the ESXi VMkernel that are currently active. This statistic applies to only worlds and devices. |
| QUED | Number of commands in the ESXi VMkernel that are currently queued. This statistic applies to only worlds and devices. |
| %USD | Percentage of the queue depth used by ESXi VMkernel active commands. This statistic applies to only worlds and devices. |
| LOAD | Ratio of ESXi VMkernel active commands plus ESXi VMkernel queued commands to queue depth. This statistic applies to only worlds and devices. |
| CMDS/s | Number of commands issued per second. |
| READS/s | Number of read commands issued per second. |
| WRITES/s | Number of write commands issued per second. |
| MBREAD/s | Megabytes read per second. |
| MBWRTN/s | Megabytes written per second. |
| DAVG/cmd | Average device latency per command in milliseconds. |
| KAVG/cmd | Average ESXi VMkernel latency per command in milliseconds. |
| GAVG/cmd | Average guest operating system latency per command in milliseconds. |
| QAVG/cmd | Average queue latency per command in milliseconds. |
| DAVG/rd | Average device read latency per read operation in milliseconds. |
| KAVG/rd | Average ESXi VMkernel read latency per read operation in milliseconds. |
| GAVG/rd | Average guest operating system read latency per read operation in milliseconds. |

Table 8-11. Storage Device Panel Statistics (continued)

| Column | Description |
| --- | --- |
| QAVG/rd | Average queue read latency per read operation in milliseconds. |
| DAVG/wr | Average device write latency per write operation in milliseconds. |
| KAVG/wr | Average ESXi VMkernel write latency per write operation in milliseconds. |
| GAVG/wr | Average guest operating system write latency per write operation in milliseconds. |
| QAVG/wr | Average queue write latency per write operation in milliseconds. |
| ABRTS/s | Number of commands cancelled per second. |
| RESETS/s | Number of commands reset per second. |
| PAECMD/s | Number of PAE commands per second. This statistic applies to only paths. |
| PAECP/s | Number of PAE copies per second. This statistic applies to only paths. |
| SPLTCMD/s | Number of split commands per second. This statistic applies to only paths. |
| SPLTCP/s | Number of split copies per second. This statistic applies to only paths. |

The following table displays the interactive commands you can use with the storage device panel.

Table 8-12. Storage Device Panel Interactive Commands

| Command | Description |
| --- | --- |
| e | Expand or roll up storage world statistics. This command allows you to view storage resource utilization statistics separated by individual worlds belonging to an expanded storage device. You are prompted for the device name. The statistics are per world per device. |
| P | Expand or roll up storage path statistics. This command allows you to view storage resource utilization statistics separated by individual paths belonging to an expanded storage device. You are prompted for the device name. |
| t | Expand or roll up storage partition statistics. This command allows you to view storage resource utilization statistics separated by individual partitions belonging to an expanded storage device. You are prompted for the device name. |
| r | Sort by READS/s column. |
| w | Sort by WRITES/s column. |
| R | Sort by MBREAD/s column. |
| T | Sort by MBWRTN column. |
| N | Sort first by DEVICE column, then by PATH, WORLD, and PARTITION column. This is the default sort order. |
| L | Changes the displayed length of the DEVICE column. |

# Virtual Machine Storage Panel

This panel displays virtual machine-centric storage statistics.

By default, statistics are aggregated on a per-resource-pool basis. One virtual machine has one corresponding resource pool, so the panel displays statistics on a per-virtual-machine basis. You can also view statistics on per-VSCSI-device basis.

Table 8-13. Virtual Machine Storage Panel Statistics

| Column | Description |
| --- | --- |
| ID | Resource pool ID or VSCSI ID of VSCSI device. |
| GID | Resource pool ID. |
| VMNAME | Name of the resource pool. |
| VSCSINAME | Name of the VSCSI device. |
| NDK | Number of VSCSI devices |
| CMDS/s | Number of commands issued per second. |
| READS/s | Number of read commands issued per second. |
| WRITES/s | Number of write commands issued per second. |
| MBREAD/s | Megabytes read per second. |
| MBWRTN/s | Megabytes written per second. |
| LAT/rd | Average latency (in milliseconds) per read. |
| LAT/wr | Average latency (in milliseconds) per write. |

The following table lists the interactive commands you can use with the virtual machine storage panel.

Table 8-14. Virtual Machine Storage Panel Interactive Commands

| Command | Description |
| --- | --- |
| e | Expand or roll up storage VSCSI statistics. Allows you to view storage resource utilization statistics broken down by individual VSCSI devices belonging to a group. You are prompted to enter the group ID. The statistics are per VSCSI device. |
| r | Sort by READS/s column. |
| w | Sort by WRITES/s column. |
| R | Sort by MBREAD/s column. |
| T | Sort by MBWRTN/s column. |
| N | Sort first by VMNAME column, and then by VSCSINAME column. This is the default sort order. |

## Network Panel

The Network panel displays server-wide network utilization statistics.

Statistics are arranged by port for each virtual network device configured. For physical network adapter statistics, see the row in the table that corresponds to the port to which the physical network adapter is connected. For statistics on a virtual network adapter configured in a particular virtual machine, see the row corresponding to the port to which the virtual network adapter is connected.

Table 8-15. Network Panel Statistics

| Column | Description |
|---|---|
| PORT-ID | Virtual network device port ID. |
| UPLINK | Y means the corresponding port is an uplink. N means it is not. |
| UP | Y means the corresponding link is up. N means it is not. |
| SPEED | Link speed in Megabits per second. |
| FDUPLX | Y means the corresponding link is operating at full duplex. N means it is not. |
| USED-BY | Virtual network device port user. |
| DTYP | Virtual network device type. H means HUB and S means switch. |
| DNAME | Virtual network device name. |
| PKTTX/s | Number of packets transmitted per second. |
| PKTRX/s | Number of packets received per second. |
| MbTX/s | MegaBits transmitted per second. |
| MbRX/s | MegaBits received per second. |
| %DRPTX | Percentage of transmit packets dropped. |
| %DRPRX | Percentage of receive packets dropped. |
| TEAM-PNIC | Name of the physical NIC used for the team uplink. |
| PKTTXMUL/s | Number of multicast packets transmitted per second. |
| PKTRXMUL/s | Number of multicast packets received per second. |
| PKTTXBRD/s | Number of broadcast packets transmitted per second. |
| PKTRXBRD/s | Number of broadcast packets received per second. |

The following table displays the interactive commands you can use with the network panel.

Table 8-16. Network Panel Interactive Commands

| Command | Description |
|---|---|
| T | Sorts by Mb Tx column. |
| R | Sorts by Mb Rx column. |
| t | Sorts by Packets Tx column. |
| r | Sorts by Packets Rx column. |
| N | Sorts by PORT-ID column. This is the default sort order. |
| L | Changes the displayed length of the DNAME column. |

## Interrupt Panel

The interrupt panel displays information about the use of interrupt vectors.

Table 8-17. Interrupt Panel Statistics

| Column | Description |
|--------|-------------|
| VECTOR | Interrupt vector ID. |
| COUNT/s | Total number of interrupts per second. This value is cumulative of the count for every CPU. |
| COUNT_x | Interrupts per second on CPU x. |
| TIME/int | Average processing time per interrupt (in microseconds). |
| TIME_x | Average processing time per interrupt on CPU x (in microseconds). |
| DEVICES | Devices that use the interrupt vector. If the interrupt vector is not enabled for the device, its name is enclosed in angle brackets (< and >). |

# Using Batch Mode

Batch mode allows you to collect and save resource utilization statistics in a file.

After you prepare for batch mode, you can use `esxtop` or `resxtop` in this mode.

## Prepare for Batch Mode

To run in batch mode, you must first prepare for batch mode.

**Procedure**

1 Run `resxtop` (or `esxtop`) in interactive mode.

2 In each of the panels, select the columns you want.

3 Save this configuration to a file (by default `~/.esxtop50rc`) using the W interactive command.

**Results**

You can now use `resxtop` (or `esxtop`) in batch mode.

## Use esxtop or resxtop in Batch Mode

After you have prepared for batch mode, you can use `esxtop` or `resxtop` in this mode.

**Procedure**

1 Start `resxtop` (or `esxtop`) to redirect the output to a file.

For example:

```
esxtop –b > my_file.csv
```

The filename must have a `.csv` extension. The utility does not enforce this, but the post-processing tools require it.

2 Process statistics collected in batch mode using tools such as Microsoft Excel and Perfmon.

**Results**

In batch mode, `resxtop` (or `esxtop`) does not accept interactive commands. In batch mode, the utility runs until it produces the number of iterations requested (see command-line option n, below, for more details), or until you end the process by pressing Ctrl+c.

## Batch Mode Command-Line Options

You can use batch mode with command-line options.

Table 8-18. Command-Line Options in Batch Mode

| Option | Description |
| --- | --- |
| *a* | Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default `~/.esxtop50rc` configuration file or a user-defined configuration file. |
| b | Runs `resxtop` (or `esxtop`) in batch mode. |
| c *filename* | Load a user-defined configuration file. If the –c option is not used, the default configuration filename is ~/.esxtop41rc. Create your own configuration file, specifying a different filename, using the W single-key interactive command. |
| d | Specifies the delay between statistics snapshots. The default is five seconds. The minimum is two seconds. If a delay of less than two seconds is specified, the delay is set to two seconds. |
| n | Number of iterations. `resxtop` (or `esxtop`) collects and saves statistics this number of times, and then exits. |
| server | The name of the remote server host to connect to (required, `resxtop` only). |
| vihost | If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the ESXi host, this option is not used. Note that the host name needs to be the same as what appears in the vSphere Web Client. |
| portnumber | The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. (`resxtop` only) |
| username | The user name to be authenticated when connecting to the remote host. You are prompted by the remote server for a password, as well (`resxtop` only). |

# Using Replay Mode

In replay mode, `esxtop` replays resource utilization statistics collected using `vm-support`.

After you prepare for replay mode, you can use esxtop in this mode. See the `vm-support` man page.

In replay mode, `esxtop` accepts the same set of interactive commands as in interactive mode and runs until no more snapshots are collected by `vm-support` to be read or until the requested number of iterations are completed.

## Prepare for Replay Mode

To run in replay mode, you must prepare for replay mode.

**Procedure**

**1** Run `vm-support` in snapshot mode in the ESXi Shell.

Use the following command.

```
vm-support -S -d duration -I interval
```

**2** Unzip and untar the resulting tar file so that `esxtop` can use it in replay mode.

**Results**

You can now use `esxtop` in replay mode.

## Use esxtop in Replay Mode

You can use `esxtop` in replay mode.

Replay mode can be run to produce output in the same style as batch mode (see the command-line option **b**, below).

**Note** Batch output from `esxtop` cannot be played back by `resxtop`.

Snapshots collected by vm-supported can be replayed by `esxtop`. However, vm-support output generated by ESXi can be replayed only by `esxtop` running on the same version of ESXi.

**Procedure**

◆ To activate replay mode, enter the following at the command-line prompt.

```
esxtop -R vm-support_dir_path
```

## Replay Mode Command-Line Options

You can use replay mode with command-line options.

The following table lists the command-line options available for `esxtop` replay mode.

Table 8-19. Command-Line Options in Replay Mode

| Option | Description |
|---|---|
| R | Path to the vm-support collected snapshot's directory. |
| a | Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default `~/.esxtop50rc` configuration file or a user-defined configuration file. |
| b | Runs `esxtop` in Batch mode. |
| c filename | Load a user-defined configuration file. If the -c option is not used, the default configuration filename is ~/.esxtop50rc. Create your own configuration file and specify a different filename using the W single-key interactive command. |

## Table 8-19. Command-Line Options in Replay Mode (continued)

| Option | Description |
| --- | --- |
| d | Specifies the delay between panel updates. The default is five seconds. The minimum is two seconds. If a delay of less than two seconds is specified, the delay is set to two seconds. |
| n | Number of iterations `esxtop` updates the display this number of times and then exits. |

# Using the vimtop Plug-In to Monitor the Resource Usage of Services

# 9

You can use the `vimtop` utility plug-in to monitor vSphere services that run in the vCenter Server Appliance.

`vimtop` is a tool similar to `esxtop`, which runs in the environment of the vCenter Server Appliance. By using the text-based interface of `vimtop` in the appliance shell, you can view overall information about the vCenter Server Appliance, and a list of vSphere services and their resource usage.

This chapter includes the following topics:

- Monitor Services by Using vimtop in Interactive Mode
- Interactive Mode Command-Line Options
- Interactive Mode Single-Key Commands for vimtop

## Monitor Services by Using vimtop in Interactive Mode

You can use the `vimtop` plug-in to monitor services in real time.

The default view of the `vimtop` interactive mode consists of the overview tables and the main table. You can use single-key commands in interactive mode to switch the view from processes to disks or network.

### Procedure

1   From an SSH client program, log in to the vCenter Server Appliance shell.

2   Run the `vimtop` command to access the plug-in in interactive mode.

## Interactive Mode Command-Line Options

You can use various command-line options when you run the `vimtop` command to enter the plug-in interactive mode.

Table 9-1. Interactive Mode Command-Line Options

| Option | Description |
| --- | --- |
| –h | Prints help for the `vimtop` command-line options. |
| –v | Prints the `vimtop` version number. |
| –c *filename* | Loads a user-defined `vimtop` configuration file. If the –c option is not used, the default configuration file is `/root/vimtop/vimtop.xml`.<br><br>You can create your own configuration file, specifying a different file name and path by using the `W` single-key interactive command. |
| –n *number* | Sets the number of performed iterations before the `vimtop` exits interactive mode. `vimtop` updates the display *number* number of times and exits. The default value is 10000. |
| –p / –d *seconds* | Sets the update period in seconds. |

# Interactive Mode Single-Key Commands for vimtop

When running in interactive mode, `vimtop` recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table.

Table 9-2. Interactive Mode Single-Key Commands

| Key Names | Description |
| --- | --- |
| h | Show a help menu for the current panel, giving a brief summary of commands, and the status of secure mode. |
| i | Show or hide the top line view of the overview panel of the `vimtop` plug-in. |
| t | Show or hide the Tasks section, which displays information in the overview panel about the tasks currently running on the vCenter Server instance . |
| m | Show or hide the Memory section in the overview panel. |
| f | Show or hide the CPU section which displays information in the overview panel about all available CPUs. |
| g | Show or hide the CPUs section which displays information in the overview panel about the top 4 physical CPUs. |
| spacebar | Immediately refreshes the current pane. |
| p | Pause the displayed information about the services resource usage in the current panels. |
| r | Refresh the displayed information about the services resource usage in the current panels. |
| s | Set refresh period. |
| q | Exit the interactive mode of the `vimtop` plug-in . |
| k | Displays the Disks view of the main panel. |
| o | Switch the main panel to Network view. |
| Esc | Clear selection or return to the Processes view of the main panel. |
| Enter | Select a service to view additional details. |
| n | Show or hide names of the headers in the main panel. |
| u | Show or hide the measurement units in the headers in the main panel. |

## Table 9-2. Interactive Mode Single-Key Commands (continued)

| Key Names | Description |
| --- | --- |
| left, right arrows | Select columns. |
| up, down arrows | Select rows. |
| <,> | Move a selected column. |
| Delete | Remove selected column. |
| c | Add a new column to the current view of the main panel. Use spacebar to add or remove columns from the displayed list. |
| a | Sort the selected column in ascending order. |
| d | Sort the selected column in descending order. |
| z | Clear the sort order for all columns. |
| l | Set width for the selected column. |
| x | Return the column widths to their default values. |
| + | Expand selected item. |
| - | Collapse selected item. |
| w | Write the current setup to a `vimtop` configuration file. The default file name is the one specified by `-c` option, or `/root/vimtop/vimtop.xml` if the `-c` option is not used. You can also specify a different file name on the prompt generated by the `w` command. |

# Monitoring Networked Devices with SNMP and vSphere

Simple Network Management Protocol (SNMP) is commonly used by management programs to monitor a variety of networked devices.

vSphere systems run SNMP agents, which can provide information to a management program in at least one of the following ways:

- In response to a GET, GETBULK, or GETNEXT operation, which is a specific request for information from the management system.

- By sending a notification which is an alert sent by the SNMP agent to notify the management system of a particular event or condition.

Management Information Base (MIB) files define the information that can be provided by managed devices. The MIB files define managed objects, described by object identifiers (OIDs) and variables arranged in a hierarchy.

vCenter Server and ESXi have SNMP agents. The agent provided with each product has different capabilities.

This chapter includes the following topics:

- Using SNMP Traps with vCenter Server
- Configure SNMP for ESXi
- SNMP Diagnostics
- Monitor Guest Operating Systems with SNMP
- VMware MIB Files
- SNMPv2 Diagnostic Counters

## Using SNMP Traps with vCenter Server

The SNMP agent included with vCenter Server can be used to send traps when vCenter Server starts and when an alarm is triggered on vCenter Server. The vCenter Server SNMP agent functions only as a trap emitter and does not support other SNMP operations, such as receiving GET, GETBULK, and GETNEXT requests.

vCenter Server can send SNMP v1 traps to other management programs. You must configure your management server to interpret the SNMP traps sent by vCenter Server.

To use the vCenter Server SNMP traps, configure the SNMP settings on vCenter Server and your management client software to accept the traps from vCenter Server.

The traps sent by vCenter Server are defined in `VMWARE-VC-EVENT-MIB.mib`.

## Configure SNMP Settings for vCenter Server

If you plan to use SNMP with vCenter Server, you must use the vSphere Web Client to configure the SNMP settings.

### Prerequisites

- Verify that the vSphere Web Client is connected to a vCenter Server instance.
- Verify that you have the domain name or IP address of the SNMP receiver, the port number of the receiver, and the community string.

### Procedure

1  In the vSphere Web Client, navigate to a vCenter Server instance.

2  Click the **Manage** tab, and click **Settings**.

3  Click **Edit**.

4  Click **SNMP receivers** to edit their settings.

5  Enter the following information for the primary receiver of the SNMP traps.

| Option | Description |
| --- | --- |
| **Primary Receiver URL** | Enter the domain name or IP address of the receiver of SNMP traps. |
| **Enable receiver** | Select the check box to enable the SNMP receiver. |
| **Receiver port** | Enter the port number of the receiver to which the SNMP agent sends traps. If the port value is empty, vCenter Server uses the default port, 162. |
| **Community string** | Enter the community string that is used for authentication. |

6  (Optional) Enter information about other SNMP receivers in the **Receiver 2 URL**, **Receiver 3 URL**, and **Receiver 4 URL** options, and select **Enabled**.

7  Click **OK**.

### Results

The vCenter Server system is now ready to send traps to the management system you have specified.

### What to do next

Configure your SNMP management software to receive and interpret data from the vCenter Server SNMP agent. See Configure SNMP Management Client Software.

# Configure SNMP for ESXi

ESXi includes an SNMP agent that can send notifications (traps and informs) and receive GET, GETBULK, and GETNEXT requests.

In ESXi 5.1 and later releases, the SNMP agent adds support for version 3 of the SNMP protocol, offering increased security and improved functionality, including the ability to send informs. You can use `esxcli` commands to enable and configure the SNMP agent. You configure the agent differently depending on whether you want to use SNMP v1/v2c or SNMP v3.

As an alternative to configuring SNMP manually using `esxcli` commands, you can use host profiles to configure SNMP for an ESXi host. See the *vSphere Host Profiles* documentation for more information.

**Note**  For information on configuring SNMP for ESXi 5.0 or earlier or ESX 4.1 or earlier, see the documentation for the appropriate product version.

- Configure the SNMP Agent for Polling

   If you configure the ESXi SNMP agent for polling, it can listen for and respond to requests from SNMP management client systems, such as GET, GETNEXT and GETBULK requests.

- Configure ESXi for SNMP v1 and v2c

   When you configure the ESXi SNMP agent for SNMP v1 and v2c, the agent supports sending notifications and receiving GET requests.

- Configure ESXi for SNMP v3

   When you configure the ESXi SNMP agent for SNMP v3, the agent supports sending informs as well as traps. SNMP v3 also provides stronger security than v1 or v2c, including key authentication and encryption.

- Configure the Source of Hardware Events Received by the SNMP Agent

   You can configure the ESXi SNMP agent to receive hardware events either from IPMI sensors or CIM indications.

- Configure the SNMP Agent to Filter Notifications

   You can configure the ESXi SNMP agent to filter out notifications if you don't want your SNMP management software to receive those notifications.

- Configure SNMP Management Client Software

   After you have configured a vCenter Server instance or an ESXi host to send traps, you must configure your management client software to receive and interpret those traps.

## Configure the SNMP Agent for Polling

If you configure the ESXi SNMP agent for polling, it can listen for and respond to requests from SNMP management client systems, such as GET, GETNEXT and GETBULK requests.

By default, the embedded SNMP agent listens on UDP port 161 for polling requests from management systems. You can use the `esxcli system snmp set` command with the `--port` option to configure an alternative port. To avoid conflicting with other services, use a UDP port that is not defined in `/etc/services`.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

**Prerequisites**

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

**Procedure**

**1**   Run the `esxcli system snmp set` command with the `--port` option to configure the port.

For example, run the following command:

```
esxcli system snmp set --port port
```

Here, *port* is the port the SNMP agent uses to listen for polling requests.

**Note**   The port you specify must not be already in use by other services. Use IP addresses from the dynamic range, port 49152 and up.

**2**   (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

## Configure ESXi for SNMP v1 and v2c

When you configure the ESXi SNMP agent for SNMP v1 and v2c, the agent supports sending notifications and receiving GET requests.

In SNMP v1 and v2c authentication is performed by using community strings. Community strings are namespaces which contain one or more managed objects. This form of authentication does not secure the communication between the SNMP agent and the management system. To secure the SNMP communication in your environment, use SNMP v3.

**Procedure**

**1**   Configure SNMP Communities

To enable the ESXi SNMP agent to send and receive SNMP v1 and v2c messages, you must configure at least one community for the agent.

**2**   Configure the SNMP Agent to Send SNMP v1 or v2c Notifications

You can use the ESXi SNMP agent to send virtual machine and environmental notifications to management systems.

## Configure SNMP Communities

To enable the ESXi SNMP agent to send and receive SNMP v1 and v2c messages, you must configure at least one community for the agent.

An SNMP community defines a group of devices and management systems. Only devices and management systems that are members of the same community can exchange SNMP messages. A device or management system can be a member of multiple communities.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

**Prerequisites**

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

**Procedure**

◆   Run the `esxcli system snmp set` command with the `--communities` option to configure an SNMP community.

For example, to configure public, East, and West network operation centers communities, run the following command:

```
esxcli system snmp set --communities public,eastnoc,westnoc
```

Each time you specify a community with this command, the settings you specify overwrite the previous configuration. To specify multiple communities, separate the community names with a comma.

## Configure the SNMP Agent to Send SNMP v1 or v2c Notifications

You can use the ESXi SNMP agent to send virtual machine and environmental notifications to management systems.

To send SNMP v1/v2c notifications with the SNMP agent, you must configure the target (receiver) unicast address, community, and an optional port. If you do not specify a port, the SNMP agent sends traps to UDP port 162 on the target management system by default.

To configure SNMP v3 traps, see Configure SNMP v3 Targets.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

Procedure

**1** Run the `esxcli system snmp set` command with the `--targets` option:

```
esxcli system snmp set --targets target_address@port/community
```

Here, *target_address* is the address of the target system, *port* is the port number to send the notifications to, and *community* is the community name.

Each time you specify a target with this command, the settings you specify overwrite all previously specified settings. To specify multiple targets, separate them with a comma.

For example, run the following command for configuring the targets 192.0.2.1@163/westnoc and 2001:db8::1@163/eastnoc:

```
esxcli system snmp set --targets 192.0.2.1@163/westnoc,2001:db8::1@163/eastnoc
```

**2** (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

**3** (Optional) Send a test trap to verify that the agent is configured correctly by running the `esxcli system snmp test` command.

The agent sends a `warmStart` trap to the configured target.

## Configure ESXi for SNMP v3

When you configure the ESXi SNMP agent for SNMP v3, the agent supports sending informs as well as traps. SNMP v3 also provides stronger security than v1 or v2c, including key authentication and encryption.

An inform is a notification that the sender will resend up to 3 times or until the notification is acknowledged by the receiver.

**Procedure**

**1** Configure the SNMP Engine ID

Every SNMP v3 agent has an engine ID which serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

**2** Configure SNMP Authentication and Privacy Protocols

SNMP v3 optionally supports authentication and privacy protocols.

**3** Configure SNMP Users

You can configure up to 5 users who can access SNMP v3 information. User names must be no more than 32 characters long.

**4** Configure SNMP v3 Targets

Configure SNMP v3 targets to allow the ESXi SNMP agent to send SNMP v3 traps and informs.

## Configure the SNMP Engine ID

Every SNMP v3 agent has an engine ID which serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

If you do not specify an engine ID, when you enable the SNMP agent, an engine ID is automatically generated.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

**Prerequisites**

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

**Procedure**

◆ Run the `esxcli system snmp set` command with the `--engineid` option to configure the SNMP engine ID.

For example, run the following command:

```
esxcli system snmp set --engineid id
```

Here, *id* is the engine ID and it must be a hexadecimal string between 5 and 32 characters long.

## Configure SNMP Authentication and Privacy Protocols

SNMP v3 optionally supports authentication and privacy protocols.

Authentication is used to ensure the identity of users. Privacy allows for encryption of SNMP v3 messages to ensure confidentiality of data. These protocols provide a higher level of security than is available in SNMP v1 and v2c, which use community strings for security.

Both authentication and privacy are optional. However, you must enable authentication in order to enable privacy.

The SNMP v3 authentication and privacy protocols are licensed vSphere features and might not be available in some vSphere editions.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

### Prerequisites

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

### Procedure

1  (Optional) Run the `esxcli system snmp set` command with the `--authentication` option to configure authentication.

   For example, run the following command:

   ```
   esxcli system snmp set --authentication protocol
   ```

   Here, *protocol* must be either **none** (for no authentication), **SHA1**, or **MD5**.

2  (Optional) Run the `esxcli system snmp set` command with the `--privacy` option to configure privacy.

   For example, run the following command:

   ```
   esxcli system snmp set --privacy protocol
   ```

   Here, *protocol* must be either **none** (for no privacy) or **AES128**.

## Configure SNMP Users

You can configure up to 5 users who can access SNMP v3 information. User names must be no more than 32 characters long.

While configuring a user, you generate authentication and privacy hash values based on the user's authentication and privacy passwords and the SNMP agent's engine ID. If you change the engine ID, the authentication protocol, or the privacy protocol after configuring users, the users are no longer valid and must be reconfigured.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

Prerequisites

- Verify that you have configured the authentication and privacy protocols before configuring users.

- Verify that you know the authentication and privacy passwords for each user you plan to configure. Passwords must be at least 7 characters long. Store these passwords in files on the host system.

- Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

Procedure

1  If you are using authentication or privacy, get the authentication and privacy hash values for the user by running the `esxcli system snmp hash` command with the `--auth-hash` and `--priv-hash` flags.

   For example, run the following command:

   ```
   esxcli system snmp hash --auth-hash secret1 --priv-hash secret2
   ```

   Here, *secret1* is the path to the file containing the user's authentication password and *secret2* is the path to the file containing the user's privacy password.

   Alternatively, you can pass the `--raw-secret` flag and specify the passwords directly on the command line.

   For example, you can run the following command:

   ```
   esxcli system snmp hash --auth-hash authsecret --priv-hash privsecret --raw-secret
   ```

   The produced output might be the following:

   ```
   Authhash: 08248c6eb8b333e75a29ca0af06b224faa7d22d6
   Privhash: 232ba5cbe8c55b8f979455d3c9ca8b48812adb97
   ```

   The authentication and privacy hash values are displayed.

**2** Configure the user by running the `esxcli system snmp set` command with the `--users` flag.

For example, you can run the following command:

```
esxcli system snmp set --users userid/authhash/privhash/security
```

The command accepts the following parameters:

| Parameter | Description |
|---|---|
| *userid* | The user name. |
| *authhash* | The authentication hash value. |
| *privhash* | The privacy hash value. |
| *security* | The level of security enabled for that user, which can be *auth* (for authentication only), *priv* (for authentication and privacy), or *none* (for no authentication or privacy). |

For example, run the following command to configure user1 for access with authentication and privacy:

```
esxcli system snmp set --users user1/08248c6eb8b333e75a29ca0af06b224faa7d22d6/
232ba5cbe8c55b8f979455d3c9ca8b48812adb97/priv
```

You must run the following command to configure user2 for access with no authentication or privacy:

```
esxcli system snmp set --users user2/-/-/none
```

**3** (Optional) Test the user configuration by running the following command:

```
esxcli system snmp test --users username --auth-hash secret1 --priv-hash secret2
```

If the configuration is correct, this command returns the following message: "User *username* validated correctly using engine id and security level: *protocols*". Here, *protocols* indicates the security protocols configured.

## Configure SNMP v3 Targets

Configure SNMP v3 targets to allow the ESXi SNMP agent to send SNMP v3 traps and informs.

SNMP v3 allows for sending both traps and informs. An inform is a message that the sender will resend a maximum of three times, waiting 5 seconds between each attempt, unless the message is acknowledged by the receiver.

You can configure a maximum of three SNMP v3 targets, in addition to a maximum of three SNMP v1/v2c targets.

To configure a target, you must specify a hostname or IP address of the system that will receive the traps or informs, a user name, a security level, and whether to send traps or informs. The security level can be either **none** (for no security), **auth** (for authentication only), or **priv** (for authentication and privacy).

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

**Prerequisites**

■ Ensure that the users who will access the traps or informs are configured as SNMP users for both the ESXi SNMP agent and the target management system.

■ If you are configuring informs, you need the engine ID for the SNMP agent on the remote system that will receive the informs.

■ Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

**Procedure**

1  (Optional) If you are configuring informs, configure the remote users by running the `esxcli system snmp set` command with the `--remote-users` option.

   For example, run the following command:

   ```
   esxcli system snmp set --remote-users userid/auth-protocol/auth-hash/priv-protocol/priv-hash/
   engine-id
   ```

   The command accepts the following parameters:

   | Parameter | Description |
   | --- | --- |
   | *userid* | The user name. |
   | *auth-protocol* | The authentication protocol, **none** (for no authentication), **MD5**, or **SHA1**. |
   | *auth-hash* | The authentication hash or **–** if the authentication protocol is **none**. |
   | *priv-protocol* | The privacy protocol, **AES128** or **none**. |
   | *priv-hash* | The privacy hash, or **–** if the privacy protocol is **none**. |
   | *engine-id* | The engine ID of the SNMP agent on the remote system that will receive the informs. |

2  Run the `esxcli system snmp set` command with the `--v3targets` option.

   For example, run the following command:

   ```
   esxcli system snmp set --v3targets hostname@port/userid/secLevel/message-type
   ```

   The parameters of the command are as follows.

   | Parameter | Description |
   | --- | --- |
   | *hostname* | The host name or IP address of the management system that will receive the traps or informs. |
   | *port* | The port on the management system that will receive the traps or informs. If you do not specify a port, the default port, 162, is used. |

| Parameter | Description |
|---|---|
| *userid* | The user name. |
| *secLevel* | The level of authentication and privacy you have configured. Use **auth** if you have configured authentication only, **priv** if you have configured both authentication and privacy, and **none** if you have configured neither. |
| *message-type* | The type of the messages received by the management system. Use **trap** or **inform**. |

**3**  (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

**4**  (Optional) Send a test notification to verify that the agent is configured correctly by running the `esxcli system snmp test` command.

The agent sends a `warmStart` notification to the configured target.

## Configure the Source of Hardware Events Received by the SNMP Agent

You can configure the ESXi SNMP agent to receive hardware events either from IPMI sensors or CIM indications.

IPMI sensors are used for hardware monitoring in ESX/ESXi 4.x and earlier. The conversion of CIM indications to SNMP notifications is available in ESXi 5.0 and later.

If you run ESXCLI commands through vCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *vSphere Command-Line Interface Concepts and Examples*.

**Prerequisites**

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

**Procedure**

**1**  Run the `esxcli system snmp set --hwsrc` *source* command to configure the source for hardware events.

Here, *source* is **sensors** or **indications**, for hardware event received from IPMI sensors or CIM indications respectively.

**2**  (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

## Configure the SNMP Agent to Filter Notifications

You can configure the ESXi SNMP agent to filter out notifications if you don't want your SNMP management software to receive those notifications.

**Prerequisites**

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with vSphere Command-Line Interfaces* for more information on how to use ESXCLI.

**Procedure**

**1** Run the `esxcli system snmp set` command to filter notifications:

```
esxcli system snmp set --notraps oid_list
```

Here, *oid_list* is a list of OIDs for the notifications to filter, separated by commas. This list replaces any OIDs that were previously specified using this command.

For example, to filter out `coldStart` (OID **1.3.6.1.4.1.6876.4.1.1.0**) and `warmStart` (OID **1.3.6.1.4.1.6876.4.1.1.1**) traps, run the following command:

```
esxcli system snmp set --notraps 1.3.6.1.4.1.6876.4.1.1.0,1.3.6.1.4.1.6876.4.1.1.1
```

**2** (Optional) If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

**Results**

The traps identified by the specified OIDs are filtered out of the output of the SNMP agent, and are not sent to SNMP management software.

**What to do next**

To clear all notification filters, run the `esxcli system snmp set --notraps reset` command.

## Configure SNMP Management Client Software

After you have configured a vCenter Server instance or an ESXi host to send traps, you must configure your management client software to receive and interpret those traps.

To configure your management client software, specify the communities for the managed device, configure the port settings, and load the VMware MIB files. See the documentation for your management system for specific instructions for these steps.

**Prerequisites**

To complete this task, download the VMware MIB files from https://kb.vmware.com/s/article/1013445.

**Procedure**

1   In your management software, specify the vCenter Server instance or ESXi host as an SNMP-based managed device.

2   If you are using SNMP v1 or v2c, set up appropriate community names in the management software.

    These names must correspond to the communities set for the SNMP agent on the vCenter Server instance or ESXi host.

3   If you are using SNMP v3, configure users and authentication and privacy protocols to match those configured on the ESXi host.

4   If you configured the SNMP agent to send traps to a port on the management system other than the default UDP port 162, configure the management client software to listen on the port you configured.

5   Load the VMware MIBs into the management software so you can view the symbolic names for vCenter Server or the host variables.

    To prevent lookup errors, load these MIB files in the following order before loading other MIB files:

    a   `VMWARE-ROOT-MIB.mib`

    b   `VMWARE-TC-MIB.mib`

    c   `VMWARE-PRODUCTS-MIB.mib`

**Results**

The management software can now receive and interpret traps fromvCenter Server or ESXi hosts.

## SNMP Diagnostics

You can use SNMP tools to diagnose configuration problems.

- Run the `esxcli system snmp test` command from the vSphere CLI set to prompt the SNMP agent to send a test `warmStart` trap.

- Run the `esxcli system snmp get` command to display the current configuration of the SNMP agent.

- The `SNMPv2-MIB.mib` file provides several counters to aid in debugging SNMP problems. See SNMPv2 Diagnostic Counters.

- The `VMWARE-ESX-AGENTCAP-MIB.mib` file defines the capabilities of the VMware SNMP agents by product version. Use this file to determine if the SNMP functionality that you want to use is supported.

# Monitor Guest Operating Systems with SNMP

You can use SNMP to monitor guest operating systems or applications running in virtual machines.

The virtual machine uses its own virtual hardware devices. Do not install agents in the virtual machines that are intended to monitor physical hardware.

**Procedure**

◆ Install the SNMP agents you normally would use for that purpose in the guest operating systems.

# VMware MIB Files

VMware MIB files define the information provided by ESXi hosts and vCenter Server to SNMP management software.

You can download these MIB files from https://kb.vmware.com/s/article/1013445.

The table VMware MIB Files lists the MIB files provided by VMware and describes the information that each file provides.

Table 10-1. VMware MIB Files

| MIB File | Description |
| --- | --- |
| VMWARE—ROOT—MIB.mib | Contains VMware's enterprise OID and top level OID assignments. |
| VMWARE—ESX—AGENTCAP—MIB.mib | Defines the capabilities of the VMware agents by product versions. This file is optional and might not be supported by all management systems. |
| VMWARE—CIMOM—MIB.mib | Defines variables and trap types used to report on the state of the CIM Object Management subsystem. |
| VMWARE—ENV—MIB.mib | Defines variables and trap types used to report on the state of physical hardware components of the host computer. Enables conversion of CIM indications to SNMP traps. |
| VMWARE—OBSOLETE—MIB.mib | For use with versions of ESX/ESXi prior to 4.0. Defines OIDs that have been made obsolete to maintain backward compatibility with earlier versions of ESX/ESXi. Includes variables formerly defined in the files VMWARE-TRAPS-MIB.mib and VMWARE-VMKERNEL-MIB.mib. |
| VMWARE—PRODUCTS—MIB.mib | Defines OIDs to uniquely identify each SNMP agent on each VMware platform by name, version, and build platform. |
| VMWARE—RESOURCES—MIB.mib | Defines variables used to report information on resource usage of the VMkernel, including physical memory, CPU, and disk utilization. |
| VMWARE—SYSTEM—MIB.mib | The VMWARE—SYSTEM—MIB.mib file is obsolete. Use the SNMPv2-MIB to obtain information from sysDescr.0 and sysObjec ID.0. |
| VMWARE—TC—MIB.mib | Defines common textual conventions used by VMware MIB files. |

### Table 10-1. VMware MIB Files (continued)

| MIB File | Description |
|---|---|
| VMWARE–VC–EVENTS–MIB.mib | Defines traps sent by vCenter Server. Load this file if you use vCenter Server to send traps. |
| VMWARE–VMINFO–MIB.mib | Defines variables for reporting information about virtual machines, including virtual machine traps. |

The table Other MIB Files lists MIB files included in the VMware MIB files package that are not created by VMware. These can be used with the VMware MIB files to provide additional information.

### Table 10-2. Other MIB Files

| MIB File | Description |
|---|---|
| ENTITY–MIB.mib | Allows the description of relationships between physical entities and logical entities managed by the same SNMP agent. See RFC 4133 for more information. |
| HOST–RESOURCES–MIB.mib | Defines objects that are useful for managing host computers. |
| HOST–RESOURCES–TYPES.mib | Defines storage, device, and filesystem types for use with HOST–RESOURCES–MIB.mib. |
| IEEE8021–BRIDGE–MIB | Defines objects for managing devices that support IEEE 802.1D. |
| IEEE8023–LAG–MIB | Defines objects for managing devices that support IEEE 802.3ad link aggregation. |
| IEEE8021–Q–BRIDGE–MIB | Defines objects for managing Virtual Bridged Local Area Networks. |
| IF–MIB.mib | Defines attributes related to physical NICs on the host system. |
| IP–MIB.mib | Defines objects for managing implementations of the Internet Protocol (IP) in an IP version independent manner. |
| IP–FORWARD–MIB.mib | Defines objects for managing IP forwarding. |
| LLDP–V2–MIB.mib | Defines objects for managing devices using Linked Layer Discovery Protocol (LLDP). |
| SNMPv2–CONF.mib | Defines conformance groups for MIBs. |
| SNMPv2–MIB.mib | Defines the SNMP version 2 MIB objects. |
| SNMPv2–SMI.mib | Defines the Structure of Managment Information for SNMP version 2. |
| SNMPv2–TC.mib | Defines textual conventions for SNMP version 2. |
| TCP–MIB.mib | Defines objects for managing devices using the TCP protocol. |
| UDP–MIB.mib | Defines objects for managing devices using the UDP protocol. |

# SNMPv2 Diagnostic Counters

The SNMPv2–MIB.mib file provides a number of counters to aid in debugging SNMP problems.

Table 10-3. Diagnostic Counters from SNMPv2-MIB lists some of these diagnostic counters.

Table 10-3. Diagnostic Counters from SNMPv2-MIB

| Variable | ID Mapping | Description |
|---|---|---|
| snmpInPkts | snmp 1 | The total number of messages delivered to the SNMP entity from the transport service. |
| snmpInBadVersions | snmp 3 | The total number of SNMP messages that were delivered to the SNMP entity and were for an unsupported SNMP version. |
| snmpInBadCommunityNames | snmp 4 | The total number of community-based SNMP messages delivered to the SNMP entity that used an invalid SNMP community name. |
| snmpInBadCommunityUses | snmp 5 | The total number of community-based SNMP messages delivered to the SNMP entity that represented an SNMP operation that was not allowed for the community named in the message. |
| snmpInASNParseErrs | snmp 6 | The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. |
| snmpEnableAuthenTraps | snmp 30 | Indicates whether the SNMP entity is permitted to generate authenticationFailure traps. The value of this object overrides any configuration information. It therefore provides a means of disabling all authenticationFailure traps. |
| snmpSilentDrops | snmp 31 | The total number of Confirmed Class PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate Response Class PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request. |
| snmpProxyDrops | snmp 32 | The total number of Confirmed Class PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in a manner other than a time-out such that no Response Class PDU could be returned. |

# System Log Files

<span style="float:right; font-size:3em; color:#999;">11</span>

In addition to lists of events and alarms, vSphere components generate assorted logs.

These logs contain additional information about activities in your vSphere environment.

This chapter includes the following topics:

- View System Log Entries
- View System Logs on an ESXi Host
- External System Logs
- Export System Log Files
- Managing ESXi Log Files
- Configure Syslog on ESXi Hosts
- Configuring Logging Levels for the Guest Operating System
- Collecting Log Files
- Viewing Log Files with the Log Browser

## View System Log Entries

You can view system logs generated by vSphere components.

These instructions apply only to vCenter Server management nodes.

**Procedure**

**1** In the vSphere Web Client navigate to a vCenter Server.

**2** From the Monitor tab, click **System Logs**.

**3** From the drop-down menu, select the log and entry you want to view.

## View System Logs on an ESXi Host

You can use the direct console interface to view the system logs on an ESXi host. These logs provide information about system operational events.

**Procedure**

**1**   From the direct console, select **View System Logs**.

**2**   Press a corresponding number key to view a log.

vCenter Server agent (vpxa) logs appear if the host is managed by vCenter Server.

**3**   Press Enter or the spacebar to scroll through the messages.

**4**   (Optional) Perform a regular expression search.

a   Press the slash key (/).

b   Type the text to find.

c   Press Enter

The found text is highlighted on the screen.

**5**   Press q to return to the direct console.

# External System Logs

VMware technical support might request several files to help resolve any issues you have with the product. This section describes the types and locations of log files found on various ESXi component systems.

**Note**   On Windows systems, several log files are stored in the Local Settings directory, which is located at `C:\Documents and Settings\<user name>\Local Settings\`. This folder is hidden by default.

## ESXi System Logs

You might need the ESXi system log files to resolve technical issues.

The ESXi system logs can be found in the `/var/run/log` directory.

## vSphere Client System Logs

You might need the vSphere Client system log files to resolve technical issues.

Table 11-1. vSphere Client System Logs lists log files associated with the vSphere Client machine.

Table 11-1. vSphere Client System Logs

| Component | Location |
|---|---|
| vSphere Client Installation log | Temp directory on the vSphere Client machine.<br>Pre-Windows 2008 example: `C:\Documents and Settings\Local Settings\Temp\vminst.log` or `vim-vic-msi.log`<br>Windows 2008 and Windows 7 example: `C:\Users\`*user_name*`\AppData\Local\Temp\vminst.log` or `vim-vic-msi.log` |
| vSphere Client Service log | `\vpx` directory in the Application Data directory on the vSphere Client machine.<br>Pre-Windows 2008 example: `C:\Documents and Settings\`*user_name*`\Local Settings\Application Data\VMware\vpx\viclient-x.log`<br>Windows 2008 and Window 7 example: `C:\Users\`*user_name*`\Local Settings\AppData\Local\VMware\vpx\viclient-x.log`<br>`x(=0, 1, ... 9)` |

# Export System Log Files

When the vSphere Web Client is connected to vCenter Server, you can select hosts from which to download system log files.

To save diagnostic data for ESXi hosts and vCenter Server, the vSphere Web Client must be connected to the vCenter Server system.

Required privileges:

■ To view diagnostic data: **Read-Only User**

■ To manage diagnostic data: **Global.Diagnostics**

Procedure

1  In the inventory, navigate to a vCenter Server instance and click the **Monitor** tab.

2  Click **System Logs**.

3  Click **Export System Logs**.

4  If you are connected to vCenter Server, select the object for which you want to export data.

   Selecting an object selects all of its child objects.

5  If you are connected to vCenter Server, select **Include information from vCenter Server and vSphere Client** to download vCenter Server and vSphere Client log files and host log files, and click **Next**.

6  If the selected host supports manifest driven exports of system log files, select the system log files to collect. Select the specific system log files to download.

   If the host does not support manifest exports of log files, all system log files are exported.

7  Select **Gather performance data** to include performance data information in the log files. Click **Generate Log Bundle**.

   You can update the duration and interval time you want the data collected.

**8** Once the log bundle generates, click **Download Log Bundle** to save it to your local computer.

**9** Specify the location to which to save the log files.

The host or vCenter Server generates a `tar.tgz` file containing the log files. The **Recent Tasks** panel shows the Generate diagnostic bundles task in progress.

**10** Click **Save**.

The Downloading Log Bundles dialog box appears when the Generating Diagnostic Bundle task is finished. The download status of each bundle appears in the dialog box.

Some network errors can cause download failures. When you select an individual download in the dialog box, the error message for that operation appears under the name and location of the log bundle file.

**11** Verify the information in the Summary and click **Finish** to download the log files.

Results

Diagnostic bundles containing log files for the specified objects are downloaded to the location specified.

# Managing ESXi Log Files

Log files are an important component of troubleshooting attacks and obtaining information about breaches of host security. Logging to a secure, centralized log server can help prevent log tampering. Remote logging also provides a long-term audit record.

Take the following measures to increase the security of the host.

- Configure persistent logging to a datastore. By default, the logs on ESXi hosts are stored in the in-memory file system. Therefore, they are lost when you reboot the host, and only 24 hours of log data is stored. When you enable persistent logging, you have a dedicated record of server activity available for the host.

- Remote logging to a central host allows you to gather log files onto a central host, where you can monitor all hosts with a single tool. You can also do aggregate analysis and searching of log data, which might reveal information about things like coordinated attacks on multiple hosts.

- Configure remote secure syslog on ESXi hosts using a remote command line such as vCLI or PowerCLI, or using an API client.

- Query the syslog configuration to make sure that a valid syslog server has been configured, including the correct port.

# Configure Syslog on ESXi Hosts

All ESXi hosts run a syslog service (`vmsyslogd`), which logs messages from the VMkernel and other system components to log files.

You can use the vSphere Web Client or the `esxcli system syslog` vCLI command to configure the syslog service.

For more information about using vCLI commands, see *Getting Started with vSphere Command-Line Interfaces*.

Procedure

1  In the vSphere Web Client inventory, select the host.

2  Click the **Manage** tab.

3  In the System panel, click **Advanced System Settings**.

4  Locate the **Syslog** section of the Advanced System Settings list.

5  To set up logging globally, select the setting to change and click the Edit icon.

| Option | Description |
| --- | --- |
| **Syslog.global.defaultRotate** | Sets the maximum number of archives to keep. You can set this number globally and for individual subloggers. |
| **Syslog.global.defaultSize** | Sets the default size of the log, in KB, before the system rotates logs. You can set this number globally and for individual subloggers. |
| **Syslog.global.LogDir** | Directory where logs are stored. The directory can be located on mounted NFS or VMFS volumes. Only the `/scratch` directory on the local file system is persistent across reboots. The directory should be specified as [*datastorename*] *path_to_file* where the path is relative to the root of the volume backing the datastore. For example, the path [`storage1`] `/systemlogs` maps to the path /vmfs/volumes/storage1/systemlogs. |
| **Syslog.global.logDirUnique** | Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by **Syslog.global.LogDir**. A unique directory is useful if the same NFS directory is used by multiple ESXi hosts. |
| **Syslog.global.LogHost** | Remote host to which syslog messages are forwarded and port on which the remote host receives syslog messages. You can include the protocol and the port, for example, `ssl://hostName1:1514`. UDP (default), TCP, and SSL are supported. The remote host must have syslog installed and correctly configured to receive the forwarded syslog messages. See the documentation for the syslog service installed on the remote host for information on configuration. |

6  (Optional) To overwrite the default log size and log rotation for any of the logs.

   a  Click the name of the log you that want to customize.

   b  Click the Edit icon and enter the number of rotations and log size you want.

7  Click **OK**.

Results

Changes to the syslog options take effect immediately.

# Configuring Logging Levels for the Guest Operating System

Virtual machines can write support and troubleshooting information into a virtual machine log file stored on a VMFS volume. The default settings for virtual machines are appropriate for most situations.

If your environment relies heavily on using VMotion, or if the defaults do not seem suitable for other reasons, you can modify the logging settings for virtual machine guest operating systems.

New log file creation happens as follows:

- Each time you power on or resume a virtual machine, and each time you migrate a virtual machine with VMotion, a new log file is created.

- Each time an entry is written to the log, the size of the log is checked. If `vmx.log.rotateSize` is set to a nondefault value, and the size is over the limit, the next entry is written to a new log. If the maximum number of log files exists, the oldest log file is deleted.

The default for `vmx.log.rotateSize` is zero (0), which means new logs are created during power on, resume, and so on. You can ensure new log file creation happens more frequently by limiting the maximum size of the log files with the `vmx.log.rotateSize` configuration parameter.

VMware recommends saving 10 log files, each one limited to no less than 2MB. These values are large enough to capture sufficient information to debug most problems. If you need logs for a longer time span, you can set `vmx.log.keepOld` to 20.

## Change the Number of Virtual Machine Log Files

You can change the number of the log files for all virtual machines on an ESXi host or for individual virtual machines.

This procedure discusses limiting the virtual machine log file number on an individual virtual machine.

To limit the number of log files for *all* virtual machines on a host, edit the `/etc/vmware/config` file. If the `vmx.log.KeepOld` property is not defined in the file, you can add it. For example, to keep ten log files for each virtual machine, add the following to `/etc/vmware/config`:

```
vmx.log.keepOld = "10"
```

You can use a PowerCLI script to change this property on all the virtual machines on a host.

You can use the `log.keepOld` parameter to affect all log files, not just the virtual machine log files.

**Prerequisites**

Turn off the virtual machine.

**Procedure**

**1** Find the virtual machine in the vSphere Web Client inventory.

a    Select a data center, folder, cluster, resource pool, or host.

b    Click the **Related Objects** tab and click **Virtual Machines**.

**2** Right-click the virtual machine and click **Edit Settings**.

**3** Select **VM Options**.

**4** Click **Advanced** and click **Edit Configuration**.

**5** Add or edit the `vmx.log.keepOld` parameter to the number of files to keep for this virtual machine.

For example, to keep 20 log files and begin deleting the oldest files as new ones are created, enter `20`.

**6** Click **OK**.

## Control When to Switch to New Virtual Machine Log Files

The `vmx.log.rotateSize` parameter specifies the log file size at which the switch to new log files happens for the logs for individual virtual machines. Use this parameter together with the `vmx.log.keepOld` parameter to ensure acceptable log file sizes without losing critical logging information.

The `vmx.log.keepOld` parameter determines how many virtual machine log file instances the ESXi host retains before overwriting the first log file. The default value of `vmx.log.keepOld` is ten, a suitable number to properly log complex operations such as VMotion. You must increase this number significantly when you change the value of `vmx.log.rotateSize`.

This procedure discusses changing the virtual machine rotate size on an individual virtual machine.

To limit the rotate size for *all* virtual machines on a host, edit the `/etc/vmware/config` file. If the `vmx.log.KeepOld` property is not defined in the file, you can add it. You can use a PowerCLI script to change this parameter for selected virtual machines on a host.

You can use the `log.rotateSize` parameter to affect all log files, not just the virtual machine log files.

You can change the value of `vmx.log.rotateSize` for all virtual machine s from the vSphere Web Client or by using a PowerCLI script.

**Prerequisites**

Turn off the virtual machine.

**Procedure**

1 Find the virtual machine in the vSphere Web Client inventory.

    a    Select a data center, folder, cluster, resource pool, or host.

    b    Click the **Related Objects** tab and click **Virtual Machines**.

2 Right-click the virtual machine and click **Edit Settings**.

3 Select **VM Options**.

4 Click **Advanced** and click **Edit Configuration**.

5 Add or edit the `vmx.log.rotateSize` parameter to the maximum file size before log information is added to a new file, or to the first log file if you have more log files than the `vmx.log.keepOld` parameter specifies.

    Specify the size in bytes.

6 Click **OK**.

# Collecting Log Files

VMware technical support might request several files to help resolve technical issues. The following sections describe script processes for generating and collecting some of these files.

## Set Verbose Logging

You can specify how verbose log files will be.

You can only set verbose logging for vpxd logs.

**Procedure**

1 Select **Administration > vCenter Server Settings**.

2 Select **Logging Options**.

3 Select **Verbose** from the pop-up menu.

4 Click **OK**.

## Collect vSphere Log Files

You can collect vSphere log files into a single location.

Procedure

◆ View the log file using one of the following methods.

| Task | Action |
|------|--------|
| **View the `viclient-*.log` file** | Change to the directory, `%temp%`. |
| **Download the log bundle from vSphere Web Client connected to a vCenter Server system** | To download the log bundle, do the following:<br><br>a   Select **Administration > System Configuration**.<br><br>b   From the Objects tab, select **Actions > Export Support Bundles...**<br><br>The log bundle is generated as a `.zip` file. By default, the vpxd logs within the bundle are compressed as `.tgz` files. You must use `gunzip` to uncompress these files. |
| **Generate vCenter Server log bundles from a vCenter Server system** | Select **Start > Programs > VMware > Generate vCenter Server log bundle**.<br><br>You can use this to generate vCenter Server log bundles even when you are unable to connect to the vCenter Server using the vSphere Client.<br><br>The log bundle is generated as a `.zip` file. By default, the vpxd logs within the bundle are compressed as `.tgz` files. You must use `gunzip` to uncompress these files. |

## Collect ESXi Log Files

You can collect and package all relevant ESXi system and configuration information, a well as ESXi log files. This information can be used to analyze the problems.

Procedure

◆ Run the following script on the ESXi Shell: `/usr/bin/vm-support`

The resulting file has the following format: `esx-`*date-unique-xnumber*`.tgz`

## ESXi Log File Locations

ESXi records host activity in log files, using a syslog facility.

| Component | Location | Purpose |
|-----------|----------|---------|
| VMkernel | `/var/log/vmkernel.log` | Records activities related to virtual machines and ESXi. |
| VMkernel warnings | `/var/log/vmkwarning.log` | Records activities related to virtual machines. |
| VMkernel summary | `/var/log/vmksummary.log` | Used to determine uptime and availability statistics for ESXi (comma separated). |
| ESXi host agent log | `/var/log/hostd.log` | Contains information about the agent that manages and configures the ESXi host and its virtual machines. |
| vCenter agent log | `/var/log/vpxa.log` | Contains information about the agent that communicates with vCenter Server (if the host is managed by vCenter Server). |

| Component | Location | Purpose |
|-----------|----------|---------|
| Shell log | `/var/log/shell.log` | Contains a record of all commands typed into the ESXi Shell as well as shell events (for example, when the shell was enabled). |
| Authentication | `/var/log/auth.log` | Contains all events related to authentication for the local system. |
| System messages | `/var/log/syslog.log` | Contains all general log messages and can be used for troubleshooting. This information was formerly located in the messages log file. |
| Virtual machines | The same directory as the affected virtual machine's configuration files, named vmware.log and vmware*.log. For example, `/vmfs/volumes/`*`datastore`*`/`*`virtual machine`*`/`<br>`vmware.log` | Contains virtual machine power events, system failure information, tools status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and so on. |

## Configure Log Filtering on ESXi Hosts

The log filtering capability lets you modify the logging policy of the syslog service that is running on an ESXi host. You can create log filters to reduce the number of repetitive entries in the ESXi logs and to denylist specific log events entirely.

Log filters affect all log events that are processed by the ESXi host vmsyslogd daemon, whether they are recorded to a log directory or to a remote syslog server.

When you create a log filter, you set a maximum number of log entries for the log messages that are generated by one or more specified system components and that match a specified phrase. You must enable the log filtering capability and reload the syslog daemon to activate the log filters on the ESXi host.

**Important** If you set a limit to the amount of logging information, you might be unable to properly troubleshoot potential system failures. If a log rotate occurs after the maximum number of log entries is reached, you might lose all instances of a filtered message.

Procedure

1 Log in to the ESXi Shell as root.

2 In the `/etc/vmware/logfilters` file, add the following entry to create a new log filter.

```
numLogs | ident | logRegexp
```

where:

- *numLogs* sets the maximum number of log entries for the specified log messages. After reaching this number, the specified log messages are filtered and ignored. Use `0` to filter and ignore all the specified log messages.

- *ident* specifies one or more system components to apply the filter to the log messages that these components generate. For information about the system components that generate log messages, see the values of the `idents` parameters in the syslog configuration files that are located in the `/etc/vmsyslog.conf.d` directory. Use a comma-separated list to apply a filter to more than one system component. Use **\*** to apply a filter to all system components.

- *logRegexp* specifies a case-sensitive phrase with Python regular expression syntax to filter the log messages by their content.

For example, if you want to set a limit of maximum two log entries from the hostd component for messages that resemble the `SOCKET connect failed, error 2: No such file or directory` phrase with any error number, add the following entry:

```
2 | hostd | SOCKET connect failed, error .*: No such file or directory
```

**Note**  A line starting with # denotes a comment and the rest of the line is ignored.

3   In the /etc/vmsyslog.conf file, add the following entry to enable the log filtering capability.

```
enable_logfilters = true
```

4   Run the `esxcli system syslog reload` command to reload the syslog daemon and apply the configuration changes.

## Turn Off Compression for vpxd Log Files

By default, vCenter Server vpxd log files are rolled up and compressed into `.gz` files. You can turn off this setting to leave the vpxd logs uncompressed.

**Procedure**

1   Log in to the vCenter Server using the vSphere Web Client.

2   Select **Administration > vCenter Server Settings**.

3   Select **Advanced Settings**.

4   In the **Key** text box, type `log.compressOnRoll`.

5   In the **Value** text box, type `false`.

6   Click **Add**, and click **OK**.

## ESXi VMkernel Files

If the VMkernel fails, an error message appears and then the virtual machine reboots. If you specified a VMware core dump partition when you configured your virtual machine, the VMkernel also generates a core dump and error log.

More serious problems in the VMkernel can freeze the machine without an error message or core dump.

# Viewing Log Files with the Log Browser

You can view, search, and export one or more vCenter Server and ESXi log files at a time using the log browser.

You can also export, manage, and view different log types.

**Note** The log browser cannot be used for the Platform Services Controller in the vCenter Server Appliance or vCenter Server on Windows.

## Retrieve Logs

When you retrieve logs for a host or vCenter Server, you can use these logs to view, search, filter, and compare with other system logs.

**Procedure**

1   Navigate to the host or vCenter Server that contain the logs you want to retrieve.

2   Click the **Monitor** tab.

3   Click **Log Browser**.

4   (Optional) If no logs for the host or vCenter Server are available, click **Retrieve now** to retrieve the logs for that object.

    The retrieved logs are based on a current snapshot of the system. Retrieving logs can take a few minutes. You can perform other tasks while the logs are being retrieved.

5   (Optional) Click **Refresh** to retrieve newer logs.

6   Select the type of log you want to browse.

**Results**

The log displays in the browser.

## Search Log Files

You can search the log files by text or by time.

**Prerequisites**

If the log is unavailable, you must retrieve it. See Retrieve Logs.

**Procedure**

1   Navigate to the Log Browser.

2   Select the type of log you want to browse.

3   In the **Actions** menu, select either **Find by Text** or **Find by Time**.

4   In the search area at the bottom of the Log Browser type the text or select the time you want to search.

**5**  Click **Next** to view the next line containing the text or time searched or **Previous** to view the previous line.

The Log Browser displays the line (in the 3rd row) that contains the text or time you searched.

## Filter Log Files

The Log Browser displays filtered searches.

**Procedure**

**1**  Navigate to the Log Browser and select a log file to browse.

**2**  Select the number of **Adjacent** lines you want the Log Browser to display.

**3**  Type the text that you want filtered in the search box.

When displaying adjacent lines, groups of consecutive entries are highlighted with a different background color.

**Results**

The log browser displays the lines in the log that contain the word you typed, along with the number of adjacent (before and after) lines.

## Create Advanced Log Filters

You can create and save log file filters to your local system.

**Procedure**

**1**  Navigate to the Log Browser and retrieve a log file object.

**2**  Click **Advanced filter**.

**3**  Enter the conditions you want included in the filter.

**4**  Type a filter name.

**5**  Click **Save** to save the filter.

The filter is saved on the vSphere Web Client server and is available the next time you start the vSphere Web Client.

**6**  Click **Filter** to view the results in the Log Browser.

**What to do next**

You can load saved filters from your local system by clicking **Save to local system**. The filters are saved in XML format. You can also load filters from an XML file from your local system by clicking **Load from local system**.

# Adjust Log Times

You might want to adjust the times in the log files to a different time zone or to compare multiple log files.

**Procedure**

**1** Navigate to the Log Browser and retrieve a log file from an object.

**2** In the **Actions** menu, select **Adjust by Time**.

**3** Select **Add** or **Subtract** and adjust the **days**, **hours**, **minutes**, **seconds**, or **milliseconds** from the original time stamps in the log.

The adjusted time stamp appears.

**4** (Optional) Click **Reset** to adjust the time stamp back to the original times.

**5** Click **Apply**.

**Results**

The log browser displays the log entries with the adjusted times.

# Export Logs from the Log Browser

You can export log files using the Log Browser.

**Procedure**

**1** Navigate to the Log Browser and select an object to browse.

**2** Select **Action > Export**.

**3** Select the type of file that you want to export.

Selecting the log file bundle downloads a potentially large file.

**4** Click **Export**.

After a few seconds a new web browser displays.

**5** Enter the location where you want to save the file.

**Results**

The log file is downloaded to your local machine and you can close the new browser window.

# Compare Log Files

You can open multiple windows in the Log Browser to compare log files.

**Procedure**

**1** Navigate to the Log Browser and retrieve a log file from an object.

**2** Select **Actions > New Browser Window** to open a window in the Log Browser.

**3** In the Log Browser window, retrieve another log file to view.

You can perform the same actions with the log file opened in the new window as you can with the original Log Browser window.

## Manage Logs Using the Log Browser

From the Log Browser, you can update, remove, and see a list of available log file bundles.

To manage log file bundles, you must access the Log Browser from the vSphere Web Client home.

### Procedure

**1** From the vSphere Web Client home, click **Log Browser**.

**2** Click the **Manage** tab.

**3** Select an object's log file in the list of sources.

**4** (Optional) Click **Remove** to delete the log file bundle.

Deleting the log file bundle reclaims disk space on the vSphere Web Client server.

All log files generated from that bundle are deleted.

**5** (Optional) Click **Update** to update the list of log file bundles.

You can view the log bundles created by other vSphere Web Client sessions.

### Results

The log does not appear in the retrievable objects list.

## Browse Log Files from Different Objects

You can browse multiple log files coming from different objects within the Log Browser at the same time. This is helpful if you want to simultaneously compare log files.

### Procedure

**1** From the vSphere Web Client home, click **Log Browser**.

**2** Click the **View** tab.

**3** Select an object (ESXi host or vCenter Server) to view its logs.

**4** Open a new browser window by selecting **Actions > New Browser Window** and select another object to view its logs.