

# VirtualCenter Monitoring and Performance Statistics

VMware Infrastructure 3 includes a system for gathering extensive metrics on performance, resource utilization, and basic statistics of all parts of the environment, including virtual machines, ESX Server hosts, resources pools, and host clusters. These metrics can be viewed using the VI Client, and are also available to external programs via the VMware Infrastructure SDK. All measurements may be obtained in real-time, or from a historic record.

The real-time metrics are obtained by connecting directly to the source of the data and transmitting the desired measurements as they are collected. The metrics are collected every 20 seconds for ESX Server 3 hosts, and every 60 seconds for ESX Server 2 hosts. They can be viewed through a VI Client session with VirtualCenter or directly with an ESX Server host; in the latter case, only metrics pertaining to that host or the virtual machines on it are available.

The historical metrics are drawn from the VirtualCenter database, and can only be viewed with a VI Client connected to the VirtualCenter server, or a program that uses the VMware Infrastructure SDK. VirtualCenter aggregates the value of each metric over successively larger time intervals, using smaller intervals to generate values which are then rolled up into larger intervals, using either an average, summation, latest, maximum or minimum value over the interval.

The historical data available depends upon the following factors:

1. The Statistics Collection Level
2. The Statistics Collection Interval

The former determines which metrics are stored, and can be set between 1 and 4, with 4 being the maximum amount of data. Each higher number setting increases the amount of data collected. The latter determines both the time for which the data is kept, as well as the period over which the data is averaged. Both of these times can be configured, and you can have multiple Statistics Collection Intervals so that different levels of aggregation can be kept for different amounts of time. More information on these two factors may be found in the Basic System Administration Guide in Chapter 16.

The type of information collected at each level is described in the following table.

Statistics Collection Level	Description
Level 1	Includes basic metrics – Average Usage for CPU, Memory, Disk, and Network; System Uptime, System Heartbeat, and DRS metrics. Does not include statistics for devices.
Level 2	Includes all metrics for CPU, Memory, Disk, and Network counters (average, summation, and latest rollup types; does not include maximum and minimum rollup types); System Uptime, System Heartbeat, and DRS metrics. Does not include statistics for devices.
Level 3	Includes all metrics (including device metrics) for all counter groups (average, summation, and latest rollup types; does not include maximum and minimum rollup types).
Level 4	Includes all metrics supported by VirtualCenter, including maximum and minimum rollup types.

*Table 1: Description of Statistics Collection Levels*

Although a thorough discussion of performance and monitoring metrics is beyond the scope of this document, the following gives some guidelines as to when you might choose the different collection levels.

Level 1: General information on utilization of the ESX Server hosts. This data can be used by a datacenter administrator for planning and forecasting. It allows you to see at a high level if the hardware resources are being used fully, or if there are certain resources that are more constrained than others (for example, memory).

Level 2: Full memory data, including detailed information on a per-virtual machine level on memory: swap in/out, memory sharing, etc. Some examples of how to use this information include

- a. Evaluation of virtual machine co-location: determine which virtual machines can be shared effectively on the same server because of complementary memory sharing
- b. Capacity planning: observe how much memory is active, to determine whether or not it is possible to put another virtual machine on the same host

Level 3: Information on a per-device level: per NIC, per HBA, and per core. Also included is detailed information on how much time a virtual machine spends occupying the CPU. This data can be used for tasks such as

- a. Determining effectiveness of VSMP (by comparing ready and wait times for each virtual CPU)
- b. Diagnosing device performance, or comparing the performance of two different devices

Level 4: Adds minimum and maximum rollup types. This information can be used, for example, to see if a device is being saturated.

Level 1 has very low overhead on both the VirtualCenter server as well as the ESX Server hosts. Levels 2 – 4 have slightly greater overhead on ESX, but can adversely impact VirtualCenter performance if there are more than 10 ESX Server hosts. In particular Level 4 can quickly fill the VirtualCenter database, so it should only be turned on for limited periods of time.

In general, for day to day operation, you should keep the Statistics Collection Level at 1 or 2. Basic performance data will be gathered historically, and if you need to see more detailed information, you can view the statistics in real-time, which enables you to see every available metric without modifying the collection level. In certain cases, you might want to record the more detailed metrics over a longer period of time. For this, you could set a higher level for the desired duration, and restore the level to a lower number when you are finished.

The Statistics Collection Level may be changed by a user with Administrator privileges using the VI Client. The setting is found in the VirtualCenter Management Server Configuration window (available from the Administration menu), in the section “Statistics”. A modification of the level takes effect immediately, but you must wait until the next round of data collection to see the change. In the VI Client, you can select which metrics to view in the Performance charts. However, metrics for which data was not collected in the specified display interval will sometimes appear as having a value of zero. Therefore, you need to know which Statistics Collection Level was active during a given period to distinguish between metrics whose value is truly zero versus metrics for which no data was collected. For metrics that have multiple rollup types, the VI Client allows you to select which one to display; in the case of metrics for which the VI Client doesn’t specify the rollup type, you must consult the tables below to determine which is the type is available for display.

The rest of this document lists the metrics available for monitoring from both VirtualCenter Server as well as directly from an ESX Server host. It presents tables divided into categories of statistics:

- CPU on page 4
- Memory on page 5
- Disk on page 7
- Network on page 7
- Resource on page 9
- Cluster on page 10
- System on page 12

The metrics from all of these groups are available via the VMware Infrastructure SDK. However, the Resource metrics are not available from the VI Client. The tables indicate the Statistics Collection Level at which historical records for the metric are kept and the rollup type recorded at that level. Note that historical records for any metric are automatically available at all higher levels too.

The tables use “VC” to indicate metrics available in VirtualCenter Server and “H” to indicate metrics available from a standalone ESX Server host. It also indicates the VMware Infrastructure 3 Inventory Object for which the metric applies: Host Cluster, Resource Pool, Virtual Machine, or ESX Server Host.

## CPU

This table contains counters that measure CPU performance.

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
usage	CPU usage as percentage over the collected interval			VC H	VC H		1		4
usagemhz	CPU usage in MHz over the collected interval	VC	VC H	VC H	VC H		1		4
reservedCapacity	The value in MHz of the reservation property of the host's resource pool, or the sum of the reservation properties of the resource pool's (immediate) children, whichever is larger. Note that children's sum of reservations can be larger than that of the parent only if the parent is marked as reservationExpandable.				VC H		2		
wait	CPU time in msec spent on wait state			VC H				3	
ready	CPU time in msec spent on ready state			VC H				3	
system	CPU time in msec spent on system processes			VC H				3	
extra	CPU time in msec that is extra			VC H				3	
used	CPU time in msec that is used			VC H				3	
guaranteed	CPU time in msec that is guaranteed			VC H				3	

## Memory

This table contains counters that measure Memory performance.

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
usage	Memory usage as percentage of total available memory	VC	VC	VC H	VC H		1		4
consumed	Amount in kB of host memory consumed by the virtual machine for guest memory.		VC	VC H			2		4
zero	Amount in kB of zero memory	VC	VC H	VC H	VC H		2		4
state	Memory state	VC			VC H	2			
heapfree	Free space in kB in memory heap	VC			VC H		2		4
reservedCapacity	Amount of memory in MB that is reserved capacity.				VC H		2		
granted	Amount in kB of memory granted	VC	VC H	VC H	VC H		2		4
vmmemctl	Amount in kB of memory used by memory control	VC	VC H	VC H	VC H		2		4
vmmemctltarget	Amount in kB of memory that can be used by memory control		VC	VC H			2		4
unreserved	Amount in kB of memory that is unreserved	VC			VC H		2		4
swapunreserved	Amount in kB of memory that is unreserved by swap	VC					2		4
active	Amount in kB of memory that is actively used	VC	VC H	VC H	VC H		2		4

<b>Resource</b>		<b>Inventory Object</b>				<b>Collection Level and Rollup Type</b>			
<b>Name</b>	<b>Item Measured</b>	<b>Cluster</b>	<b>Res Pool</b>	<b>VM</b>	<b>Host</b>	<b>Latest</b>	<b>Avg</b>	<b>Sum</b>	<b>Max/Min</b>
shared	Amount in kB of memory that is shared	VC	VC H	VC H	VC H		2		4
swapped	Amount in kB of memory that is used by swap	VC			VC H		2		4
swaptarget	Amount in kB of memory that can be swapped		VC	VC H			2		4
swapped	Amount in kB of memory that is swapped		VC H	VC H			2		4
swpin	Amount in kB of memory that is swapped in		VC	VC H			2		4
swapout	Amount in kB of memory that is swapped out		VC	VC H			2		4
overhead	Amount in kB of memory that is an overhead	VC	VC H	VC H	VC H		2		4
heap	Amount in kB of memory that is allocated for heap	VC			VC H		2		4
sharedcommon	Amount in kB of memory that is shared by common				VC H		2		4

## Disk

This table contains counters that measure Disk performance.

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
usage	The sum of the data read and written for all of the disk instances of the host or virtual machine in kB/sec.			VC H	VC H		1		4
read	Amount of data read from the disk in the period in kB/sec, per HBA			VC H	VC H		3		
write	Amount of data written to the disk in the period in kB/sec, per HBA			VC H	VC H		3		
numberWrite	Number of disk writes in the period, per HBA			VC H	VC H			3	
numberRead	Number of disk reads in the period, per HBA			VC H	VC H			3	

## Network

This table contains counters that measure network performance.

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
usage	The sum of data transmitted and received for all the NIC instances of the host or virtual machine.			VC H	VC H		1		4
received	Rate of data received in the period in kB/sec, per NIC			VC H	VC H		3		
transmitted	Rate of data transmitted in the period in kB/sec, per NIC			VC H	VC H		3		
packetRx	Number of packets transmitted in the period, per NIC			VC H	VC H			3	
packetTx	Number of packets received in the period, per NIC			VC H	VC H			3	



## Resource

This table contains performance counters that relate to the time spent by a virtual machine on a CPU as a percentage of the stated time interval. The active time is defined as the time that a virtual machine spends occupying a CPU. The run time is the time that the virtual machine spends running on the CPU. The scheduling limit is the amount of time that the virtual machine did not run because this would violate the “limit” resource setting of the virtual machine or resource pool (i.e. even though the virtual machine is ready to run, it is prevented from doing so). The statistics include both average and peak times expressed as percentages over intervals of 1 minute, 5 minutes and 15 minutes. At the host level, the metric is a sum over all virtual machines currently running on that host. Note that this set of metrics is only available via the VMware Infrastructure SDK.

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
samplePeriod	sample period in msec			VC H	VC H	3			
sampleCount	sample count			VC H	VC H	3			
maxLimited1	scheduling limit in percentage over the past one minute			VC H	VC H	3			
maxLimited5	scheduling limit in percentage over the past 5 minutes			VC H	VC H	3			
maxLimited15	scheduling limit in percentage over the past 15 minutes			VC H	VC H	3			
runpk1	peak run time in percentage over the past 1 minute			VC H	VC H	3			
runpk5	peak run time in percentage over the past 5 minutes			VC H	VC H	3			
runpk15	peak run time in percentage over the past 15 minutes			VC H	VC H	3			
runav1	average run time in percentage over the past 1 minute			VC H	VC H	3			
runav5	average run time in percentage over the past 5 minutes			VC H	VC H	3			
runav15	average run time in percentage over the			VC H	VC H	3			

<u>Resource</u>		<u>Inventory Object</u>				<u>Collection Level and Rollup Type</u>			
<b>Name</b>	<b>Item Measured</b>	<b>Cluster</b>	<b>Res Pool</b>	<b>VM</b>	<b>Host</b>	<b>Latest</b>	<b>Avg</b>	<b>Sum</b>	<b>Max/Min</b>
	past 15 minutes								
actpk1	peak active time in percentage over the past 1 minute			VC H	VC H	3			
actpk5	peak active time in percentage over the past 5 minutes			VC H	VC H	3			
actpk15	peak active time in percentage over the past 15 minutes			VC H	VC H	3			
actav1	average active time in percentage over the past 1 minute			VC H	VC H	3			
actav5	average active time in percentage over the past 5 minutes			VC H	VC H	3			
actav15	average active time in percentage over the past 15 minutes			VC H	VC H	3			

## Cluster

These metrics pertain to VMware host cluster features, such as DRS and HA

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
effectivecpu	DRS effective CPU resources available in MHz	VC					1		
effectivemem	DRS effective memory resources available in kB	VC					1		
failover	HA number of failures that can be tolerated	VC				1			
cpufairness	An integer between 1 and 100 representing the percentage of CPU resources allocated.	VC				1			
memfairness	An integer between 1 and 100 representing the percentage of memory resources allocated.	VC				1			

## System

This table contains performance counters that relate to system availability

Resource		Inventory Object				Collection Level and Rollup Type			
Name	Item Measured	Cluster	Res Pool	VM	Host	Latest	Avg	Sum	Max/Min
resourceCpuUsage	CPU usage in MHz			VC H	VC H		3		4
uptime	Number of seconds since startup			VC H	VC H			1	
heartbeats	Number of heartbeats in this period			VC H	VC H			1	