ESX Instance Workload Data

Contents

- Overview
- ESX Platform Metrics
- ESX Workload Service Monitor
- ESX Advanced Metrics Service Monitor

Overview

In addition to system performance metrics from the ESX platform, Uptime Infrastructure Monitor also collects workload data for each instance running on an ESX server. These metrics show the consumption of the various physical resources by each ESX instance.

ESX Platform Metrics

Uptime Infrastructure Monitor collects the following workload data (as documented in the VMware specifications) from an instance running on an ESX system:

Metric	Explanation	Group/VMware Performance Counter		
CPU	The percentage of CPU time that is being used by a VMware instance. This is a percentage of the available maximum amount of CPU time.	cpu/usagemhz		
Memo ry	The amount of physical memory, in kilobytes, that is being used by a VMware instance.	mem/active		
Netwo rk IO	The amount of the network I/O capacity, in kilobits per second, that is being used by a VMware instance.	net/usage		
Disk IO	The amount of the disk I/O capacity, in kilobytes per second, that is being used by a VMware instance.	disk/usage		
% Ready	The amount of time that one or more instances running on an ESX v3 server is ready to run, but cannot run because it cannot access the processor on the ESX v3 server.	cpu/ready		
% Used	The percentage of CPU time that an instance running on an ESX v3 server is using.	cpu/used		

ESX Workload Service Monitor

The following metrics are gathered by the ESX Workload service monitor, which is deployed to monitor instances running on an existing ESX system:

Metric	Explanation
CPU Usage	The amount of processor power, measured in megahertz (MHz), that the instances on the ESX server are consuming.
Network Bandwidth Usage	The amount of network traffic in and out of the server, measured in megabits per second (Mbit/s).
Disk Usage	The amount of data being written to the server's hard disk, measured in kilobytes per second (kB/s).
Memory Usage	The amount of overall system memory, measured in megabytes (MB).
Percent Ready	The percentage of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server.
Percent Used	The percentage of CPU time that an instance running on an ESX server is using.

ESX Advanced Metrics Service Monitor

The following metrics are gathered by the ESX Advanced Metrics service monitor, which is deployed to monitor instances running on an existing ESX system:

Metric	Explanation	Guest or Host Metric
Percent Wait	The percentage of CPU time that an instance running on an ESX server spent in a wait state.	guest
Memory Balloon	The amount of memory, in KB, allocated by the virtual machine memory control driver (vmmemctl, installed with VMware Tools). It's a VMware exclusive memory-management driver that controls ballooning.	guest
Memory Balloon Target	The target value, in KB, set by VMkernal for the virtual machine's memory balloon size. In conjunction with the vmmemctl metric, this metric is used by VMkernel to inflate and deflate the balloon for a virtual machine.	guest
Memory Overhead	The amount of machine/host memory, in KB, allocated to a virtual machine beyond its reserved amount.	guest
Memory Swap In	The amount of data, in KB, that has been swapped in to machine memory from the swap file since the virtual machine was powered on.	guest
Memory Swap Out	The amount of data, in KB, that the VMkernel memory has written to the virtual machine's swap file from machine memory. Refers to VMkernel swapping and not to guest OS swapping.	guest
Memory Zero	The amount of guest physical memory, in KB, that contains only 0s, thus can be safely used by other virtual machines for their 0 pages memory requirement.	guest
Memory Swap Used	The average amount of memory, in KB, that is used by swap. Swap is the sum of memory swapped for all VMs powered on, and vSphere services on the host.	host
Memory Swap Target	The amount of memory, in KB, available for swapping, where the target size for a virtual machine swap file is calculated by the VMkernel.	guest
Disk Total Latency	The average amount of time, in milliseconds, taken during the collection interval to process a SCSI command issued by the Guest OS to the virtual machine. The sum of kernelCommandLatency and physical deviceCommandLatency.	host
Disk Kernel Latency	The average amount of time, in milliseconds, spent by VMkernel processing each SCSI command.	host
Disk Device Latency	The average amount of time, in milliseconds, taken to complete a SCSI command from the physical device.	host
Disk Queue Latency	The average amount of time, in milliseconds, spent in the VMkernel queue, per SCSI command, during the collection interval.	host
Disk Commands Aborted	The number of SCSI commands aborted during the collection interval.	host
Disk Commands Issued	The number of SCSI commands issued during the collection interval.	host
Disk Bus Resets	The number of SCSI-bus reset commands issued during the collection interval.	host