

Graphing VMware vSphere Components

- [Overview](#)
- [Viewing the Status of a vSphere Element](#)
- [CPU Performance Graphs](#)
 - [CPU Workload](#)
 - [CPU Usage](#)
 - [Multi-CPU Usage](#)
 - [Wait and Ready Time](#)
- [Memory Usage Graphs](#)
 - [Memory Workload](#)
 - [Memory Usage](#)
 - [Memory Profile](#)
 - [Memory Guest Paging](#)
 - [Memory Swap](#)
- [Process Graphs](#)
 - [Process Workload](#)
 - [Detailed Process Information](#)
 - [Number of Processes](#)
- [Network Graphs](#)
 - [Network Workload](#)
 - [Network I/O](#)
 - [Network Errors](#)
- [Disk Graphs](#)
 - [Disk Workload](#)
 - [Disk I/O](#)
 - [Disk Errors](#)
 - [Disk Latency](#)
 - [Disk Storage Capacity](#)
 - [Disk Storage Profile](#)
- [VMware Management Graphs](#)
 - [VM Instance Motion](#)
 - [Generating an Instance Motion Graph](#)
 - [Power Consumption](#)
 - [Power States](#)

Overview

By generating graphs, you can display performance and usage statistics for Uptime Infrastructure Monitor Elements. Generating graphs helps you diagnose problems through root-cause analysis, as well as review the overall health and performance of monitored Elements. With vSync, you can generate graphs for VMware vSphere components that are monitored by Uptime Infrastructure Monitor, allowing you to diagnose and review all components of your infrastructure—whether virtual assets or physical ones—from the same view.

When viewing the Graphing tab for any VMware vSphere component, you see different graphing options depending on the type of VMware vSphere object, or in the case of VMs, the operating system that is running.

Viewing the Status of a vSphere Element

You can view the status of your VMware vCenter servers, ESX servers, and VMs using Quick Snapshots. The Quick Snapshot summarizes both the recent and current performance of key hardware and process information for a VMware vSphere component that exists in Uptime Infrastructure Monitor as an Element, and can help administrators identify potential issues.



If there are not 24 hours' worth of data available, Uptime Infrastructure Monitor uses data from as far back as possible to generate charts.

The specific contents of the Quick Snapshot depend on the VMware vSphere component currently in focus:

- VMware vCenter server
- ESX server or cluster
- virtual machine
- virtual machine using WMI or the Uptime Infrastructure Monitor agent to collect metrics

The Quick Snapshot is typically used as a preliminary step toward root cause analysis. When you first acknowledge an issue by clicking an Element name on either **Global Scan**, or the **My Alerts** section of **My Portal**, you are shown the Quick Snapshot for that Element. From here, you can work with the information provided in the charts and tables and begin further investigation:


- clicking the expand arrow at the top-right of a chart enlarges it
- in the enlarged chart, click-dragging a start and end point along the timeline expands that specific range
- when viewing an enlarged chart, you print or export it by clicking the context menu icon at the top-right, then making the appropriate choice
- at any zoom range, hovering the mouse pointer along the timeline displays the value for that precise interval
- when more than one metric is displayed, clicking metrics in the legend toggles them on and off, allowing you to focus on a specific metric

vCenter Server Quick Snapshot Contents

The following information is displayed in a VMware vCenter server Quick Snapshot:

Datacenter Summary	
datacenters that are monitored are listed in alphabetical order, along with resource information for the datacenter as a whole	
CPU Capacity Trend	<ul style="list-style-type: none"> the CPU usage trend of the datacenter over the last 24 hours maximum and minimum total CPU usage per hour clicking opens a pop-up displaying the CPU Workload graph for the datacenter showing CPU usage in MHz
Memory Capacity Trend	<ul style="list-style-type: none"> the memory usage trend of the datacenter over the last 24 hours maximum and minimum total CPU usage per hour clicking opens a pop-up display the Memory Workload graph for the datacenter, showing the memory consumed metric
Total Active CPU	the total CPU cycles, in GHz, that are available, whether they are currently used
Total Active Memory	the total available memory, in GB, that is available, whether it is currently in use
Running VMs and Hosts Count	<ul style="list-style-type: none"> the current total number of VMs and ESX hosts running within the datacenter clicking opens a custom graph in a pop-up window, displaying the number of VMs and ESX hosts running during the time period
Active Outages	<ul style="list-style-type: none"> the current number of Uptime Infrastructure Monitor service monitors that are attached to any Element within the datacenter that are presently in a WARN or CRIT state clicking displays the Service Status page for the datacenter
Top Clusters / Top ESX Servers	
the top five clusters and ESX servers are respectively listed in order of current CPU usage	
CPU Trend 24h	<p>the CPU usage trend of the cluster or ESX server over the last 24 hours, expressed as a percentage of total available CPU cycles</p> <p>clicking opens a pop-up displaying a CPU Workload graph showing actual CPU usage over the last 24 hours</p>
Current %	the current percentage of total available CPU cycles consumed by the cluster or ESX server
the top five clusters and ESX servers are respectively listed in order of memory used	
Memory Trend 24h	<p>the memory usage trend of the cluster or ESX server over the last 24 hours, expressed as a percentage of total available memory</p> <p>clicking opens a pop-up displaying a Memory Workload graph showing actual memory usage over the last 24 hours</p>
Current GB	the current percentage of total available memory currently used by the cluster or ESX server
Top Resource Pools	
the top five resource pools are listed in order of current CPU usage	
CPU Trend 24h	<p>the CPU usage trend of the resource pool over the last 24 hours, expressed in raw GHz</p> <p>clicking opens a pop-up displaying a CPU Workload graph showing actual CPU usage over the last 24 hours</p>
Current GHz	the amount of CPU cycles, in raw GHz, currently consumed by the resource pool
the top five resource pools are listed in order of memory used	
Memory Trend 24h	<p>the memory usage trend of the resource pool over the last 24 hours, expressed in raw GBs</p> <p>clicking opens a pop-up displaying a Memory Workload graph showing actual memory usage over the last 24 hours</p>
Current GB	the amount of total memory, in raw GB, currently used by the resource pool

ESX Server and Cluster Quick Snapshot Contents

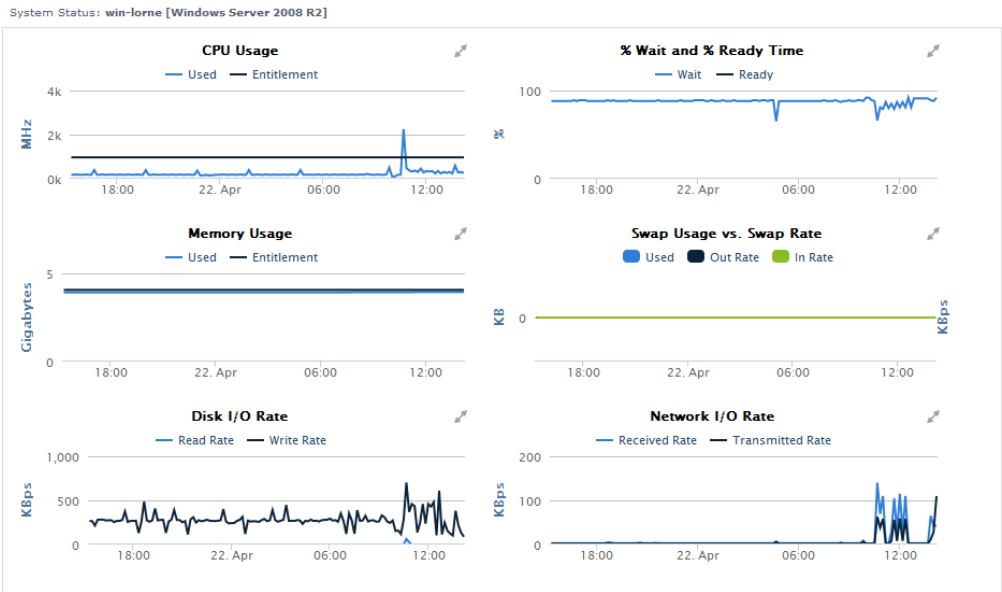


The following information is displayed in an ESX server or cluster Quick Snapshot:

CPU Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">total CPU capacity, in GHZ, which acts as the top line of the graphthe amount of CPU, in GHZ, used by VMsthe amount of CPU, in GHZ, reserved for VMs
VM state	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">total number of VMs powered ontotal number of suspended VMstotal number of VMs powered off
Memory Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">total memory capacity, in GB, which acts as the top line of the graphthe amount of memory, in GB, used by VMsthe amount of memory, in GB, reserved for VMs
Swap Usage vs. Swap Rate	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">total swap memory used by VMs, in MB or GB, whichever is more appropriatethe rate, in MBps, at which VMs are swapping memory from diskthe rate, in MBps, at which VMs are swapping memory to disk


Disk I/O Rate	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">the rate, in MBps, at which VMs are reading data from diskthe rate, in MBps, at which VMs are writing data to disk
Network I/O Rate	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">the rate, in MBps, at which VMs are receiving data from the networkthe rate, in MBps, at which VMs are transmitting data over the network
Top VMs by CPU	<ul style="list-style-type: none">top five VMs running on the ESX host, ordered by CPU usage, in MHzthe Details link is a shortcut to a top-10 workload graph showing the 24-hour trend
Top VMs by Memory	<ul style="list-style-type: none">top five VMs running on the ESX host, ordered by memory usage, in MBthe Details link is a shortcut to a top-10 workload graph, showing the 24-hour trend
Top VMs by Disk I/O	<ul style="list-style-type: none">top five VMs running on the ESX host, ordered by disk usage, in KBpsthe Details link is a shortcut to a top-10 workload graph, showing the 24-hour trend
Top VMs by Network I/O	<ul style="list-style-type: none">top five VMs running on the ESX host, ordered by network I/O throughput, in KBpsthe Details link is a shortcut to a top-10 workload graph, showing the 24-hour trend

Virtual Machine Quick Snapshot Contents



The following information is displayed in a virtual machine Quick Snapshot:

CPU Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none">CPU usage in MHzif allocated to the VM, the amount of reserved CPU cycles, in MHzif allocated to the VM, the amount of entitled CPU cycles, in MHz
-----------	---

% Wait and % Ready Time	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> wait time: amount of time during the interval, as a percentage, that the VM had scheduled CPU time, but gave nothing to process ready time: amount of time during the interval, as a percentage, that the VM was ready to process, but was not scheduled CPU time by the host <div>  It is possible to be presented with values that exceed 100%. The underlying data used in the Quick Snapshot are migrated from VMware vSphere via vSync; VMware conventions include percent-based metrics that can be greater than 100%. For example, refer to the VMware Technical Note, <i>Performance Counters</i>, at http://www.vmware.com/files/pdf/technote_PerformanceCounters.pdf. </div>
Memory Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> memory usage in MB or GB, whichever is more appropriate if allocated to the VM, the amount of reserved memory, in MB if allocated to the VM, the amount of entitled memory, in MB
Swap Usage vs. Swap Rate	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> total swap memory, in KB, used during the interval the rate, in KBps, at which the VM is swapping memory from disk the rate, in KBps, at which the VM is swapping memory to disk
Disk I/O Rate	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> the rate, in KBps, at which the VM is reading data from disk the rate, in KBps, at which the VM is writing data to disk
Network I/O Rate	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> the rate, in KBps, at which the VM is receiving data from the network the rate, in KBps, at which the VM is transmitting data over the network
Top 10 Processes	<p>note that this process list is only available if the VM Element's metrics are reported via WMI or the Uptime Infrastructure Monitor Agent</p> <ul style="list-style-type: none"> top 10 processes on the VM, ranked by CPU usage the Details link is a shortcut to the Detailed Process Information table for the VM Element

Viewing a Quick Snapshot for a vSphere Component

To display the Quick Snapshot page for any VMware vSphere component, do the following:

1. In the *My Infrastructure* panel, locate the VMware vCenter, ESX server, or VM Element whose Quick Snapshot you would like to view.
2. Click the gear icon beside the Element.
3. In the Element's *Configure* pop-menu, click *Graph Performance*.

Note that when you are viewing an Element's profile, you can always access its Quick Snapshot by clicking the *Graphing* tab, then clicking *Quick Snapshot* in the tree panel.

CPU Performance Graphs

Uptime Infrastructure Monitor uses the following graphs to chart the performance of one or more CPUs on a system:

CPU Workload

CPU Workload graphs help you gauge the demand placed on your computing resources, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers or clusters are consuming the most CPU cycles, or for an ESX server, which VMs are creating the largest workload.

CPU Workload Metrics

For CPU Workload graphs for a VMware vCenter server, the amount of MHz used can be generated for the following Element types:

- Datacenters
- Clusters
- ESX servers
- virtual machines
- resource pools
- vApps

For CPU Workload graphs for an ESX server, the amount of MHz used can be generated for the following Element types:

- virtual machines
- resource pools
- vApps

Graphing CPU Workload for a vCenter or ESX Server

To generate a CPU Workload graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the Tree panel, click *CPU Workload* (for a VMware vCenter server), or *Workload* in the *CPU* section (for an ESX server).
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#).
4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
5. In the first step, select a VMware vSphere component whose metrics you want to graph.
6. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
7. Click *Generate Graph*.
A pop-up window appears, displaying the CPU Workload graph you have configured.

CPU Usage

The CPU Usage graph shows how CPU resources are managed and used across all instances on an ESX server, or an individual VM.

CPU Usage Metrics

The following metrics are displayed in a CPU Usage graph:

Reservation	the minimum amount of CPU resources reserved for the VM, or all VMs running on an ESX host
Entitlement	the amount of CPU resources allocated to a VM, or all VMs running on an ESX host, on the assumption that all VMs on a host are completely busy, and load is perfectly balanced across hosts
Usage MHz	the amount of CPU used, in MHz, by the VM, or all powered on VMs running on the ESX host
Usage %	the amount of CPU resources, as a percentage of total available CPU, used by the cluster, ESX host, or VM
Usage % (agent)	the amount of CPU resources, as a percentage of total available CPU, used by a VM that is reporting detailed metrics via WMI or the Uptime Infrastructure Monitor agent

Graphing CPU Usage for an ESX Server or VM

To generate a CPU Usage graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose CPU usage you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the Tree panel, click *Usage* in the *CPU* section.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#).
4. Select the metrics you want to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the CPU Usage graph you have configured.

Multi-CPU Usage

The Multi-CPU Usage graph charts the performance statistics for ESX hosts or VMs with more than one CPU. These statistics indicate whether a system is effectively balancing tasks between CPUs, or if processes are forced off CPUs in certain circumstances. You can also use this graph to determine whether there are too many system interrupts that are using a CPU or that are overloading a CPU.

Multi-CPU Usage Metrics

The following metrics are displayed in a Multi-CPU Usage graph:

ESX Server	
Usage MHz	usage of total CPU resources by physical cores, as a percentage, on the ESX host
Virtual Machine	
Usage MHz	the amount of CPU usage, in MHz, by CPU, on the VM
% System Time	the amount, as a percentage, spent on system processes on each vCPU in the VM
% Wait Time	the amount, as a percentage, vCPUs on the VM spent in wait state

% Used	the amount, as a percentage, of used vCPU time
Virtual Machine with WMI- or Agent-Based Metrics	
% Total	total amount of user, privileged, and interrupt time, by vCPU
% Usr	the amount of time, as a percentage, of user processes, by vCPU
% Sys	the amount of time, as a percentage, of system process use, by vCPU
% Wait	the amount of time, as a percentage, of wait time use
Interrupts	number of CPU interrupts

i For the % Wait metric, it is possible that you see values that exceed 100%. The underlying data used for the CPU graphs are migrated from VMware vSphere via vSync; VMware conventions include percent-based metrics that can be greater than 100%. For example, refer to the VMware Technical Note, *Performance Counters*, at http://www.vmware.com/files/pdf/technote_PerformanceCounters.pdf.

Graphing Multi-CPU Usage for an ESX Server or VM

To generate a multi-CPU usage graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose multi-CPU workloads you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *CPU* section of the Tree panel, click *Multi-CPU Usage*.
3. Select the start and end dates and times for which the graph chart data, then click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. In the next step, select whether to graph the highest resource consumers, or specific VMs.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
5. Click *Generate Graph*.
A pop-up window appears, displaying the Multi-CPU Workload graph you have configured.

Wait and Ready Time

CPU wait time is the amount of time a VM is given scheduled time, but there is nothing to process, resulting in an idle CPU.

CPU ready time is the amount of time that a VM is ready for processing, but could not get scheduled to run on the physical CPU.

The Wait and Ready Time graph helps you determine if a guest is waiting too often on a host, or a host is waiting too often on a guest.

Wait and Ready Time Metrics

The following metrics are displayed in a Wait and Read Time graph:

% Wait	the amount of time during the interval, as a percentage, that the VM or all VMs on an ESX host, resource pool or vApp had scheduled CPU time, but gave nothing to process
% Ready	the amount of time during the interval, as a percentage, that the VM or all VMs on an ESX host, resource pool or vApp were ready to process, but were not scheduled CPU time by the host

i For these metrics, it is possible to be presented with values that exceed 100%. The underlying data used in these graphs are migrated from VMware vSphere via vSync; VMware conventions include percent-based metrics that can be greater than 100%. For example, refer to the VMware Technical Note, *Performance Counters*, at http://www.vmware.com/files/pdf/technote_PerformanceCounters.pdf.

Graphing Wait and Ready Time for an ESX Server or VM

To generate a percent-wait or percent-ready graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose wait and ready time you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *CPU* section of the Tree panel, click *Wait & Read Time*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
5. In the first step, select a VMware vSphere component whose metrics you want to graph.
6. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
7. Choose either the % *Wait*, or % *Ready* metric to include in the graph.
8. Click *Generate Graph*.
A pop-up window appears, displaying the CPU Wait and Ready Time graph you have configured.

Graphing Wait and Ready Time for an Agent-Based VM

To generate a percent-wait or percent-ready graph for a VM that receives metrics via WMI or the Uptime Infrastructure Monitor agent, do the following:

- 1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VM whose wait and ready time you want to graph, click its corresponding gear icon, then click *Graph Performance*.
- 2. In the Tree panel, click *Wait & Ready Time* in the *CPU* section.
- 3. Select the start and end dates and times for which the graph charts data. For more information, [Understanding Dates and Times](#)
- 4. Select one of both of the *% Ready* and *% Wait* metrics to include in the graph.
- 5. Click *Generate Graph*.
- 6. A pop-up window appears, displaying the graph you have configured.

Memory Usage Graphs

Uptime Infrastructure Monitor uses the following graphs to chart memory usage for a system:

Memory Workload

Memory Workload graphs help you gauge the demand placed on the memory allocated to your physical and virtual, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers or clusters are using the most memory, or for an ESX server, which VMs are the recipients of overcommitted granted memory.

Memory Workload Metrics

The following metrics can be used to generate a Memory Workload graph:

VMware vCenter Server	
Consumed	the amount of memory, in MB, that is used by or reserved for a VM, or all VMs that are part of a resource pool or vApp, or the total physical memory used by the datacenter, cluster, or ESX host
Ballooned	the amount of guest physical memory reclaimed through ballooning from a VM, or all VMs that are part of a resource pool or vApp, or the total memory balloon for all VMs in a datacenter, cluster, or ESX host
ESX Server	
Usage	the amount of memory used, as a percentage of total available memory, by all VMs on the ESX host, or in a resource pool or vApp
Active	the estimated amount of active memory in use, in KB, by VMs, or all VMs in a resource pool or vApp
Consumed	the amount of memory, in MB, that issued by or reserved for a VM, or all the VMs that are part of a resource pool or vApp
Granted	the amount of physical memory allocated to machine memory for a VM, or all the VMs that are part of a resource pool or vApp
Ballooned	the amount of guest physical memory reclaimed through ballooning from a VM, or all VMs that are part of a resource pool or vApp

Graphing Memory Workload for a vCenter or ESX Server

To generate a Memory Workload graph, do the following:

- 1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click *Graph Performance*.
- 2. In the Tree panel, click *Memory Workload* (for a VMware vCenter server), or *Workload* in the *Memory* section (for an ESX server).
- 3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
- 4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
- 5. In the first step, select a VMware vSphere component whose metrics you want to graph.
- 6. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
- 7. Select the metric to include in the graph.
- 8. Click *Generate Graph*.
A pop-up window appears, displaying the Memory Workload graph you have configured.

Memory Usage

The Memory Usage graph displays the amount of memory used on an ESX server, or a VM.

Memory Usage Metrics

The following metrics are displayed in a Memory Usage graph:

ESX Server Metrics

Reservation	the minimum amount of memory reserved for VMs running on an ESX host
Consumed	the amount of memory, in MB, that is used by or reserved for VMs running on an ESX host
Active	the estimated amount of active memory in use, in MB, by VMs running on an ESX host
Granted	the amount of memory allocated to VMs running on an ESX host, on the assumption that all are completely busy, and load is perfectly balanced across hosts
VM metrics	
Reservation	the minimum amount of memory reserved for the VM
Entitlement	the amount of memory, in MB, allocated to a VM, on the assumption that all VMs on a host are completely busy, and load is perfectly balanced across hosts
Consumed	the amount of memory, in MB, that is used by or reserved for a VM
Active	the estimated amount of active memory in use, in MB, by the VMs
Granted	the amount of physical memory allocated to machine memory for a VM
Used MB (Agent)	the amount of memory used by a VM that is reporting detailed metrics via WMI or the Uptime Infrastructure Monitor agent

Graphing Memory Usage for an ESX server or VM

To generate a Memory Usage graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose memory usage you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the Tree panel, click *Usage* in the *Memory* section.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the metrics you want to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the Memory Usage graph you have configured.

Memory Profile

The Memory Profile graph displays trends in allocated memory for VMs on an ESX server. Use this graph to ensure VMware vSphere's memory resource management techniques are working correctly.

Memory Profile Metrics

The following metrics are displayed in a Memory Profile graph:

ESX Server Metrics	
Shared Common	the amount of machine memory that is shared by all powered on VMs on the ESX host
Swapped	the total amount of guest physical memory swapped out to swap files of all VMs running on the ESX host
Ballooned	the total memory balloon of all powered on VMs
Overhead	the total amount of machine memory used to run all VMs running on the ESX host
VM metrics	
Shared	the amount of guest physical memory shared with other VMs
Swapped	the amount of guest physical memory swapped out to the VMs swap file
Ballooned	the amount of guest physical memory reclaimed through ballooning from a VM
Overhead	the amount of machine memory used to run the VM

Graphing a Memory Profile for an ESX Server or VM

To generate a Memory Profile graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose Memory Profile graph you want to create, click its corresponding gear icon, then click *Graph Performance*.
2. In the Tree panel, click *Profile* in the *Memory* section.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the metrics you want to include in the graph.

- Click *Generate Graph*.

A pop-up window appears, displaying the Memory Profile graph you have configured.

Memory Guest Paging

This graph indicates whether a VM guest is short of memory.

Guest Paging Metrics

The following metrics are displayed in a Guest Paging graph:

Page In/sec	number of memory pages transferred from disk to RAM
Page Out/sec	number of memory pages transferred from RAM to disk
Page Frees/sec	number of memory pages removed from the swap to and from RAM and secondary storage

Graphing Guest Memory Paging Rates on an Agent-Based VM

To generate a guest memory paging rate graph for a VM that receives metrics via WMI or from the Uptime Infrastructure Monitor agent, do the following:

- Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VM whose paging rates you want to create, click its corresponding gear icon, then click *Graph Performance*.
- In the *Memory* section of the Tree panel, click *Guest Paging*.
- Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
- Select the paging metrics you want to include in the graph.
- Click *Generate Graph*.
A pop-up window appears, displaying the Memory Profile graph you have configured.

Memory Swap

The Memory Swap graph charts the amount of memory held in swap files on ESX servers, as well as the rates at which data is swapping between memory and secondary storage.

Memory Swap Metrics

The following metrics are displayed in a Memory Swap graph:

ESX Server Metrics	
Swapped	total guest physical memory, in MB, swapped out to swap files of VMs in resource pools, vApps, or on the ESX host
Swap Total Rate	the total swap in rate and swap out rate, in MB, for VMs in resource pools, vApps, or on the ESX host
VM Metrics	
Swap Used	the amount of guest physical memory, in MB, swapped out to the VM's swap file
Swap In Rate	the rate at which memory, in MB, is swapped into active memory
Swap Out Rate	the rate at which memory, in MB, is swapped out to disk
Swap Used MB (Agent)	the amount of guest physical memory, in MB, swapped out to the VM's swap file, as reported via WMI or the Uptime Infrastructure Monitor agent

Creating a Memory Swap Graph for an ESX server or VM

To generate a percent-wait or percent-ready graph, do the following:

- Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose memory swap statistics you want to graph, click its corresponding gear icon, then click *Graph Performance*.
- In the *Memory* section of the Tree panel, click *Swap*.
- Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
- Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
- In the first step, select a VMware vSphere component whose metrics you want to graph.
- In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
- Select the metric you wish to include in the graph.
- Click *Generate Graph*.
A pop-up window appears, displaying the Memory Swap graph you have configured.

Creating a Memory Swap Graph for an Agent-Based VM

To generate a Memory Swap graph for a VM that receives metrics via WMI or the Uptime Infrastructure Monitor agent, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VM whose memory swap information you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the Tree panel, click *Swap* in the *Memory* section.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the metrics you want to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the Memory Swap graph you have configured.

Process Graphs

Uptime Infrastructure Monitor uses the following graphs to report system process information for VMware vSphere-monitored VMs whose metrics are reported via WMI or with the Uptime Infrastructure Monitor agent:

Process Workload

The Process Workload graph determines the demand that network and local services are putting on a VM whose metrics are reported via WMI or with the Uptime Infrastructure Monitor agent. The graph charts an aggregate amount of performance information for a given user, group, or process.

Process Workload Metrics

The following metrics are displayed in a Process Workload graph:

% CPU	the percentage of the VM's CPU time that is taken up by a user, group, or process
Total Memory Size	<ul style="list-style-type: none">the amount of the page file and virtual memory that is taken up by a user, group, or processon Windows VMs, Memory Size is called Virtual Bytes.
RSS Memory Size	<ul style="list-style-type: none">the Run Set Size, which is the amount of physical memory that is used by a user, group, or processon Windows systems, RSS is called Working Set

Graphing Process Workloads for an Agent-Based VM

To generate a Process Workload graph for a VM that receives metrics via WMI or the Uptime Infrastructure Monitor agent, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VM whose process workload you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Process* section of the Tree panel, click *Workload*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. Click one of the *Quick Graphs* options to display a pre-configured workload graph in a pop-up window, or skip this step to manually configure a graph.
5. In the first step, select an object whose workload you want to graph.
6. In the next step, select whether to graph the highest resource consumers, or specific processes running on the VM.
If you select *Specific*, a process selection dialog appears, requiring you to build a list.
7. Select one or more process metrics to include in the graph.
8. Click *Generate Graph*.
A pop-up window appears, displaying the Process Workload graph you have configured.

Detailed Process Information

Detailed process information provides insight into how various user and system processes are consuming system resources.

Detailed Process Information Statistics

The following information is displayed in a Detailed Process Information table:

Process	the name of the process, which is taken from its executed path name
PID	the number that identifies the process
PPID	<ul style="list-style-type: none">the number that identifies the parent process. The PPID can help identify possible relationships between processeson Windows systems, the PPID is called the Creating Process ID

UID	<ul style="list-style-type: none"> the ID of the user or account that is consuming CPU time on Windows systems, the UID is called the Owner
GID	<ul style="list-style-type: none"> the ID of the group that is consuming CPU time on Windows systems, the GID is called the Group Name
Memory Used	<ul style="list-style-type: none"> the amount of memory, expressed as a percentage of total available memory, consumed by a process on Windows systems, Memory Used is called Virtual Bytes the Memory Used value can be misleading because shared memory between processes is counted multiple times (e.g., if five Oracle processes are using 10% of available memory, this does not indicate that Oracle is consuming 50% of system memory)
RSS	<ul style="list-style-type: none"> Run Set Size - the amount of physical memory that is used on Windows systems, RSS is called the Working Set
CPU %	<ul style="list-style-type: none"> the percentage of the CPU time used by the process, calculated by dividing total used CPU Time by the process' running time; if applicable, the result is further divided by the number of CPUs for the Element on which the process is running. on Windows systems, the CPU % is called % Processor Time
Mem %	the percentage of the memory used by the process
Runtime	the total runtime of the process
Children Runtime	the runtime of all child processes
Start Time	the time at which the process started, which can be used to determine its lifetime

Listing Detailed Process Information for an Agent-Based VM

To list Detailed Process Information for a VM Element, do the following:

- Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VM whose process information you want to display, click its corresponding gear icon, then click *Graph Performance*.
- In the Tree panel, click *Detailed Proc Info* in the *Process* section.
- Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
- Click *Display Process Information*.
A pop-up window appears, displaying the process information for the specified time interval.

Number of Processes

The Number of Processes graph displays the VM's process activity with the following details:

- number of processes
- running processes
- blocked processes
- waiting processes

Analyzing processes helps you determine whether there is enough CPU capacity for the processes that are run on a system. If the size of the blocked or waiting queue is disproportionate to the running queue, then either the system does not have enough CPUs or is too I/O bound.

A blocked process signals a disk bottleneck. If the number of blocked processes approaches or exceeds the number of processes in the run queue, you should tune the disk subsystem. Whenever there are any blocked processes, all CPU idle time is treated as wait for I/O time. If database batch jobs are running on the system that is monitored, there always are some blocked processes. However, you can increase the throughput of batch jobs by removing disk bottlenecks.

Process Count Statistics

The following WMI- or agent-based information is displayed in a Process Count table:

Running	the number of processes that are currently running on a system, as reported by the system kernel
Blocked	the number of processes that are unable to complete, usually due to I/O operation
Waiting	the number of processes awaiting execution
Total Processes	the total number of running, blocked, and waiting processes

Graphing Process Counts for an Agent-Based VM

To generate a Number of Processes graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VM whose process counts you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Process* section of the Tree panel, click *Number of Processes*.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the metrics you want to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the process counts graph you have configured.

Network Graphs

Uptime Infrastructure Monitor uses the following graphs to chart network performance and usage:

Network Workload

Network Workload graphs help you gauge the demand placed on the network, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers or clusters are transmitting and receiving the most data, or for an ESX server, which VMs' virtual NICs are the busiest.

Network Workload Metrics

For Network Workload graphs for a VMware vCenter server, the sum of KBps transmitted and received through NICs and vNICs can be generated for the following Element types:

- Datacenters
- Clusters
- ESX servers
- virtual machines
- resource pools
- vApps

For Network Workload graphs for an ESX server, the sum of KBps transmitted and received through virtual NICs can be generated for the following Element types:

- virtual machines
- resource pools
- vApps

Graphing Network Workload for a vCenter or ESX Server

To generate a Network Workload graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click *Graph Performance* .
2. In the Tree panel, click *Network Workload* (for a VMware vCenter server), or *Workload* in the *Network* section (for an ESX server).
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
5. In the first step, select a VMware vSphere component whose metrics you want to graph.
6. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
7. Click *Generate Graph*.
A pop-up window appears, displaying the Network Workload graph you have configured.

Network I/O

The I/O graph displays rates and total sums of data that are moving in and out of physical and virtual network interfaces over a specified time period. It can be used to find where I/O bottlenecks are occurring, and whether there are any bursts in network traffic.

Network I/O Metrics

The following metrics are displayed in a Network I/O graph:

ESX Server Metrics	
Total Rate	the sum of data transmitted and data received, in KBps, through the physical NICs on the ESX host
Transmit Rate	the average rate, in KBps, at which data was transmitted through each physical NIC on the ESX host
Receive Rate	the average rate, in KBps, at which data was received through each physical NIC on the ESX host
VM Metrics	

Total Rate	the sum of data transmitted and data received, in KBps, through the vNICs on the VMs on the ESX host
Transmit Rate	the average rate, in KBps, at which data was transmitted through each vNIC on the VM
Receive Rate	the average rate, in KBps, at which data was received through each vNIC on the VM
Total Rate (Agent)	the sum of data transmitted and data received, in KBps, through the guest NICs on the VMs on the ESX host, as reported via WMI or the Uptime Infrastructure Monitor agent
Send Rate (Agent)	the average rate, in KBps, at which data was transmitted through the guest NICs on the VMs on the ESX host, as reported via WMI or the Uptime Infrastructure Monitor agent
Receive Rate (Agent)	the average rate, in KBps, at which data was received through the guest NICs on the VMs on the ESX host, as reported via WMI or the Uptime Infrastructure Monitor agent

Graphing Network I/O Rates for an ESX Server or VM

To generate a Network I/O graph for an ESX server or VM, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose network rates you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Network* section of the Tree panel, click *I/O*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. If you are graphing network I/O for a VM that receives metrics via WMI or from the Uptime Infrastructure Monitor agent, select whether to graph network rates through virtual NICs or guest NICs.
5. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
6. Select the network metric to include in the graph.
7. Click *Generate Graph*.
8. A pop-up window appears, displaying the network I/O rate graph you have configured.

Network Errors

The Network Errors graph displays the rate of errors with physical NICs on ESX servers, and vNICs on VMs.

The most common types of errors for physical connections include collisions in a hubbed environment or the presence of full-duplex handshake errors between a system and a switch. Additionally, the following communication line problems can cause network errors:

- excessive noise
- cabling problems
- problems with backbone connections

Network Error Metrics

The following metrics are displayed in a Network Errors graph

ESX Server Metrics	
Total Dropped/min	the total number of dropped packets through the physical NICs on the ESX host, per minute, during the time period
Transmit Dropped /min	the number of packets outbound through the physical NICs on the ESX host that were dropped per minute during the time period
Receive Dropped/min	the number of packets inbound through the physical NICs on the ESX host that were dropped per minute during the time period
VM metrics	
Total Errors (Agent)	the sum of in errors and out errors through a guest NIC on the VMs on the ESX host
In Errors (Agent)	the number of packets received, but unable to be decoded due to a missing header or trailer
Out Errors (Agent)	the number of packets that were not sent, due to problems transmitting the packet or formatting the packet for transmission
Collisions (Agent)	the simultaneous presence of signals from two nodes on the network (e.g., from two nodes beginning to transmit over a network at the same time)
TCP Retransmits (Agent)	the number of TCP retransmits that occurred during the time period, due to lost or broken packets

Graphing Network Error Rates for an ESX Server or VM

To generate a network error graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose network error rate you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Network* section of the Tree panel, click *Errors*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
5. Select the error metric to include in the graph.
6. Click *Generate Graph*.
A pop-up window appears, displaying the network error rate graph you have configured.

Disk Graphs

Uptime Infrastructure Monitor uses the following graphs to chart disk details, performance, and usage:

Disk Workload

Disk Workload graphs help you gauge the demand placed on your physical or guest storage, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers are experiencing the heaviest disk read/write activity, or for an ESX server, which VMs are writing the most to disk, prompting further investigation into VM read and write rates.

Disk Workload Metrics

For Disk Workload graphs for a VMware vCenter server, the aggregated disk I/O rate, in KBps, can be generated for the following Element types:

- Datacenters
- Clusters
- ESX servers
- virtual machines
- resource pools
- vApps

For Disk Workload graphs for an ESX server, the aggregated disk I/O rate, in KBps, can be generated for the following Element types:

- virtual machines
- resource pools
- vApps

Graphing Disk Workload for a vCenter or ESX Server

To generate a Disk Workload graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the Tree panel, click *Disk Workload* (for a VMware vCenter server), or *Workload* in the *Disk* section (for an ESX server).
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
5. In the first step, select a VMware vSphere component whose metrics you want to graph.
6. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
7. Click *Generate Graph*.
A pop-up window appears, displaying the Disk Workload graph you have configured.

Disk I/O

The Disk I/O charts the rates at which physical disks on ESX servers, and virtual disks on VMs are reading and writing data. VMs that are receiving metrics via WMI or the Uptime Infrastructure Monitor agent report deeper metrics including wait times and data transfer speeds. These metrics can help you determine which disks are busiest.

Disk I/O Metrics

The following metrics are displayed in a Disk I/O graph:

ESX server and VM Element Metrics	
Total RW Rate	the total read rate and write rate for virtual disks on VM, or LUNs on the ESX host
Read Rate	the rate, in KBps, at which data is read from each virtual disk on the VM, or each LUN on the ESX host
Write Rate	the rate, in KBps, at which data is written to each virtual disk on the VM, or each LUN on the ESX host
Total RW/min	the total number of read requests and write request from each virtual disk on the VM, or LUN on the ESX host
Read Requests/min	the number of times data was read from each virtual disk on the VM, or LUN on the ESX host
Write Requests/min	the number of times data was written to each virtual disk on the VM, or LUN on the ESX host

VM w/ agent Metrics on Guest Disks	
% Busy (Agent)	the percentage of guest disk capacity in use
Avg Queue Requests (Agent)	the average number of processes that are waiting to access the guest disk
Transfers/sec (Agent)	the total number of disk transfer requests processed per second
Blocks/sec (Agent)	the total number of data blocks written to and from the disk
Avg Service Time (Agent)	the average time, in milliseconds, required to perform a task
Avg Wait Time (Agent)	the average time, in milliseconds, that a transaction is waiting in a queue (note that the wait time is directly proportional to the length of the queue)

Graphing Disk I/O Rates for an ESX Server or VM

To generate a Disk I/O graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose disk I/O rate you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Disk* section of the Tree panel, click *I/O*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
5. In the first step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
6. Select the disk I/O metric to include in the graph.
7. Click *Generate Graph*.
A pop-up window appears, displaying the disk I/O graph you have configured.

Graphing Disk I/O Rates for an Agent-Based VM

To generate a Disk I/O graph for a VM that receives metrics via WMI or the Uptime Infrastructure Monitor agent, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose disk I/O rates you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Disk* section of the Tree panel, click *I/O*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. Select whether to graph a disk I/O rate on the VM disk, or a guest disk.
5. In the next step, select whether to graph the highest resource consumers, or specific VMs.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
6. Select the disk I/O metric to include in the graph.
7. Click *Generate Graph*.
A pop-up window appears, displaying the CPU Wait and Ready Time graph you have configured.

Disk Errors

The Disk Errors graph charts the error rates on physical disks on ESX servers, or virtual disks on VMs. High error rates indicates performance issues with the underlying hardware.

Disk Error Metrics

The following metrics are displayed in a Disk Errors graph:

Bus Resets/min	the number of SCSI bus reset commands issued on the virtual disk on the VM, or host disk
Commands Aborted/min	the number of SCSI commands on the virtual disk on the VM, or host disk, that were aborted

Graphing Disk Error Rates for an ESX Server or VM

To generate a disk error rate graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose disk error rates you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Disk* section of the Tree panel, click *Errors*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
4. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
5. If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
6. Select the disk error metric to include in the graph.

7. Click *Generate Graph*.

A pop-up window appears, displaying the disk error rate graph you have configured.

Disk Latency

The Disk Latency graph indicates the health and performance of physical storage on an ESX server.

Disk Latency Metrics

The following metrics are displayed in a Disk Latency graph:

ESX Server Metrics	
Device Latency	the average amount of time, in milliseconds, required to complete a SCSI command from the physical device on the ESX host
Device Read Latency	the average amount of time, in milliseconds, required to complete reading from the physical device on the ESX host
Device Write Latency	the average amount of time, in milliseconds, required to complete writing to the physical device on the ESX host
Kernal Latency	the average amount of time, in milliseconds, VMKernel spent processing each SCSI command on the ESX host
Kernal Read Latency	the average amount of time, in milliseconds, VMKernel spent processing each SCSI read command on the ESX host
Kernal Write Latency	the average amount of time, in milliseconds, VMKernel spent processing each SCSI write command on the ESX host
Queue Latency	the average amount of time, in milliseconds, spent in the VMKernel queue, per SCSI command
Queue Read Latency	the average amount of time, in milliseconds, taken per SCSI read command
Queue Write Latency	the average amount of time, in milliseconds, taken per SCSI write command
Command Latency	the average amount of time, in milliseconds, to process a SCSI command issued by the guest OS to the VM
Command Read Latency	the average amount of time, in milliseconds, to process a read command issued from the guest OS to the VM
Command Write Latency	the average amount of time, in milliseconds, to process a write command issued from the guest OS to the VM

Graphing Disk Latency for an ESX Server

To generate a percent-wait or percent-ready graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server whose disk latency you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Disk* section of the Tree panel, click *Latency*.
3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#).
4. Click the *Quick Graph* option to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
5. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
6. Select the latency metric to include in the graph.
7. Click *Generate Graph*.
A pop-up window appears, displaying the disk latency graph you have configured.

Disk Storage Capacity

A Disk Storage Capacity graph charts the amount of total and used space, in gigabytes, on an Element's disk. This includes VMware vCenter servers, ESX servers, and VMs. For VMs that are receiving metrics from the Uptime Infrastructure Monitor agent, the capacity of the file systems are also available.

The current capacity and usage for each datastore is displayed, as well as the week-over-week change, and estimated time to fill. You can select a datastore to graph to see datastore usage trends over time.

Disk Storage Capacity Details

The following metrics are displayed in a Disk Storage Capacity graph:

Total Capacity	the total amount of space available on the system or VM
Used	the amount of space currently used
Free	the amount of free space
Change Per Week	a positive or negative value indicating the amount of free space compared to one week ago
Time to Fill	based on the changes during the last week, the estimated amount of time before the disk is filled

Graphing Storage Capacities for vSphere Components

To generate a Disk Storage graph for a VMware vCenter Server, an ESX server, or a VM, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the Element whose storage capacity you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Disk* section of the Tree panel (for an ESX server or VM Element), or in the *Datastores* section (for a VMware vCenter Element), click *Storage Capacity*.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the datastores to include in the graph.
If the Element graphed is a WML- or agent-based VM, you can also select the guest file system to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the storage capacity graph you have configured.

Disk Storage Profile

The Disk Storage Profile gives a breakdown of all datastores for an ESX server or VM. The current capacity and usage for each datastore is displayed. You can select a datastore to graph to see datastore usage trends over time.

Disk Storage Profile Details

The following information is displayed in a Storage Profile graph:

- Total Capacity
- vDisk Usage
- Snapshot Usage
- Swap Usage
- Other Usage

Graphing a Storage Profile for an ESX Server or VM

To generate a storage profile for an ESX server or VM, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or VM whose storage profile you want to display, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Disk* section of the Tree panel, click *Storage Profile*.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the datastore you want to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the process counts graph you have configured.

VMware Management Graphs

Uptime Infrastructure Monitor uses the following graphs to track and manage VMs:

VM Instance Motion

The VMware VMotion tool enables you to move ESX instances from one server to another without any downtime or loss of data. For example, you might use VMotion to move an instance to newer and faster hardware, or to temporarily relocate the instance while performing a hardware upgrade.

The Instance Motion graph enables you to keep track of a moving VMware instance. For a given ESX instance, the graph charts which systems it is running on over a given time range.

This graph can be generated when you are viewing a VMware vCenter server, ESX server, or a VM; the topological level at which you begin to configure the graph determines which instances are available to graph.

Generating an Instance Motion Graph

To generate an Instance Motion graph, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server or ESX server whose VMs' instance motion you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Metrics* section of the Tree panel, click *Instance Motion*.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. From the drop-down menu, select a VM.
5. Click *Generate Graph*.

Power Consumption

To assist virtualization initiatives that are meant to save power costs, or to gauge the efficiency of existing virtual datacenters, power usage by watts is available as a graphing metric.

The Power Consumption graph can tell you how much power your ESX hosts are consuming by datacenter, cluster, or individual server. Additionally, you can graph the power usage of individual VMs.

You can graph power usage levels at the VMware vCenter level in order to assess current load distribution, or verify VMware’s automated distributive resource balancing is functioning.

Graphing Power Consumption for a vCenter Server

To generate a power consumption graph for a VMware vCenter server, do the following:

- 1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server whose power consumption you want to graph, click its corresponding gear icon, then click *Graph Performance*.
- 2. In the *Metrics* section of the Tree panel, click *Power Consumption*.
- 3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
- 4. Select the logical group whose power consumption you want to graph.
- 5. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
- 6. Click *Generate Graph*.
A pop-up window appears, displaying the power consumption graph you have configured.

Graphing Power Consumption for an ESX Server

To generate a power consumption graph for an ESX server, do the following:

- 1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server or cluster whose power consumption you want to graph, click its corresponding gear icon, then click *Graph Performance*.
- 2. In the *Metrics* section of the Tree panel, click *Power Consumption*.
- 3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
- 4. Click *Generate Graph*.
A pop-up window appears, displaying the power consumption of the ESX server over the specified time period.

Power States

Power States graphs allow you to graph VM activity on an ESX host, or ESX activity on a VMware vCenter server.

The power state graphs help you manage both available computing resources within your VMware vSphere clusters and datacenters, as well as power consumption in your physical datacenters.

VMware vCenter Element Power States

The following states are displayed in a Power States graph:

vCenter Server Metrics	
VMs Powered On	the virtual machine is powered on
ESX Powered On	the host is powered on
ESX Powered Off	the host was powered off by an administrator through the VMware vSphere Client
ESX Unknown	as is the case in the VMware vSphere Client, a host that is in an unknown state is assumed to be powered off by an administrator
ESX In Maintenance	the host was put in maintenance mode by an administrator
ESX In Standby	the host was put in standby mode either explicitly by an administrator, or automatically by vSphere Distributed Power Management (DPM)
ESX Server Metrics	
VMs Powered On	the virtual machine is powered on
VMs Powered Off	the virtual machine is powered off
VMs Suspended	the virtual machine is not running, but a snapshot of its running applications and processes is retained.

Graphing Power States for a vCenter Server

To generate a Power States graph, do the following:

- 1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the VMware vCenter server whose ESX server or VM power states you want to graph, click its corresponding gear icon, then click *Graph Performance*.
- 2. In the *Metrics* section of the Tree panel, click *Power States*.
- 3. Select the start and end dates and times for which the graph charts data, and click *Apply Date and Time*. For more information, see [Understanding Dates and Times](#)
- 4. Click one of the *Quick Graphs* options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.
- 5. In the first step, select whether to graph power states by datacenter or cluster.

6. In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.
If you select *Specific*, an Element selection dialog appears, requiring you to build a list.
7. Select the power states you want to include in the graph.
8. Click *Generate Graph*.
A pop-up window appears, displaying the CPU Wait and Ready Time graph you have configured.

Graphing VM Power States for an ESX Server

To generate a power status graph for an ESX server's VMs, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the ESX server whose power states you want to graph, click its corresponding gear icon, then click *Graph Performance*.
2. In the *Metrics* section of the Tree panel, click *Power States*.
3. Select the start and end dates and times for which the graph charts data. For more information, see [Understanding Dates and Times](#)
4. Select the power states you want to include in the graph.
5. Click *Generate Graph*.
A pop-up window appears, displaying the power states graph you have configured.