

Using Graphs

Overview

up.time can display the performance and availability statistics for the systems that you are monitoring in a graph. You can use the graphs to collect and display information for Elements, services, and configurations.

You have different graphing options depending on the operating system that is running on a host. The metrics that *up.time* agents capture and return to the Monitoring Station differ from operating system to operating system.



Note - If a graph is not available in the Tree panel for a given host, it is because the host does not provide the metric that the graph requires. Also, if you add a node or a virtual node, such as a router or IP address, you can only see them in the Config and the Services tabs as other metrics such as CPU and disk usage are not available from the node.

UNIX vs. Windows Performance Monitoring

In most cases, you can interpret performance data from different platforms - such as Windows, UNIX and Linux - in similar ways. When the interpretation of the data is different, the *up.time* interface displays operating system-specific information - such as the performance counters being used - as necessary.

Viewing the Status of a System

You can view the status of a system in your environment using a Quick Snapshot. The Quick Snapshot summarizes key hardware and process information for a system for the last 24 hours. If there is not 24 hours worth of data available, then *up.time* uses data from as far back as possible to generate charts.

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 scan the information provided in the charts and tables, and begin further investigation. (For example, if you notice problem while viewing the Quick Snapshot, you can generate a report to obtain more information about the problem.)

The Quick Snapshot contains the following information:

System Status Charts	Top 10 Processes	File System Statistics
<ul style="list-style-type: none">• CPU Usage• Memory Usage• Disk I/O (transfers/sec)• Network I/O rates• Outages• Disk usage	<ul style="list-style-type: none">• Process name• Process ID• % CPU usage• % memory usage	<ul style="list-style-type: none">• Device• Mount• Size• Used space• Available space• % used



Note - The components that comprise a Quick Snapshot depend on the type of Element in view. Monitored Elements typically provide the aforementioned information. For information about Quick Snapshots for VMware vSphere objects, or a network device, see [Viewing the Status of a vSphere Element](#), and [Viewing the Status of a Network Device](#), respectively.

Viewing a Quick Snapshot

In the *Global Scan* panel, click the name of the system whose information you want to graph. The Quick Snapshot is displayed by default.

Generally speaking, you can access a Quick Snapshot for an Element by clicking the *Graphing* tab, then clicking *Quick Snapshot* in the Tree panel.

Monitoring CPU Performance

up.time uses the following graphs to chart the performance of one or more CPUs on a system:

These graphs use the same input criteria, but they return different data.

Usage (% busy)

The Usage (% Busy) graph charts the percentage of a system's CPU resources that are being used over a period that you specify. This graph displays three components of CPU time: user, system, and wait I/O. Taken together, these components display the total amount of CPU usage. On a system with multiple CPUs, the numbers are averages across all CPUs.

CPU Usage in Windows

The key CPU usage metric in Windows is *% Usr Time*, which monitors the amount of time the CPU spends processing a thread that is not idle. If usage is consistently at 80% to 90%, you may need to upgrade the CPU or add more processors.

You should monitor a separate instance of this counter for each processor on systems with multiple CPUs. The value returned by the counter represents the sum of processor time on a specific processor.

Note - To determine the average for all processors, monitor the System: %Total Processor Time metric.

Optionally, you can monitor the following metrics:

Processor: % Privileged Time

The percentage of time that the CPU spends executing Windows kernel commands. If this metric is consistently high you should consider using a faster or more efficient disk subsystem.

Processor: %User Time

The percentage of time that the CPU spends executing user processes.

Processor: % Interrupt Time

The time that the CPU spends managing hardware requests. This metric enables you to determine the level of device activity.

System: Processor Queue Length

The number of threads that are waiting for processor time.

CPU Usage in UNIX and Linux

In UNIX and Linux, *up.time* graphs the following metrics:

User Time per CPU

The amount of time that the CPU spends in user mode. During user time, the CPU is processing application threads or threads that support tasks which are specific to applications.

System Time per CPU

The amount of time that the kernel spends processing system calls. If all of the CPU time is spent in system time, there could be a problem with the system kernel, or the system is spending too much time processing I/O interrupts.

Wait I/O Time per CPU

The amount of waiting time that a runnable process for a device takes to perform an I/O operation. Wait I/O problems are frequently related to problems with a disk.

Run Queue Length

The Run Queue Length graph counts the number of processes that are not currently running, and which are waiting to be served by the CPU. If several processes are trying to use CPU time, you might need to install a faster processor, or add another processor if you are using a multiprocessor system.

A long queue increases the time that a request waits before it is carried out by the CPU. However, it does not affect the time that is required to process each request once the CPU starts carrying out the request.

up.time counts the number of processes that are waiting in queue at a particular point in time. If the run queue or load average is greater than four times the number of CPUs, then processes must wait too long for the CPU to process the requests.

Run Queue Occupancy

The Run Queue Occupancy graph charts the percentage of time that one or more services or processes are waiting to be served by the CPU.

If the run queue occupancy is close to 100% and the run queue length is considered low, the CPU is not necessarily overloaded. While there may always be services waiting to be processed, the CPU may still be able to quickly process them.

If the run queue occupancy is high and the queue is long, then there is a capacity problem. However, a system should always have some idle time. Having consistently low idle time usually means that your system is working near its maximum capacity.

Generating a CPU Performance Graph

To generate a CPU performance graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - [Usage \(% busy\)](#)
 - [Run Queue Length](#)
 - [Run Queue Occupancy](#)
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph*.

Multi-CPU Usage

The Multi-CPU Usage graph charts the performance statistics for systems with more than one CPU. These statistics indicate whether or not a system is effectively balancing tasks between CPUs, or if processes are being forced off CPUs in certain circumstances. You can also use this graph to determine whether or not there are too many system interrupts that are using a CPU or that are overloading a CPU.

up.time can also collect and chart information for systems running Net-SNMP that have two or more CPUs. However, if the system was recently added to *up.time*, or if the *HOST-RESOURCES* MIB - which is used to collect data from the system - has not been properly installed and configured, *up.time* cannot collect CPU performance data. You must either wait until *up.time* is able to collect performance data, or check whether or not the *HOST-RESOURCES* MIB is properly installed and configured on the system that is being monitored.



Note - If there is only one CPU on the system, the following message is displayed instead of a graph:

This system is currently listed as only having one CPU.

Generating a Multi-CPU Usage Graph

To generate a Multi-CPU Usage graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the Tree panel, click the *Graphing* tab.
3. Click *Multi-CPU Usage*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click one of the following options:
 - **User %**
The percentage of CPU user processes that are in use. For Windows systems, this option is *% User Time*.
 - **System %**
The percentage of CPU kernel processes that are in use. For Windows systems, this option is *% System Time*.
 - **% Privileged Time**
On Windows systems, the percentage of time that the CPU spends executing kernel commands.
 - **Wait I/O %**
The percentage of time that a process which can be run must wait for a device to perform an I/O operation.
 - **SMTX**
The number of read or write locks that a thread was not able to acquire on the first attempt, as reported by the *mpstat* command.



Note - While it is trying to acquire locks, the thread is active but is not performing any tasks.

- **XCAL**
The number of interprocess cross-calls.

In a multi-processor environment, one processor sends cross-calls to another processor to get that processor to do work. Cross-calls can also be used to ensure consistency in virtual memory. Heavy file system activity - such as NFS - can result in a high number of cross-calls.
 - **Interrupts**
The number of CPU interrupts. For Windows systems, this option is
 - **% Interrupt Time**
Interrupts are a mechanism that a device uses to signal to the kernel that it needs attention, and that immediate processing is required on its behalf.
 - **Interrupts/sec**
On Windows systems, rate at which CPU handles interrupts from applications or hardware each second. If the value for Interrupts/sec is high, there could be problems with the hardware on the system.
 - **Total %**
On Solaris systems, the total amount of User %, System %, and Wait I/O %.
On Windows systems, this option is *% Total* and is the total amount of % User Time, % Privileged Time, and % Interrupt Time.
6. Select the CPUs to graph from the *Choose CPUs to graph* list.
 7. Click *Generate Graph*.

Graphing Memory Usage

up.time uses the following graphs to chart memory usage on a system:

These graphs use the same input criteria, but they return different data.

Used

This graph charts the amount of memory being used on a system. Used memory is the amount of physical memory occupied by the operating system, system library files, and applications.

Cache Hit Rate

This graph indicates how effectively buffers are controlling the flow of data between disks and the system.

CPU cache is a small store of free memory that is used by frequently-performed tasks for repeated fast disk access. The cache hit rate measures how often the system accesses the CPU cache.

The cache hit rate calculations are taken from the following metrics:

The number of transfers between the system buffers and various disks.
The number of times the system buffer was accessed.

Cache read efficiency should be close to 100%. Cache write efficiency should be approximately 66%. However, low percentages do not always indicate performance problems.

Paging Statistics

This graph indicates whether or not a system is short of memory. *up.time* checks whether or not the *pgscan* rate and *page-out* statistics are consistently high. Use the following equation to calculate the scan rate threshold:

$$\text{scan threshold} = \text{handspreadpages} \div \text{residence time}$$

The *handspreadpages* variable is fixed at 8192 on UltraSPARC systems with more than 256 MB of memory. The *residence time* variable is generally fixed at 30 seconds. Therefore, the default scan rate threshold is 273 .

You should also examine the swap device for excessive activity. To identify the device, check the file */etc/vfstab* for the *tmpfs* file system. You can also use the *swap -l* command to list the physical partitions that are being used for swap on the system.

Free Swap

When a program requires more memory than is physically available, information that is not being used is written to a temporary buffer on the hard disk, called *swap* . The Free Swap graph charts the amount of available free swap space, as a percentage of total available free swap space.

Microsoft Windows writes data to the Windows Page File when it needs additional memory. The Windows Page File can range in size from 20 million bytes to over 200 million bytes. The *VPaging File(_Total)% Usage* performance counter extracts page file information.

On Solaris, swap space is separated into:

Physical swap space

The actual space on a disk available for swapping.

Virtual swap space

The amount of physical swap space and the amount of memory that is available for swapping.

If the amount of swap space drops to zero, then the system cannot create new processes or store information in the */tmp* file system.

Linux swaps data to a dedicated swap partition.

Generating a Memory Usage Graph

To generate a memory usage graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - [Used](#)
 - [Cache Hit Rate](#)
 - [Paging Statistics](#)
 - [Free Swap](#)
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph* .

Graphing Processes

up.time uses the following graphs to chart the activity of processes on a system:

Number of Processes

This graph charts the number of processes that are currently running on a system. The process count is taken from the system kernel, and can be used to determine process usage trends.

Process Running, Blocked, Waiting

This graph indicates whether or not there is enough CPU capacity for the processes that are being 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 being monitored, there will always be some blocked processes. However, you can increase the throughput of batch jobs by removing disk bottlenecks.

Process Creation Rate

This graph determines whether or not there are runaway processes on a system or if a forking-based process (like a Web server) is spawning too many processes over a specified period of time.

Generating a Process Graph

To generate a process graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - [Number of Processes](#)
 - [Process Running, Blocked, Waiting](#)
 - [Process Creation Rate](#)
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph* .

Graphing TCP Retransmits

The TCP Retransmits graph indicates whether or not data is being transmitted over a network. Using TCP, information is transmitted in pieces called *packets* . A packet consists of:

A header

Contains transmission information, such as the IP addresses of the sender and receiver, the protocol that is being used, and the packet number.

A payload

Contains the data that is being sent.

A trailer

Contains data that denotes the end of the packet, as well as error correction information.

TCP retransmits indicate that certain network services may not be completing properly because of a high load on a network or a system. A lost packet can indicate network congestion, and requires the sender to reduce the transmission rate and to retransmit the packet. A slower transmission rate combined with retransmitted packets reduces network performance.

Generating a TCP Retransmits Graph

To generate a TCP retransmits graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *TCP Retransmits*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph*.

Graphing User Activity

up.time uses the following graphs to chart the activity of users on a system:

Login History

The number of times or frequency at which a user has logged into a system during any 30 minute time interval.

Sessions

The number of sessions or number of distinct users who are logged into a system during any 30 minute time interval.

Using these graphs, an administrator can identify user load and whether or not there is any correlation between user logins or number of sessions and problems with the performance of the system. These graphs use the same input criteria, but they return different data.

Generating a User Activity Graph

To generate a user activity graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click either *Login History* or *Sessions*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph*.



Note - If there is no data to graph, the message No Data found for the given time range appears in the graph window.

Workload Graphs

The three workload graphs determine the demand that network and local services are putting on a system. The graphs chart an aggregate amount of performance information for a given user, group, or process.

You can generate the following workload graphs:

Workload - User

The demand that network and local services are putting on the system, based on the IDs of the users who are logged into a system.

Workload - Group

The demand that network and local services are putting on the system, based on the IDs of the user groups that are logged into a system.

Workload - Process Name

The demand that network and local services are putting on a system, based on the processes that are running.

These graphs use the same input criteria, but they return different data.

Each workload graph captures the following metrics:

CPU %

The percentage of CPU time that is taken up by a user, group, or process.

Memory Size

The amount of the page file and virtual memory that is taken up by a user, group, or process.

On Windows systems, Memory Size is called *Virtual Bytes*.

RSS

The Run Set Size, which is the amount of physical memory that is being used by a user, group, or process. On Windows systems, RSS is called *Working Set*.

 *Note - Graphs generated for SNMP agents only chart the memory metric.*


Generating a Workload Graph

To generate a workload graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - [Workload - User](#)
 - [Workload - Group](#)
 - [Workload - Process Name](#)
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click one of the following metrics:
 - [See CPU %](#).
 - [See Memory Size](#). or Virtual Bytes (on UNIX and Windows, respectively)
 - [See RSS](#). or Working Set (on UNIX and Windows, respectively)

 *Note - You can only graph one metric at a time.*

6. Select one or more of the available users, groups, or processes from the list.
If you are generating a workload graph by processes, (i.e., *Workload - Process Name* graph), enter a regular expression in the *Process Selection Regex* field to automatically add matching process names for graphing, and avoid dealing with ungainly lists of system processes.

 *Note - The list of available process will vary by server and by operating system.*

7. Click *Add*.
8. Click *Generate Graph*

Workload Top 10 Graphs

The three Workload top 10 graphs chart the 10 processes that are consuming the most CPU resources. Consumption of CPU resources is tracked via one of the following: a user ID, a group ID, or the name of a process. Workload Top 10 graphs enable you to quickly determine which processes are consuming the most CPU resources over a specified time period.

Each graph uses the same input criteria, but they return different data.

Generating a Workload Top 10 Graph

To generate a Workload Top 10 graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - [Workload Top 10 - User](#)
 - [Workload Top 10 - Group](#)
 - [Workload Top 10 - Process Name](#)
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click one of the following options:
 - CPU %
 - Memory Size
 - RSS

Graphs generated for SNMP agents only chart the memory size metric.
6. Click *Generate Graph*.

LPAR Workload Graphs

up.time can collect workload information from logical partitions (LPARs) that are running on pSeries servers. The following graphs visualize the workload information for all LPARs on a server:

[Workload - CPU](#)

The amount of CPU time that is being used by the LPAR.

[Workload - Memory](#)

The total amount of memory being used by an LPAR.

[Workload - Disk](#)

The amount of data that has been transferred to and from the disk.

[Workload - Network](#)

The amount of data that has been transferred over the network interface used by the LPAR.

You can also graph the CPU entitlement of individual LPARs using the CPU Utilization graph. See "[LPAR CPU Utilization Graphs](#)" for more information.

Generating an LPAR Workload Graph

To generate an LPAR Workload graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the pSeries server which is hosting the LPARs whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - Workload - CPU
 - Workload - Memory
 - Workload - Disk
 - Workload - Network
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph* .

LPAR CPU Utilization Graphs

Using the CPU Utilization graph, you can better determine the CPU entitlements of the LPARs on a system. The entitlements indicate the amount of CPU power that is assigned to an individual LPAR. For example, an entitlement of 0.5 indicates that an LPAR is assigned half of the processing power of a CPU.

You can use the graphs to give you a clearer view of how much you may need to increase an LPAR's entitlement. Instead of using trial and error to determine optimum entitlements, you can use actual data to determine accurate entitlements.

To generate an LPAR CPU Utilization graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the pSeries server which is hosting the LPAR whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Under the *LPAR Workload* heading, click *Workload - CPU Utilization* .
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Select the name of the LPAR whose information you want to graph.

If the message *There are no LPARs for this date range* is displayed, do one of the following:

- Click the *Update List* button.
- Change the date range.
- Click *Generate Graph* .

Network Graphs

Network graphs track the performance and reliability of your computing network. You can generate the following network graphs:

The I/O and Errors graphs use the same input criteria, but return different data. NetFlow graphs are available if *up.time* is integrated with Scrutinizer.

I/O

The I/O graph charts the average amount of data that is moving in and out of a network interface over a specified time period. *up.time* also identifies bursts of network traffic.

The I/O graph captures the following statistics:

[In bytes](#)

The number of bytes received over the network interface each second.

[Out bytes](#)

The number of bytes sent by the network interface each second.

Errors

The Errors graph charts the number of network interface errors that occur each second. The most common types of errors include collisions in a hubbed environment or the presence of full-duplex handshake errors between a system and a switch.

As well, the following communication line problems can cause network errors:

Excessive noise.

Cabling problems.

Problems with backbone connections.

The Errors graph captures the following statistics:

[In Errors](#)

A data packet was received but could not be decoded because either the header or trailer of the packet was not available.

[Out Errors](#)

A data packet could not be sent due to problems transmitting the packet or formatting the packet for transmission.
Collisions

The simultaneous presence of signals from two nodes on the network. A collision can occur when two nodes start transmitting over a network at the same time. Packets that are involved in a collision are broken into fragments and must be retransmitted.

NetFlow

The NetFlow graphing function transfers you to your Scrutinizer instance.

For node-type Elements that are exporting data to Scrutinizer, a graph that covers a specified time frame is generated. It shows the monitored node's bi-directional throughput rates through known ports, which are determined based on use by all known applications.

For other Elements, the generated graph shows network traffic from the host, allowing you to pinpoint heavy users.


Generating a Network Graph

To generate network graphs, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click one of the following options:
 - [I/O](#)
 - [Errors](#)
 - [NetFlow](#) (available if *up.time* has been integrated with Scrutinizer)
4. For I/O and Errors graphs, select the start and end dates and times for which the graph will chart data. For NetFlow, select one of the set time frames. For more information, see [Understanding Dates and Times](#)
5. For I/O and Errors graphs, select one or more network interfaces from the *Available Interfaces* list, and then click *Add*.
6. Click *Generate Graph*.

Viewing the Status of a Network Device

The Quick Snapshot for a network device summarizes both the recent (24-hour) and current performance of SNMP-based devices, and can help administrators identify potential issues.

 If there are not 24 hours' worth of data available, *up.time* will use data from as far back as possible to generate charts.

Network Device Quick Snapshot ContentsThe Quick Snapshot is typically used as a preliminary step toward root-cause analysis. When you first acknowledge an issue by clicking the network device's Element name in either *Global Scan* or the *My Alerts* section of *My Portal* , you are shown its Quick Snapshot. From here, you can scan the information provided in the charts and tables (e.g., overloaded ports, or excessively long round-trip times), and begin further investigation.

The following information is displayed in a network device's Quick Snapshot.

Performance Charts	
% Packet Loss	<ul style="list-style-type: none">• line graph showing the percentage of transmitted packets that were lost• shows performance for the last 24 hours, in increments dependent on the network device's Platform Performance Gatherer's polling intervals (default: 10 minutes)
Average Round-Trip Time	<ul style="list-style-type: none">• line graph showing the number of milliseconds for ping to be returned from network device• shows performance for the last 24 hours, in increments dependent on the network device's Platform Performance Gatherer's polling interval (default: 10 minutes)
Port Status	
Port Name	the name of the port on the network device
Port Type	the interface type (i.e., Ethernet or Virtual/VLAN)
Usage	the percentage of the port's maximum throughput that was used during the most recent time interval
In Rate	the average throughput of inbound packets, in Mbps, during the most recent time interval
In Usage	the percentage of the port's maximum throughput that was used by inbound packets during the most recent time interval
Out Rate	the average throughput of outbound packets, in Mbps, during the most recent time interval
Out Usage	the percentage of the port's maximum throughput that was used by outbound packets, during the most recent time interval
Errors	the average number of errors per second, during the most recent time interval

Discards	the average number of packets discarded per second, during most recent the time interval
Status	the current status of the port, based on information retrieved from the network device's Platform Performance Gatherer service

Viewing a Quick Snapshot for a Network Device

To display the Quick Snapshot page for a network device Element, do the following:

1. In the *My Infrastructure* panel, locate the network device 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 a network device Element's profile, you can always access its Quick Snapshot by clicking the *Graphing* tab, then clicking *Quick Snapshot* in the tree panel.

Graphing Network Device Performance

up.time allows you to generate graphs to display the performance of the following:

Network Device Port I/O

The I/O graph displays the average amount of data moving in and out of a network device's ports over a specified time period. This can help you confirm bursts in network traffic, and identify ports that are receiving and transmitting large amounts of data in relation to their maximum throughput.

You can generate top-10-port graphs based on a specific criterion, or focus on a specific port on your network device, and create a graph that includes multiple metrics.

Network I/O Metrics

The following metrics can be used when generating a Network I/O graph for a network device Element:

Total Rate	the combined incoming and outgoing data rates, in Mbps, for the port during the time period
Usage	the percentage of the port's maximum throughput that was used by inbound and outbound packets, during the time interval
In Rate	the average throughput of inbound packets, in Mbps, during the time interval
In Usage	the percentage of the port's maximum throughput that was used by inbound packets during the time interval
Out Rate	the average throughput of outbound packets, in Mbps, during the time interval
Out Usage	the percentage of the port's maximum throughput that was used by outbound packets, during the time interval

Graphing Network I/O Rates for a Network Device

To generate a Network I/O graph for a network device, do the following:

1. Go to the Element's Quick Snapshot page.
For example, in the *My Infrastructure* panel, find the network device Element 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 will chart 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 next step, select whether you will generate a *Top 10* ports graph, or a graph for a *Specific* port.
If you select *Specific*, an Element selection dialog will appear, requiring you to select a specific port on the network device.
6. Select the network metric to include in the graph.
If you are graphing I/O for a *Specific* port, you can include multiple metrics in the graph.
7. Click *Generate Graph*.
A pop-up window will appear, displaying the network I/O rate graph you have configured.

Network Device Port Errors

The network device Errors graph displays the number of errors or discards that occur each second. The following communication line problems can cause network errors:

- excessive noise
- cabling problems
- problems with backbone connections

Network Error Metrics

The following metrics can be used when generating a Network Error graph:

Errors	the total number of errors per second during the time period
In Errors	the number of packets received, but unable to be decoded, per second, due to a missing header or trailer
Out Errors	the number of packets that were not sent, per second, due to problems transmitting the packet or formatting the packet for transmission
Discards	the total number of packets dropped per second, through the port, during the time period
In Discards	the number of packets inbound through the port that were dropped per second, during the time period
Out Discards	the number of packets outbound through the port that were dropped per second, during the time period

Graphing Network Error Rates for a Network Device

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 network device Element 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 *Errors*.
3. Select the start and end dates and times for which the graph will chart 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 next step, select whether you will generate a *Top 10* ports graph, or a graph for a *Specific* port.
If you select *Specific*, an Element selection dialog will appear, requiring you to select a specific port on the network device.
6. Select the network metric to include in the graph.
If you are graphing network errors for a *Specific* port, you can include multiple metrics in the graph.
7. Click *Generate Graph*.
A pop-up window will appear, displaying the network error graph you have configured.

Disk Performance Statistics Graph

The Disk Performance Statistics graph charts a set of disk performance metrics returned by utilities - such as *perfmom* on Windows, and *iostat* or *sar* on Solaris - that are running on a system.

Requests can experience delays proportional to the length of the request queue minus the number of spindles on the disks. For optimal performance, this difference should be less than two on average.

Generating a Disk Performance Statistics Graph

To generate a Disk Performance Statistics graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *Disk Performance Statistics*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#).
5. Select one of the following options:
 - **Percent Busy**
The percentage of the disk capacity that is being used.



Note - For NFS systems, 100% busy does not indicate that the server itself is saturated, but that the client always has outstanding requests to that server.

- **Average Queue**
The average number of processes that are waiting to access the disk.
The length of the queue is affected by how busy the system is and the amount of time that each transaction requires to perform a disk operation. A complete transaction must occur before the next transaction can start. Longer disk operations per transaction increases the average length of the queue.
 - **Read/Writes**
The number of read/write requests, per second, from or to a disk.
 - **Throughput (blks/s)**
The amount of disk traffic, in blocks of 512 bytes, that is flowing to and from a disk each second.
 - **Average Wait Time**
The average time, in milliseconds, that a transaction is waiting in a queue. The wait time is directly proportional to the length of the queue.
 - **Average Serve Time**
The average time, in milliseconds, required to perform a task.
 - **All of the above for one disk**
up.time graphs all of the metrics listed above for a single disk.
6. Select the disks for which you want to collect information from the list.
If you select multiple disks and selected *All of the above for one disk* in step 5, then *up.time* only graphs information for the first disk that you selected.
 7. Click *Generate Graph*.

Top 10 Disks Graph

The Top 10 Disks graph displays the ten busiest disks in your environment as of the last sample that *up.time* has taken. If there are fewer than ten disks on the system, then all of the disks on a system will be charted in the graph.

Generating a Top 10 Disks Graph

To generate a Top 10 Disks graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *Top 10 Disks*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Select one of the following options:
 - **Percent Busy**
The percentage of the disk capacity that is being used.



Note - For NFS systems, 100% busy does not indicate that the server itself is saturated, but that the client has outstanding requests to that server.

- **Average Queue**
The average number of processes that are waiting to access the disk.
The length of the queue is affected by the amount of time that each transaction requires to perform a disk operation. For both sequential and random disk transactions, a complete transaction must occur before the next transaction can begin. Longer disk operations per transactions increase the average length of the queue.
 - **Read/Writes**
The number of read/write requests per second from or to a disk.
 - **Throughput (blks/s)**
The amount of traffic, in 512 byte blocks, that is flowing to and from a disk.
 - **Average Wait Time**
The average time, in milliseconds, that a transaction is waiting in a queue. The wait time is directly proportional to the length of the queue.
 - **Average Serve Time**
The average time, in milliseconds, required to perform a task.
6. Click *Generate Graph*.

File System Capacity Graph

A File System Capacity graph charts the amount of total and used space, in kilobytes, on a server's disk. On Windows servers, *up.time* looks at the capacity of the main partition (usually the C:\ drive). On UNIX and Linux servers, *up.time* looks at the individual file systems (for example, */var*, */export*, */usr*) on all the disks on the server.



Note - If a single disk system has no partitions, then the file system capacity is the same as the disk capacity.

The File System Capacity graph visualizes the following statistics:

- **Total Size**
The total amount of space available on the system.
- **Space Used**
The amount of space on the file system that has been used.

Generating a File System Capacity Graph

To generate a File System Capacity graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *File System Capacity*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Select one or more file systems from the list.
If you are generating a graph for a Windows system, you will only be able to generate a graph for the C:\ drive.
6. Click *Generate Graph*.

VXVM Stats Graph

The VXVM Stats graph charts the amount of data written to or read from a Solaris volume that is managed by the Veritas Volume Manager. Veritas Volume Manager is storage management system that operates between a host's operating system and its filesystems or database management systems. Veritas Volume Manager enables you to manage disk drives on a system as if they were *volumes* (logical devices that appear to be physical partitions on a disk).

Depending on the options that you specify, this graph contains the following information:

the number of read and write operations to and from the volume
the number of blocks that were read and written to and from the volume
the amount of time that is required to read data from and write data to the volume

If Veritas Volume Manager is not running on a host, or if *up.time* cannot connect to the volume, an error message informing you that *up.time* cannot detect the Veritas Volume Manager appears in the *Graphing* subpanel.

In the *Info & Rescan* panel, verify that the entry *Has a Logical Volume Manager?* is set to *Yes*. If it is, then ensure that you can connect to the host from the Monitoring Station.

Generating a VXVM Stats Graph

To generate a VXVM Stats graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *VXVM Stats*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#).
5. In the *Available Disk Groups and Volumes* area, select one or more volumes on which to report.
The disk groups or volumes that appear in this area will vary from system to system. You must select at least one disk group or volume.
6. Select one of the following options:
 - *I/O Operations*
The number of times, per second, that data is written to and read from the volume.
 - *Block Throughput*
The amount of disk traffic, in blocks of 512 bytes, that is flowing to and from the volume.
 - *Average Service Times*
The average amount of time, in milliseconds, that is required for a request to be carried out.
7. If necessary, uncheck either of the *Read* or *Write* checkboxes.
Depending on the option you chose in step 6, the Read and Write options chart the following information in the graph:
 - If you selected *I/O Operations* in step 6, the number of read and write operations to and from the volume.
 - If you selected *Block Throughput* in step 6, the number of blocks that were read and written to and from the volume.
 - If you selected *Average Service Times* in step 6, the amount of time requires to read and write data to and from the volume.



Note - Select only one option if you are comparing more than one volume.

8. Click *Generate Graph*.

Novell NRM Graphs

up.time can collect data from systems that are running version 6.5 of the Novell Remote Manager (NRM). *up.time* retrieves NRM service metrics and then stores this information in the DataStore. Using the data that is collected from NRM, you can generate graphs for the following metrics:

- **Available Memory**
The amount of memory that is not allocated to any service.
- **DS Thread Usage**
The number of server threads that Novell eDirectory uses. The server thread limit ensures that server threads are available for other functions as needed.
- **Work To Do Response Time**
The amount of time that a Work To Do process requires to run from the time a process is scheduled.
- **Allocated Server Processes**
How the service processes are allocated on the NRM system.
- **Available Server Processes**
The number of available processes on the NRM system.
- **Abended Thread Count**
The number of threads that have *abended* (ended abnormally) and that are suspended because of abended recovery.
- **Packet Receive Buffers**
The status of Packet Receive Buffers (which transmit and receive packets) for the NRM system.
- **Available ECBs**
The status of available Event Control Blocks (ECBs), which are Packet Receive Buffers that have been created but which are not currently being used.
- **LAN Traffic**
Whether or not the NRM system can transmit and receive packets.
- **Available Disk Space**
The status of the available disk space on a server.
- **Disk Throughput**
The status of amount of the data being read from and written to the storage media on the server.
- **Connection Usage**
The number of connections that are being used, and the peak number of connections used on this server.

For more information about Novell NRM systems, see [Novell NRM Systems](#).

Generating a Novell NRM Graph

To generate a Novell NRM graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the Novell NRM system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab and then click one of the metrics in the list.
3. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#).
4. Click *Generate Graph*.

Instance Motion Graphs

The VMware VMotion tool enables you to move ESX instances from one server to another without any downtime or loss of data. You would use VMotion to, for example, 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 has been running on over a given time range.

Generating an Instance Motion Graph

To generate an Instance Motion graph, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the ESX instance whose motion you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *Instance Motion*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)
5. Click *Generate Graph*.

Displaying Detailed Process Information

Detailed process information provides an insight into how various user and system processes are consuming system resources. The information is not presented in a graph - it is a table that contains the following information:

- **Process**
The name of the process, which is taken from its executed path name.
- **PID**
The number that identifies the process.
- **PPID**
The number that identifies the parent process. The PPID can help identify possible relationships between processes.
On Windows systems, the PPID is called the *Creating Process ID*.
- **UID**
The ID of the user or account that has been consuming CPU time.
On Windows systems, the UID is called the *Owner*.
- **GID**
The ID of the group that has been consuming CPU time.
On Windows systems, the GID is called the *Group Name*.
- **Memory Used**
The amount of memory, expressed as a percentage of total available memory, being 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. For example, if five Oracle processes are using 10% of available memory, this does not indicate that Oracle is consuming 50% of system memory.
- **RSS**
Run Set Size - the amount of physical memory that is being used.
On Windows systems, RSS is called the *Working Set*.
- **CPU %**
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*.
- **User Time**
The amount of time (in seconds) that a particular user, group, or account has been using the CPU.
This value is not displayed for Windows systems.
- **User System Time**
The amount of time (in seconds) that a process has been consuming system time on the CPU.
This value is not displayed for Windows systems.



Note - You can get a better indication of the amount of work a process has done by dividing this amount by a sample of time - for example, five minutes.

- **Start Time**
The time at which the process started. This can be used to determine the lifetime of a process.



Note - The process information for the current date and time is displayed in the Graphing subpanel.

Generating Detailed Process Information

To display detailed process information, do the following:

1. In the *Global Scan* or *My Infrastructure* panel, click the name of the system whose information you want to graph.
2. In the *Tree* panel, click the *Graphing* tab.
3. Click *Detailed Process Information*.
4. Select the start and end dates and times for which the graph will chart data. For more information, see [Understanding Dates and Times](#)

5. Click *Display Process Information*.

A window containing a chart that lists the process information for the time period that you specified appears.

Detailed Process Information

☒ Specific Date and Time

Date Range: YYYY-MM-DD HH:MM:SS

☐ Last

From: 2008-04-16 00:00:00 24

☐ Quick Date

To: 2008-04-16 23:59:59 24

Display Process Information

AIX5 (aix5l) - Process Information - displaying latest sample dated 2008-04-16 16:03:40

Process	PID	PPID	UID	GID	Memory Used	RSS	CPU %	Mem %	Runtime	Children Runtime	Start Time
dtgreet	3902	4726	root	system	2.59 MB	840 KB	0.2	1	4h 35m	0s	2008-01-02 09:05:58
inetd	6966	4948	root	system	600 KB	644 KB	0.1	1	1h 57m	0s	2008-01-02 09:05:12
AIXPowerMgtDaemon	11616	1	root	system	1.50 MB	96 KB	0	0	0s	0s	2008-01-02 09:05:58
bioid	10070	4948	root	system	340 KB	220 KB	0	0	0s	0s	2008-01-02 09:05:47
diagd	13934	1	root	system	284 KB	304 KB	0	0	0s	0s	2008-01-02 09:06:00
getty	12902	1	root	system	692 KB	508 KB	0	0	0s	0s	2008-01-02 09:06:01
httpd-lite	13676	1	imnadm	imnadm	428 KB	336 KB	0	0	0s	0s	2008-01-02 09:06:01
IBM.AuditRMd	24510	4948	root	system	2.27 MB	2.41 MB	0	2	1s	0s	2008-04-09 22:29:29
IBM.CSMAgentRMd	16814	4948	root	system	2.18 MB	2.41 MB	0	2	3s	0s	2008-04-16 00:35:08
IBM.ERmd	17270	4948	root	system	2.72 MB	2.89 MB	0	3	0s	0s	2008-04-14 14:29:07
ksh	19538	25872	uptime	adm	576 KB	784 KB	0	1	0s	0s	2008-04-16 15:36:20
ksh	23208	15438	uptime	adm	508 KB	720 KB	0	1	0s	0s	2008-04-16 15:36:20
qdaemon	11098	4948	root	printq	396 KB	264 KB	0	0	3s	0s	2008-01-02 09:05:55
rmcd	13422	4948	root	system	2.60 MB	1.02 MB	0	1	21m 52s	0s	2008-01-02 09:06:03
rpc.lockd	10590	4948	root	system	496 KB	180 KB	0	0	0s	0s	2008-01-02 09:05:53
rpc.statd	10328	4948	daemon	sys	1.96 MB	324 KB	0	0	0s	0s	2008-01-02 09:05:49
sadc	18446	19538	root	adm	256 KB	272 KB	0	0	0s	0s	2008-04-16 15:36:21
uptmagn1	15438	6966	uptime	adm	240 KB	276 KB	0	0	0s	0s	2008-04-16 15:36:22
uptmagn1	25872	6966	uptime	adm	228 KB	264 KB	0	0	0s	0s	2008-04-16 15:36:19
writesrv	11360	4948	root	system	464 KB	184 KB	0	0	0s	0s	2008-01-02 09:05:57

6. From the dropdown list, select the date and time for which you want to view process information.