

Using Reports

Reports for Performance and Analysis

The following reports enable you to visualize the overall performance of a system in the *up.time* environment, as well as analyze the information to determine the cause of problems with those systems:

Resource Usage Report

The Resource Usage report tracks the usage of system resources and performance information for systems over a given period of time. In addition to the usage information being reported on, the report displays the following information:

- the name and description of the system
- an overview of the system configuration, including architecture, memory size, operating system version, number of CPUs, and host ID

Creating a Resource Usage Report

To create a Resource Usage report, do the following:

1. In the Reports Tree panel, click *Resource Usage*.
2. In the *Date and Time Range* area, select the dates and times on which to report. For more information, see [Understanding Dates and Times](#)
3. Select one or more of the following report options:
 - **Service Status**
The status of each service that has been assigned to the selected system or systems. The statuses are *OK*, *WARN*, *CRIT*, *MAINT*, and *UNKNOWN*.
 - **Network I/O**
The average amount of traffic, measured in megabytes per second, that is travelling through the network interfaces. The report also identifies bursts in network activity that may occur over short intervals. This information appears as a graph in the report.
 - **Free Memory**
The amount of free memory available to the system. This information appears as a graph in the report.
 - **File System Capacity**
The amount of free disk space on the system. This information appears as a graph in the report.
 - **Workload (Top 10 - RSS)**
The top 10 processes that are consuming physical memory (in KB), as measured by the run-set size (RSS) of the process. This information appears as a graph in the report.



Note - This graph does not appear when you generate this report for a VMware ESX system.

- **Resource Utilization**
The average and maximum amount of CPU and memory use.
- **Network Errors**
Any errors that have occurred with the physical network interface. The errors can be, for example, collisions or handshake errors between a system and a switch.
- **Page Scanning Statistics**
The number of file system pages scanned by the page scanning daemon. This information appears as a graph in the report.
- **Workload (Top - 10 CPU)**
The top 10 processes that are consuming CPU time, grouped by user ID, group ID, and process name. This information appears as a graph in the report.




Note - This graph does not appear when you generate this report for a VMware ESX system.

- **Multi-CPU**
The percentage of total CPU time that is being used on systems with more than one CPU.




Note - If you find the report's rendered graph too dense due to a large number of CPUs, alternatively generate a Multi-CPU Usage graph while including fewer CPUs.

- **CPU Performance Graph**
Tracks the performance of a system's CPU over a specified time period. This information appears as a graph in the report.
- **TCP Retransmits**
Any network services that may not be completing properly because of undue network or system load. This information appears as a graph in the report.
- **Disk Statistics**
The following statistics for each disk on a system:
 - percentage of the disk that is busy
 - average queue length
 - number of reads and writes per second
 - number of blocks being accessed per second
 - average wait time, in seconds
 - average service time, in seconds

 *Note - If the system for which you are creating a report for has multiple disks, a graph for each disk on the system is generated.*

- Workload (Top 10 - Memsize)
- The top 10 processes that consume system memory, based on the total memory size of the processes - including virtual pages and shared memory. This information appears as a graph in the report.

 *Note - This graph does not appear when you generate this report for a VMware ESX system.*

Optionally, click *Select All* to generate a report on all of the options listed above.

4. If you selected more than one report option and plan to report on more than one system, you can optionally click the *Group report options by system* checkbox.
Selecting this option combines the metrics for each system for which you are generating the report.
5. To generate reports for systems in specific groups, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific systems, select the systems from the *List of Systems*.
8. Select a report generation option. See [Report Generation Options](#) for details.
9. If you want to save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Multi-System CPU Report

The Multi-System CPU report charts and compares the CPU performance statistics from multiple systems in your environment. These statistics indicate whether or not the systems are exhibiting balanced behavior, or if processes are being forced off CPUs in certain circumstances.

Creating a Multi-System CPU Report

To create a Multi-System CPU report, do the following:

1. In the Reports Tree panel, click *Multi-System CPU*.
2. In the *Date and Time Range* area, select the dates and times on which to report. For more information, see [Understanding Dates and Times](#).
3. If you want the report to only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.
For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select *1:00* from the *Start* dropdown list, and *13:00* from the *End* dropdown list.
4. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
5. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
6. If you are generating reports for specific systems in your environment, select them from the *List of Systems*.
7. Select a report generation option. See [Report Generation Options](#) for details.
8. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

CPU Utilization Summary Report

The CPU Utilization Summary report generates a tabular summary of the CPU and memory consumption over a specific time period. Specifically, this report returns the following information:

- number of CPUs on the server.
- the total processor speed of all the CPUs, in MHz
- the maximum, minimum, and average CPU use, expressed as a percentage
- the maximum, minimum, and average memory use, expressed as a percentage
- the maximum, minimum, and average page scan per second, expressed as a percentage

Creating a CPU Utilization Summary Report

To create a CPU Utilization Summary report, do the following:

1. In the Reports Tree panel, click *CPU Utilization Summary*.
2. In the *Date and Time Range* area, select the dates and times on which to report. For more information, see [Understanding Dates and Times](#).
3. Select one of the following options from the *Sort by* dropdown list to sort the results that *up.time* returns:
 - Average CPU (the default)
 - Hostname
 - # of CPUs
 - CPU Speed
 - Maximum CPU
 - Minimum CPU
 - Average Memory
 - Maximum Memory
 - Minimum Memory
 - Average Page Scan
 - Maximum Page Scan

- Minimum Page Scan
4. Select *Ascending* or *Descending* from the *Sort Direction* dropdown list.
 5. Optionally, in the *Minimum sort value for inclusion* field enter a value for the sort threshold.
The report displays items from the *Sort By* list, whose value is equal to or greater than the value in this field. For example, if you chose *# of CPUs* from the *Sort by* list and set this field to 2, the report only displays systems with two or more CPUs.
 6. Select one or more of the following *CPU statistics* at which the report will look:
 - sys
The percentage of CPU time that is being use to carry out system processes.
 - usr
The percentage of CPU time that is being used to carry out user processes.
 - wio
The percentage of CPU time that could be handling processes, but which is waiting for I/O operations to complete.
 7. Select one or more of the following statistics on which to report:
 - CPU
The percentage of CPU resources that are being used.
 - Memory
The percentage of system memory that is being used.
 - Page Scans
The number of page scans per second.



Note - The statistic you select must match the sort criteria that you selected in step 4. For example, if your sort criteria is Average CPU you must also select the CPU statistic. Otherwise, an error message appears when you try to generate the report.

8. Optionally, in the *Architectures to exclude* field enter either the name of a system architecture or a regular expression that *up.time* will use to ignore certain system architectures when generating the report.
For example, if you want to exclude all Solaris systems from the report, enter *SunOS* in the field.



Note - up.time determines the architecture of a system by checking the output of the `uname -a` command on UNIX or Linux, or by analyzing one or both of the following Windows registry keys:

```
HKEY_LOCAL_MACHINE\\Software\\Microsoft\\
WindowsNT\\CurrentVersion
```

```
HKEY_LOCAL_MACHINE\\Software\\Microsoft\\
Windows\\CurrentVersion
```

9. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
10. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
11. If you are generating reports for specific systems in your environment, select them from the *List of Systems*.
12. Select a report generation option. See [Report Generation Options](#) for details.
13. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

CPU Utilization Ratio Report

The CPU Utilization Ratio report charts, in a table, the ratio of the percentage of CPU usage over a specified period of time. The ratio is derived by dividing the percentage of system time that is being used by the percentage of user time. For example, if the amount of system time that is being used is 22.12% and the amount of user time is 5.2%, then the CPU utilization ratio is 4.25.

This report contains the following information:

- the names of the hosts for which the report has been generated
- the percentage of CPU time that is being used to carry out user processes (*USR %*)
- the percentage of CPU time that is being use to carry out system processes (*SYS %*)
- the CPU utilization ratio for each host, which is derived by dividing *SYS %* by *USR %*=

Creating a CPU Utilization Ratio Report

To generate a CPU Utilization Ratio report, do the following:

1. In the Reports Tree panel, click *CPU Utilization Ratio*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. If you want the report to only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.
For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select *1:00* from the *Start* dropdown list, and *13:00* from the *End* dropdown list.

- Optionally, enter a value in the *Highlight ratios over threshold* field.
Any ratios that exceed the value in this field will be highlighted in the report. For example, if you enter 2 and a server returns a ratio of 3.5%, that ratio is highlighted.
- If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
- To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
- If you are generating reports for specific systems in your environment, select them from the *List of Systems*.
- Select a report generation option. See [Report Generation Options](#) for details.
- To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Wait I/O Report

The Wait I/O report enables you to determine the amount of time that processes spend waiting on I/O from a system device.

The Wait I/O report contains the following information:

- the names of the hosts for which the report has been generated
- the average, maximum, and minimum wait I/O times expressed as percentages

Creating a Wait I/O Report

To create a Wait I/O report, do the following:

- In the Reports Tree panel, click *Wait I/O*.
- In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
- If you want the report to only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.
For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select *1:00* from the *Start* dropdown list, and *13:00* from the *End* dropdown list.
- Optionally, enter a value in the *Highlight average WIO over threshold* field.
Any system with an average Wait I/O percentage that exceeds the value that you enter in this field will be highlighted in red in the report. As well, the following text appears in the header of the report:
Systems with an Average Wait I/O over x.x% are highlighted
Where x.x is the percentage that you entered in this field.
- If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
- To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
- If you are generating reports for specific systems in your environment, select them from the *List of Systems*.
- Select a report generation option. See [Report Generation Options](#) for details.
- Do one of the following:
 - Click the *Generate Report* button.
 - Enter a name for the report in the *Save to My Portal As* field, and optionally enter text in the *Report Description* field. Then, click *Save Report*.
The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
- To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Scheduling Reports](#) for more information on configuring a scheduled report.

Inventory Report

The Inventory report provides details about the composition of your monitored infrastructure by operating system, across physical and virtual Elements. The report contents can optionally be organized by group, and can include individual Element entries.

These different reporting options allow you to confidently assess your inventory from a variety of perspectives, and help you answer the following types of questions:

- How many older versions of the *up.time* Agent are deployed on my systems?
- What is the operating system breakdown across my infrastructure, both virtualized, and on physical systems?
- I need to upgrade a particular OS version; which systems are candidates for this deployment?

Inventory Report Information



Note - Since the Inventory report displays all monitored Elements, the report is intended for system administrators. Non-administrative up.time users who do not have permission to view all Elements will not be able to see complete inventory listings.

The following information can be displayed in an Inventory report:

Operating System Summary	
unique breakdown of physical and virtual Elements, with totals by OS, and component totals for OS versions	
Physical Elements	<p>the total number of systems-type Elements (i.e., not network devices, Applications, and SLAs)</p> <p>this total includes virtual machines that are not managed by a VMware vCenter server Element (e.g., LPARs or VMware VMs manually added to up.time, and not through vSync)</p>
Virtual Elements	the total number of Elements running on VM instances (i.e., VM instances with their own UUID)

Operating System	the detected operating system, including VMware environments
Version	the detected operating system version; build version details are listed if available
Element Name	the Element's host name
Architecture	the detected hardware platform type on which the Element's CPUs are running
Agent Version	if applicable, the version of the <i>up.time</i> Agent that is running on the Element
Added Date	the date the Element was added to <i>up.time</i> 's monitored inventory
Group	the Element's <i>My Infrastructure</i> group name

Creating an Inventory Report

To create an Inventory report, do the following:

1. In the Reports Tree panel, click *Inventory Report*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
3. Optionally select the *Don't List Individual Elements* check box to restrict the report to inventory summaries.
4. Optionally select the *Group by Selected System Groups* check box to organize the report by group.
5. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific Elements in your environment, select them from the *List of Elements*.
8. Select a report generation option. See [Report Generation Options](#) for details.
9. Do one of the following:
 - Click the *Generate Report* button.
 - Enter a name for the report in the *Save to My Portal As* field, and optionally enter text in the *Report Description* field. Then, click *Save Report*.
The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
10. To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Scheduling Reports](#) for more information on configuring a scheduled report.

Service Monitor Metrics Report

You can configure the *up.time* service monitors to retain data, which is saved to the *up.time* DataStore for later use. The Service Monitor Metrics report visualizes the retained data in a line chart.

For example, if you have configured a service monitor to retain response time data then this report charts any changes in the response time (in milliseconds) that have occurred over the time period that you specified for the report.

Creating a Service Monitor Metrics report is a two-step process:

- enter the basic parameters for the report
- select the values for the retained on which you want to report

Creating Service Monitor Metrics Reports

To create a Service Monitor Metrics report, do the following:

1. In the Reports Tree panel, click *Service Monitor Metrics*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
3. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
4. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
5. If you are generating reports for specific systems in your environment, select them from the *List of Elements*.
6. Click *Go to page 2*.
A table containing the current retained service metrics appears in the *Service Metrics* subpanel.
7. Click the checkboxes in the *Select* column to select the variables on which you want to report
8. Optionally, select one of the following:
 - Show all non-ranged metrics on one chart
This option combines all of the variables you selected in one chart. Any ranged metrics will appear in their own charts.
 - Display charts as stacked area
Each chart in the report will have two or more data series stacked on top of each other, rather than the line graph that usually appears in the report.
9. To save the report, do the following:
 - Enter a name for the report in the *Save to My Portal As* field.
 - Optionally, enter text in the *Description* field.
 - Click *Save Report*.
The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
10. To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Reports for Capacity Planning

The following reports enable you to visualize the resource usage of systems in your *up.time* environment, and then use that information to better plan, deploy, and consolidate your server resources:

Enterprise CPU Utilization Report

The Enterprise CPU Utilization report enables you to compare the processing power of different types of systems in your environment. Performing this kind of comparison is difficult because different types of systems use different processors - for example, a Windows server uses an Intel processor while a Solaris server may use a SPARC processor. The benchmarks for measuring the power of each type of processor will be different.

An Enterprise CPU Utilization report offers a quick snapshot of the overall performance of the servers in your environment. Based on the information in the report, you can then determine how best to optimize CPU capacity across your enterprise.

up.time can measure processing power using statistics called a *power units*. Power units are the number of CPUs on a system multiplied by the speed of the processors. For example, a Solaris server has four CPUs and each CPU runs at 168 Mhz. The total number of power units for the server is 672 (4 x 168). If you compare this to a Windows server with one CPU running at 2900 MHz (2,900 power units), then you can conclude that the Windows server has more processing power.

Enterprise CPU utilization is a percentage that is derived by dividing the total number of power units used by the total number of power units available. For example, if the number of power units used is 104 and the total number of available power units is 2,346 then the enterprise CPU utilization is 4.34%.

Creating an Enterprise CPU Utilization Report

To create an Enterprise CPU Utilization report, do the following:

1. In the Reports Tree panel, click *Enterprise CPU Utilization*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. If you want the report to only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section
For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select *1:00* from the *Start* dropdown list, and *13:00* from the *End* dropdown list.
4. Select one of the following options from the *Sort by* dropdown list to sort the results that *up.time* returns:
 - Hostname (the default)
 - # of CPUs
 - CPU Speed
 - Power Units Total
 - Power Units Used Total
 - Power Units Used Partial
 - CPU Utilization Total
 - CPU Utilization Partial
5. Select *Ascending* or *Descending* from the *Sort Direction* dropdown list.
6. Select one or more of the following *CPU statistics* at which the report will look:
 - sys
The percentage of CPU time that is being use to carry out system processes.
 - usr
The percentage of CPU time that is being used to carry out user processes.
 - wio
The percentage of CPU time that could be handling processes, but which is waiting for I/O operations to complete.
7. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
8. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
9. If you are generating reports for specific systems in your environment, select them from the *List of Systems*.
You should select more than one system.
10. Select a report generation option. See [Report Generation Options](#) for details.
11. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

File System Capacity Growth Report

The File System Capacity Growth report illustrates the following:

- The used, available, percentage used, and total size of the file system at the beginning and end of the reporting period. The used, available, and total size metrics are measured in megabytes.
- The percentage by which the file system has changed over the reporting period, charting the following: used space, available space, percentage used, and total size of the file system.

On Windows servers with a single disk, *up.time* looks at the capacity of the main partition (usually the C:\ drive). If the Windows server has multiple disks, this report collects information for all of the disks. On UNIX and Linux servers, *up.time* looks at individual file systems (for example, */var*, */export*, or */usr*) on all the disks in the system.

 *Note - This report ignores floppy drives, tapes drives, and CD-ROM drives.*

Creating a File System Capacity Growth Report

To create a File System Capacity Growth report, do the following:

1. In the Reports Tree panel, click *File System Capacity Growth*.

2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
If no data available for the date range, the report displays a message indicating that there is no data for the time period.
3. Optionally, in the *Exclude file system names like* field enter either the name of a file system or a regular expression that *up.time* will use to ignore certain file systems when generating the report.
For example, if you want to exclude the */boot* file system from the report, enter */boot* in the field.
4. Optionally, enter a value in the *Exclude filesystems over % full* field.
This value is expressed as a percentage. The report displays the information for file systems whose used disk space is less than the amount you enter in this field. For example, if you set this field to *45*, the report only displays file systems whose percentage used values are less than or equal to 45%.
5. Click the *Show totals for each system only* checkbox to report only on the total amount by which all file systems on all disks drives have grown, rather than displaying amounts for each file system.
6. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
7. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
8. If you are generating reports for specific systems in your environment, select them from the *List of Systems*.
9. Select a report generation option. See [Report Generation Options](#) for details.
10. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Server Virtualization Report

Many organizations have a number of production servers that are not being used to their full capacity. For example, a server could be running one or two applications and not using much of the hardware. Instead of wasting resources, you can consolidate these applications in a virtual environment, for example using VMware. This enables you to run applications on distinct servers, but without using as much hardware.

The Server Virtualization report can help you to pinpoint physical servers that can be combined on a single virtual server. The report highlights servers that are good candidates for virtualization - ones that do not fully use their CPU, memory, or disk resources.

In the report, each system will have one of the following stars beside it:

- Indicates that the system is a good candidate for virtualization. The corresponding metrics are highlighted in green.
- Indicates that the system is a reasonable candidate for virtualization. The corresponding metrics are highlighted in blue.
- Indicates that the system is a poor candidate for virtualization. The corresponding metrics are not highlighted.

As well, the metrics for Average Power Units Used (*Power Units* measure the power of CPUs by multiplying the number of CPUs on a system by their speed), Avg Disk I/O, and Avg Network I/O for each system may be highlighted.

Creating a Server Virtualization Report

To generate a Server Virtualization report, do the following:

1. In the Reports Tree panel, click *Server Virtualization*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Click the *Display Element custom fields* option to insert the content of the custom fields in the system profile into the report.
The custom fields contain additional information about the system - for example, the types of reports that should be run on this system or when maintenance is scheduled. For more information, see [Editing a System Profile](#)
4. In the Target Machine area, do the following to specify the hardware of the server on which the other servers will be consolidated:
 - Select the type of processor used on the target server from the *Architecture* dropdown list:
 - Alpha
A 64-bit processor from HP.
 - Itanium
A 64-bit processor from Intel.
 - x86
A standard 32-bit processor.
 - Sparc
The range of SPARC processor used on system that run the Solaris operating system.
 - POWER
The POWER processor, used with IBM p-series and i-series servers.
 - Select number of CPUs on the target system from the *Num CPUs* dropdown list. Then, enter the processor speed of the CPUs in the *MHz* field.
For example, if the target system has four CPUs and each have a processor speed of 1,000 MHz, select *4* from the dropdown list and enter *1000* in the field.
 - Select the type of disk interface that is used on the target server from the *Disk I/O* dropdown list:
 - ATA
 - SCSI
 - iSCSI
 - SATA
 - SATA II
 - Fibre
If none of the options above apply, enter the data transfer speed of the disk (measured in megabits per seconds) in the *MBps* field.
 - From the *Network I/O* dropdown list, select the type of disk interface that is used on the target server:
 - 10Mbit
 - 100Mbit
 - 1Gbit
 - 10Gbit
If none of the options above apply, enter the data transfer speed of the network interface (measured in megabits per seconds) in the *MBps* field.

5. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific systems in your environment, select them from the *List of Systems* .
8. Select a report generation option. See [Report Generation Options](#) for details.
9. Do one of the following:
 - Click the *Generate Report* button.
 - Enter a name for the report in the *Save to My Portal As* field, and optionally enter text in the *Report Description* field. Then, click *Save Report* .
The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
10. To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Scheduling Reports](#) for more information on configuring a scheduled report.

Using the Server Virtualization Report

The results of a Server Virtualization report can help you to determine which physical servers to combine on a single virtual server. In order to effectively use the report, you must analyze the results in more depth.

First, look at the average number of power units used by the systems that you want to consolidate on a virtual server. That figure should be less than the total number of power units available on the target system.

Next, look at the disk I/O for the individual systems. If the system is running an application that has high levels of disk usage (for example, a database), that system might not benefit from virtualization. If, however, the target system has a very fast disk, you can still consider moving the candidate system to it.

Also, consider the geographical locations of the systems for which you are generating the report. For example, the report states the four systems of a similar type are good candidates for virtualization. However, two of those system are in different parts of the country or the world. In this case, adding them to a virtual server is not a viable option.

Solaris Mutex Exception Report

Solaris system with two or more CPUs can suffer from mutex (mutual exclusion) locks when two or more threads are waiting for the same resource. During processing, the Solaris kernel maintains locks on various resources. The kernel allocates enough mutex locks to allow multiple CPUs to complete their work simultaneously. However, if two or more CPUs try to get the same lock at the same time, all but one CPU will stall.

The Solaris Mutex Exception report pinpoints multi-processor Solaris systems that have a high number of mutex stalls. The report contains the following information:

- the display name in *up.time* of the system
- the number of CPUs on the system
- the average number of mutex stalls for all the CPUs on the system, over the time period that you specified; if this value exceeds the threshold that you set, it is highlighted in red

Creating a Solaris Mutex Exception Report

To create a Solaris Mutex exception report, do the following:

1. In the Reports Tree panel, click *Solaris Mutex Exception* .
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
If no data available for the date range, the report displays a message indicating that there is no data for the time period.
3. If you want the report to only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.
For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select *8:00* from the *Start* dropdown list, and *18:00* from the *End* dropdown list.
4. Optionally, enter a value in the *Highlight average SMTX over threshold* field.
If the number of mutex stalls for a system, averaged for all of its CPUs over the defined reporting time period, exceeds the value in this field, the number will be highlighted in the report. For example, if you enter *75* and a server returns *93* , that value is highlighted.
5. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific Applications in your environment, select them from the *List of Elements* .



Note - Only Solaris systems with two or more CPUs are shown in the List of Elements.

8. Select a report generation option. See [Report Generation Options](#) for details
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the Solaris Mutex Exception Report

The following is an example of a Solaris Mutex Exception report:

Solaris Mutex Exception		
Date Range: 2008-04-16 00:00:00 to 2008-04-16 17:55:11 between the time range: 00:00 to 23:59		
Multi CPU systems with an Average SMTX over 150.0 are highlighted		
System Name	Number of CPUs	Average SMTX
Opteron (opteron)	2	158.19
The Zones (vmh-t1k-a)	24	0.83

The number of mutex stalls for the first system in the list exceeds the threshold that was set when the report was defined. Based on this information, you can generate one of the following graphs to get a better idea of the performance of the CPUs on the system:

1. Multi-CPU Usage (see [Generating a Multi-CPU Usage Graph](#) for more information)
2. Run Queue Length (see [Run Queue Length](#) for more information)
3. Run Queue Occupancy (see [Run Queue Occupancy](#) for more information)

From there, you determine how to best reduce the queue size to improve performance.

Network Bandwidth Report

The Network Bandwidth report keeps track of the amount of data moving in and out of each network interface on a system. This report helps you identify or confirm that specific systems are being overloaded, based on the amount of data they are sending or receiving; such systems could become bottlenecks for the whole network.

The amount of data moving through each interface is measured in megabytes. However, the following systems store data as packets rather than bytes:

- AIX
- FreeBSD
- IRIX
- MacOS
- Novell NRM

If you are monitoring one or more of these systems, you can specify a ratio for converting packets to bytes.

Different network interfaces have a maximum packet size called a Maximum Transmission Unit (MTU) - an ethernet interface, for example, has an MTU of 1,500 bytes. Most interfaces will not transmit packets at the MTU. The value that you specify for the bytes-per-packet conversion will be based on the observed performance of the network interface. Fifty percent of MTU is a good average to use - the default value in *up.time* is 750.

The report contains the following information:

- the display name in *up.time* of the system
- the names of each network interface on the system
- the total amount of data, measured in megabytes, that is moving in and out of each network interface

Generating a Network Bandwidth Report

To generate a Network Bandwidth report, do the following:

1. In the Reports Tree panel, click *Network Bandwidth*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
If no data available for the date range, the report displays a message indicating that there is no data for the time period.
3. To only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.

Daily Hours

Include data samples between these hours only:

Start: End:

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select 8:00 from the *Start* dropdown list, and 18:00 from the *End* dropdown list.

4. If you are monitoring systems that store network traffic data in packets rather than bytes, enter a conversion ratio in the *Bytes per Packet* field.
For example, you can specify a conversion ratio of 1,000 bytes per packet. The default is 750 bytes per packet.
5. To generate reports for groups of systems, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific Applications in your environment, select them from the *List of Elements*.
8. Select a report generation option. See [Report Generation Options](#) for details
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

The following is an example of a Network Bandwidth report:

Network Bandwidth Report			
Date Range: 2008-04-17 00:00:00 to 2008-04-17 10:59:19 between the time range: 00:00 to 23:59			
Bytes per packet: 750			
Hostname	Interface	Total MB In	Total MB Out
AIX DEV LPAR (10.1.1.57)	en0	2,186.97	2,125.59
AIX QA LPAR (10.1.1.56)	en2	417.78	336.75
AIX5 (aix5l)	(unknown)	0.00	0.00
	en0	223.55	126.35
ELinux (elinux)	eth1	0.00	0.00
	sit0	0.00	0.00
	eth0	14.49	11.50
ESX4 (vmh-esx4)	vmnic0	1,311.36	5,709.02
	vmnic1	1,422.36	435.30
ESX7 (vmh-esx7)	vmnic0	2,801.19	3,546.28
	vmnic1	0.00	84.88
Exchange (uptime-exchange)	netif0	9.66	9.66
	netif1	362.76	454.91
vmh-prod	vmnic0	55,057.18	77,215.26
	eth0	14,756.37	7,039.90
WebSphere (lab-websphere51)	netif0	932.94	932.94
	netif1	24.36	124.54

In this example, the system Filter has high levels of network traffic flowing in and out of a particular network interface. Based on this information, you can generate a Network graph (see [Network Graphs](#) for more information) to get a better idea of why network I/O is so high on the system.

Disk I/O Bandwidth Report

The Disk I/O Bandwidth report keeps track of the amount of data being read from and written to a disk on a system. The report can display the amount of data either as blocks or megabytes.

The report contains the following information:

- the display name of the system in *up.time*
- the names of each disk on the system
- where applicable, the name of the file system on the disk
- the total amount of data, measured in megabytes, that is being read from and written to the disk

Using Regular Expressions

You can use regular expressions to include or exclude disks and file systems when generating a [Disk I/O Bandwidth Report](#) (or a [File System Service Time Summary Report](#)), as shown below:

Exclude Disks: <div style="border: 1px solid #ccc; height: 40px; margin-top: 5px;"></div>	Exceptions: <div style="border: 1px solid #ccc; height: 40px; margin-top: 5px;"></div>
Note: you can enter regular expressions into these fields.	
Exclude File Systems: <div style="border: 1px solid #ccc; height: 40px; margin-top: 5px;"></div>	Exceptions: <div style="border: 1px solid #ccc; height: 40px; margin-top: 5px;"></div>
Note: you can enter regular expressions into these fields.	

Using regular expressions, you can focus on particular disks or file systems on a server and also decrease the length of your report.

The regular expression syntax used with the [Disk I/O Bandwidth Report](#) or a [File System Service Time Summary Report](#) is similar to that used with the File System Capacity Growth report. For example, if you are generating a report on an Oracle volume and only want to focus on five specific file systems, you can enter the regular expression `/u[0-4]` in the *Exceptions* field.

If, on the other hand, you are working with a UNIX system with multiple disks and want to focus on disks whose names start with `md1` but ignore those whose names start with `md2`, you can enter the regular expression `/md1.*` in the *Exceptions* field and `/md2.*` in the *Exclude Disks* field.

Generating a Disk I/O Bandwidth Report

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

1. In the Reports Tree panel, click *Disk I/O Bandwidth*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
If no data available for the date range, the report displays a message indicating that there is no data for the time period.
3. To only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section shown below:

Daily Hours

Include data samples between these hours only:

Start: 8:00 ▼ End: 21:00 ▼

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select `8:00` from the *Start* dropdown list, and `18:00` from the *End* dropdown list.

4. In the *Bytes per Block* field, specify the size of input and output blocks in bytes. The default is 512 bytes.
Optionally, click the *Output in MB* to display the I/O values in megabytes rather than blocks.
5. If you want to include or exclude certain disks, enter the following in the *Exclude Disks* and *Exceptions* fields:
 - The name of the disk.
 - A regular expression. See [Using Regular Expressions](#) for more information.
6. If you want to include or exclude certain file systems, enter the following in the *Exclude File Systems* and *Exceptions* fields:
 - The name of the file system.
 - A regular expression. See [Using Regular Expressions](#) for more information.
7. To generate reports for groups of systems, select the groups from the *List of Groups* area.
8. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
9. If you are generating reports for specific Applications in your environment, select them from the *List of Elements*.
10. Select a report generation option. See [Report Generation Options](#) for details.
11. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the Disk I/O Bandwidth Report

The following is an example of a Disk I/O Bandwidth report:

Disk I/O Bandwidth Report			
Date Range: 2008-04-17 00:00:00 to 2008-04-17 11:27:36 between the time range: 00:00 to 23:59			
Output displayed in megabytes, 512 bytes per block			
Hostname	Disk Name	File System	I/O Total
AIX DEV LPAR (10.1.1.57)	hdisk0		265 MB
AIX QA LPAR (10.1.1.56)	hdisk0		1,176 MB
AIX5 (aix51)	hdisk0		201 MB
	hdisk1		0 MB
	l		0 MB
ELinux (elinux)	hda	/boot	183 MB
Exchange (uptime-exchange)	2	E:	75,315 MB
	0	C:	100,095 MB
	1	D:	0 MB
WebSphere (lab-websphere51)	0	C:	208,791 MB
	1	D:	598 MB

In this example, the systems Brightmail and Weblogic Server have high levels of disk I/O. Based on this information, you can generate a Disk Performance Statistics graph (see [Generating a Disk Performance Statistics Graph](#) for more information) to get a better idea of why disk I/O is so high on the system.

CPU Run Queue Threshold Report

The CPU Run Queue Threshold report lists -- when a system's CPU reaches a high level of usage -- the number of jobs that were ready to run but waiting in a queue, as well as the amount of time they were waiting.

If the size of the run queue is appreciably larger than the number of available processors on a system, or the run queue is backlogged for long periods of time, you can conclude that the server is overloaded.

You can use this report to pinpoint servers that are overloaded using the following factors:

- the CPU is busier than a value that you specify
- the length of the CPU run queue is greater than the threshold that you specify

This report contains the following information:

- the display name of the system in *up.time*
- the number of CPUs on the system
- the run queue threshold
- the minimum, maximum, and average length of the run queue (i.e., the number of jobs waiting to be processed) over the period of time that you specify
- graphs that illustrate the number of minutes that the CPU run queue spent over the threshold
- optionally, a list of processes that were in the run queue during the time period that you specify

Generating a CPU Run Queue Threshold Report

To generate a CPU Run Queue Threshold report, do the following:

1. In the Reports Tree panel, click *CPU Run Queue Threshold*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
If no data available for the date range, the report displays a message indicating that there is no data for the time period.
3. To only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.

Daily Hours

Include data samples between these hours only:

Start: 8:00

End: 21:00

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select 8:00 from the *Start* dropdown list, and 18:00 from the *End* dropdown list.


4. In the *Max CPU (%)* field, specify the threshold for CPU usage.
CPU usage is considered critical when both the CPU usage and the length of the run queue exceed this threshold.
5. In the *Threshold* field, enter the number of queued up jobs that, when exceeded, is considered excessive.
Multiple CPUs are taken into account so that the defined threshold scales up with each additional CPU present on a monitored system.
6. Select any of the following statistics to include in the report:

- sys (CPU system time)
- usr (CPU user time)
- wio (CPU wait I/O time)

The statistics that you select will be added together and compared to the threshold that you specified in step 4. For example, to see when system time and user time are over 80%, select the *sys* and *usr* options and then enter *80* in the *Max CPU (%)* field.

- If you want to include a list of processes that are in the run queue in the report, click *Show Processes*.
- Click the *Maintain Graph Scale* option to keep the scale of the graphs in the reports consistent.

For example, if you have three systems, and one is 1,200 minutes over the threshold then scale of the graph is 1,200 for all of the graphs in the report.



- To generate reports for groups of systems, select the groups from the *List of Groups* area.
- To generate reports for one or more views, select the groups from the *List of Views* area. See [Working with Views](#) for more information about views.
- If you are generating reports for specific Applications in your environment, select them from the *List of Elements*.
- Select a report generation option. See [Report Generation Options](#) for details
- To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel. See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the CPU Run Queue Threshold Report

The following is an example of a CPU Run Queue Threshold report:

CPU Run Queue Threshold

Date Range: 2008-04-17 00:00:00 to 2008-04-17 12:02:02 between the time range: 00:00 to 23:59


Max CPU Percentage: 90.0%, Threshold 2.0 x number of CPUs

CPU Options: Usr, Sys, Wio

CPU Run Queue Depth

Hostname	# CPUs	Threshold	High	Low	Average	Sustained Minutes Above Threshold
AIX DEV LPAR (10.1.1.57)	1	2.00	3	1	2.17	5
AIX5 (aix5l)	1	2.00	4	0	2.20	15

Minutes over Threshold of 2.0 for AIX DEV LPAR (10.1.1.57)



In this example, the system is consistently over the run queue threshold that was specified when the report was defined. Based on this information, you can generate a CPU performance graph (see [See Monitoring CPU Performance](#) for more information) to get a better idea of why the system is exceeding the CPU run queue threshold.

File System Service Time Summary Report

The File System Service Time Summary report indicates which system disks (and file systems) are using an excessive amount of time to complete disk operations. This report helps you identify which systems may benefit from configuration changes (e.g., adding RAM, moving a file system to another hard disk, implementing a RAID).

The report contains the following information:

- the name of the systems for which the report has been generated
 - the names of the disks and file systems on the system
 - the high, low, and average service times for each disk or file system, measured in milliseconds
 - the nth percentile for each disk or file system (e.g., although a file system may have had a high service time of 100ms, its 95th percentile of 40ms means 95% of the service times were 40ms or lower)
- On a system with heavy disk usage, disks and file systems will be in the higher end of the percentile.

You can also sort the results in the report by one of six criteria that you can specify when defining the report.

Generating a File System Service Time Summary Report

To generate a File System Service Time Summary report, do the following:

1. In the Reports Tree panel, click *File System Service Time Summary*.
2. In the *Date and Time Range* area, select the dates and times on which to report.

For more information, see [Understanding Dates and Times](#)

If no data available for the date range, the report displays a message indicating that there is no data for the time period.

3. To only include data from certain hours during the day, select those hours from the dropdown lists in the *Daily Hours* section.

Daily Hours

Include data samples between these hours only:

Start: 8:00 End: 21:00

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select 8:00 from the *Start* dropdown list, and 18:00 from the *End* dropdown list.

4. Select one of the following options from the *Primary Sort by* dropdown list to sort the results that *up.time* returns:
- System Name
 - Disk
 - High Service Time (the default)
 - Low Service Time
 - Average Service Time
 - High Percentile
5. Select *Ascending* or *Descending* from the associated dropdown list.
6. Optionally, do the following:
- Select another sort criteria from the *Secondary Sort by* dropdown list.
 - Select *Ascending* or *Descending* from the associated dropdown list.
7. In the *Threshold* field, specify the threshold for file system service time.
Disk or file system service time is considered critical when it exceeds this threshold.
8. In the *Percentile* field, specify the percentage of time at which the service time for systems is below the threshold.
The default is 95, which is the lowest service time that is greater than at least 95% of all of the recorded values in the time range that you specified in step 2.
9. If you want to include or exclude certain disks, enter the following in the *Exclude Disks* and *Exceptions* fields:
- The name of the disk.
 - A regular expression. See [Using Regular Expressions](#) for more information.
You can enter one name or regular expression on a single line.
10. If you want to include or exclude certain file systems, enter the following in the *Exclude File Systems* and *Exceptions* fields:
- The name of the file system.
 - A regular expression. See [Using Regular Expressions](#) for more information.
You can enter one name or regular expression on a single line.
11. To generate reports for groups of systems, select the groups from the *List of Groups* area.
12. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
13. If you are generating reports for specific Applications in your environment, select them from the *List of Elements*.
14. Select a report generation option. See [Report Generation Options](#) for details
15. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the File System Service Time Summary Report

The following is an example of a File System Service Time Summary report:

File System Service Time Summary						
Date Range: 2008-04-17 00:00:00 to 2008-04-17 12:48:45 between the time range: 00:00 to 23:59						
Sorted by Descending High Service Time						
Showing Percentile: 95.0, Threshold 20ms						
System Name	Disk	File System	Service Time (milliseconds)			
			High	Low	Average	95th Percentile
lab-t1-4 (10.1.1.234)	c0t0d0	/	382.00	0.00	2.50	0.00
HP Integ (hpinteg)	c2t0d0		341.00	0.00	53.92	181.80
	c2t1d0		234.00	0.00	26.37	113.00
lab-t1-2 (10.1.1.232)	c0t0d0	/	185.00	0.00	119.71	173.35
The Vault (vault)	1	D:	171.00	0.00	5.66	50.00
	0	C:	160.00	0.00	7.05	44.20
AIX DEV LPAR (10.1.1.57)	hdisk0		41.00	3.00	6.51	10.00
QA RedHat Instance (qa1-rhes4-x86)	sda	/boot	29.00	0.00	3.14	9.00
qa-DC1 (10.1.0.98)	0	C:	24.00	0.00	8.66	12.00
MyMachine (dev-rmeloche)	0	C:	23.00	0.00	3.44	12.00

In this example, the disks on each system have high levels of service time, and they are in the highest percentile that exceeds the service time threshold.

Reports for Service Level Agreements

The following reports enable you to assess your organization's ability to meet, and diagnose failures in meeting service level agreements by summarizing compliance and reporting on compliance and non-compliance of an SLA's component objectives and services:

SLA Summary Report

The SLA Summary report shows whether an SLA's performance target is being met, whether performance—even through currently compliant with the defined target—may eventually fall short in the future, and how component SLOs contributed to performance. The report contains charts and a table that provide the following information:

- your defined service level target, and how closely the SLA was met over daily, weekly, or monthly intervals
- a trend line that indicates whether compliance is at risk of not being met on a future date
- an optional breakdown of how component SLOs contributed to the SLA not achieving 100% compliance

The report answers the following questions:

- Are we meeting our service targets? If we aren't, which areas of our infrastructure are failing?
- Are things getting better or worse?

For more information on SLA definitions, see [Service Level Agreements](#).

Creating an SLA Summary Report

To create an SLA Summary Report:

1. In the Reports Tree panel, click *SLA Summary*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Select a *Compliance Period* to report on.
4. Clear the *Display Outage Tables* checkbox if you want the report to display only outage graphs.
5. If you want to generate reports for one or more groups that include SLAs, select the groups from the *List of Groups* area.
6. To generate reports for one or more views that contain SLAs, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific Service Level Agreements, select them from the *List of SLAs*.
8. Select a report generation option. See [Report Generation Options for details](#)
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

SLA Detailed Report

In cases where an SLA compliance target is not being met, the SLA Detailed report breaks down both the outages of an SLA's component SLOs, and the outages of each SLOs component services. This report allows you to pinpoint when specific services experienced outages, assisting with further investigation.

The report answers the following questions:

- Were there any outages yesterday? If so, how long were they and on which systems did they happen?
- Which business users were affected by service outages?
- What kinds of transaction volumes are we processing?
- What are the most important things we can fix in order to meet our SLA targets?

For more information on SLA definitions, see [Service Level Agreements](#).

Creating an SLA Detailed Report

To create an SLA Summary Report:

1. In the Reports Tree panel, click *SLA Detailed*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Select a *Compliance Period* to report on.
4. Clear the *Display Outage Tables* checkbox if you want the report to display only outage graphs.
5. If you want to generate reports for one or more groups that include SLAs, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific Service Level Agreements, select them from the *List of SLAs*.
8. Select a report generation option. See [Report Generation Options](#) for details
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Reports for Availability

The following reports enable you to visualize the availability metrics for all your mission-critical Applications and your critical system services:

Application Availability Report

The Application Availability report tracks the availability of the Applications in your environment, as well as the monitors that are associated with the Applications. This report contains the following information:


- the name of the Application
- the service monitors that are associated with the Application
- the percentage of time that the Application and monitors are in OK, Unknown, Warning, and Critical states

For more information on Applications, see [Working with Applications](#).

Creating an Application Availability Report

To create an Application Availability report, do the following:

1. In the Reports Tree panel, click *Application Availability*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Click the *Show Details* option to generate a full listing of information about the availability of the Applications, which is broken down by individual Applications.
4. If you do not select this option, then a summary of the status of all Applications appears on a single line, as shown below:



5. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific Applications in your environment, select them from the *List of Applications*.
8. Select a report generation option. See [Report Generation Options](#) for details
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Incident Priority Report

The Incident Priority report provides information on the frequency, duration, and recovery time of critical-level events, and the overall reliability of your monitored systems. This information is presented for services that are associated with groups of Elements (whether a pre-defined group, or an manually selected list of individual Elements). Compared to the Service Monitor Outages report, the Incident Priority report, instead of providing an auditable list of outages, uses a comparative approach to indicate how efficiently systems are running in relation to each other, and furthermore, how efficiently problems are dealt with.

In order to report this efficiency, the following building blocks are available as elements in the report:

- Incidents: The total number of outages for all service monitors associated with selected Elements. Critical-level events for multiple service monitors that are associated with a single Element will each contribute to the incident count.
- Incident Top 20: The 20 systems with the highest incident counts for the given time period (incidents being the number of times service monitors associated with selected Elements were in a critical state).
- Total Downtime: The total amount of time that all service monitors associated with selected Elements were in a critical state. Multiple service monitors in a critical state that are associated with a single Element each contribute to the downtime total.
- Downtime Top 20: The 20 systems with the highest downtime totals for the given time period.
- Incident Priority Quadrant: A graph in which all selected Elements are placed on quadrants based on the total downtime, and number of incidents caused by their associated service monitors.

Note that, to provide clear results in the report, only service monitors that were manually assigned to, and are directly associated with, an Element are taken into account when downtime and incident counts are tallied. This means service monitors that may be automatically installed such as the Platform Performance Gatherer are not included; additionally, only an Application's status as a whole affects downtime and incident counts, but its component service monitors--both master and regular service monitors--do not.

Using downtime and efficiency counts, the Incident Priority report includes the following key elements:

- Mean Time Between Failure: The average amount of time that an Element's associated service monitors were all running (i.e., in non-critical states) over a given time period.
Elements whose associated service monitors experience no downtime are still included in the report, but will not include an MTBF count since they did not experience an incident during the time period.
- Mean Time to Repair: The average number of minutes any of an Element's associated service monitors were in a critical state over a given time period.
A service is considered repaired, or being repaired, when its status changes from critical to one of "MAINT", "UNKNOWN", "WARNING", or "OK".

For all report elements, a service monitor is considered to have reached a critical state--thus has caused an incident, is contributing to downtime, or is an ongoing failure--when it actually generates an alert. The period preceding the alert, during which rechecks are intermittently being performed to avoid a false positive, does not count. See [Understanding the Alert Flow](#) for information on rechecks leading to a generated alert.

Creating an Efficiency Report

To create an Efficiency report, do the following:

1. In the Reports Tree panel, click *Efficiency*.

2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
Service monitors that, based on the selected time range, are already in a critical state will be included in calculations for downtime, incident counts, and other report elements.
3. In the *Report Options* area, select the charts you want included in the report. See the previous section, [See Incident Priority Report.](#), for more information on the available charts.
4. In the *Report Options* section, select the level of granularity at which the information will be presented (i.e., daily, weekly, or monthly).
5. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific systems in your environment, select them from the *List of Elements*.
8. Select a report generation option. See [Report Generation Options](#) for details.
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Service Monitor Availability Report

The Service Monitor Availability report tracks the status of the services associated with the hosts in your environment. This report lists the percentage of time each service was in the following states over the time period that you specify: OK, Warning, Critical, Maintenance, or Unknown.

For more information on each status, see [See Understanding the Status of Services.](#)

Creating Service Monitor Availability Reports

To create Service Monitor Availability reports, do the following:

1. In the Reports Tree panel, click *Service Monitor Availability*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
4. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
5. If you are generating reports for specific systems in your environment, select them from the *List of Systems and Nodes*.
6. Select a report generation option. See [Report Generation Options](#) for details
7. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Service Monitor Outages Report

The Service Monitor Outages report lists all warning or critical events for services that have occurred over a specified time period. Use this report to determine the cause of a problem by analyzing the declining availability of a server or set of servers.

The Service Monitor Outages report contains the following information:

- the date and time at which metrics were gathered for each service
- the duration of the outage
- whether or not a notification was sent, or an action was taken
- the status of each service
- a short message about the status - for example:
UPTIME-filter - up.time agent running on filter, up.time agent 3.9 solaris 1.17

Creating a Service Monitor Outages Report

To create a Service Monitor Outages report, do the following:

1. In the Reports Tree panel, click *Service Monitor Outages*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Select one of the following options from the *Sort by* dropdown list:
 - Sample Time by Element
 - Service Name by Element
 - All Sample Times
4. From the *Sort Direction* dropdown list, select *Ascending* or *Descending*.
5. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific systems in your environment, select them from the *List of Elements*.
8. Select a report generation option. See [Report Generation Options](#) for details.
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Reports for J2EE Applications

The following reports enable you to visualize any performance problems with applications that are running a J2EE environments:

WebSphere Report

The WebSphere report charts a set of counters that provide insight into the health and performance of a WebSphere Application Server. Depending on the number of options that you select, the report can become quite long and can take considerable time to generate. For most options, the report contains charts for two or more metrics.

Creating a WebSphere Report


To create a WebSphere report, do the following:

1. In the Reports Tree panel, click *WebSphere*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Select one or more of the following report options:
 - Thread pool
A set of counters that report on the number of connection threads that have been created or destroyed, that are concurrently active or are hung, that are in the thread pool, or time that are in use.
 - JDBC Connection Pool
A set of counters that monitor the performance of JDBC data sources.
 - Enterprise Beans
A set of counters that report the following: load values, response times, and life cycle activities for enterprise Java beans.
 - JVM Runtime
A set of counters that monitor the performance of the Java Virtual Machine (JVM) that is running on the WebSphere server.
 - Transaction Manager
A set of counters that report on the status of global, local, and concurrent transactions.
 - Servlet Session Manager
A set of counters that report on usage information from the HTTP servlets that are running on the server.Optionally, click *Select All* to generate a report on all of the options listed above
4. If you selected more than one report option and plan to report on more than one system, you can optionally click the *Group report options by system* checkbox.
Selecting this option combines the metrics for each system for which you are generating the report.
5. To generate reports for systems in specific groups, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific systems, select the systems from the *List of Systems*.
8. Select a report generation option. See [Report Generation Options](#) for details.
9. If you want to save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the WebSphere Report

Since WebSphere is large and complex, it can be difficult to pinpoint the source of a problem with the server or an application running on the server. This is especially true when that problem is intermittent. Watching for problems in real time only gives you a snapshot of the problem. The *up.time* WebSphere report, on the other hand, gives you a detailed historical perspective of the problem. Using the information in the report, you can find the source of the problem.

For example, users have trouble working with an application that intensively uses a database. Checking the *Connection Pool* charts section of a WebSphere report could indicate the source of the problem - the database has reached its maximum number of connections.



You can then adjust the size of the database connection pool to allow more connections.

Or, if a WebSphere application is using a large amount of memory you could check the *JVM charts* section of the report. If there are spikes in the heap size or memory usage of the JVM, you can tune the JVM to ensure that it is working at optimal levels.

WebLogic Report

The WebLogic report charts a set of metrics (see [WebLogic](#) for details) that provide insight into the health and performance of a WebLogic server. Using the WebLogic report, you can pinpoint problem areas on your WebLogic server and quickly determine how to fix those problems.

Depending on the number of options that you select, the report can become quite long and can take considerable time to generate. For most options, the report contains charts for two or more metrics.

Creating a WebLogic Report

To create a WebLogic report, do the following:

1. In the Reports Tree panel, click *WebLogic*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
3. In the *Report Options* area, select one or more of the following options:
 - **Thread pool**
The report charts the number of pending request in the thread pool, as well as the free size of the pool.
 - **Server Stats**
The report charts the number of connection requests that WebLogic accepts before refusing additional requests, as well as the number of open sockets to the server.
 - **JDBC Connection Pool**
The report charts the number of active and leaked connections to the server, as well as the size of the connection pool, the number of connections that are waiting or delayed, and the number of failures to reconnect to the server.
 - **Enterprise Beans**
The report charts the number of Enterprise Java Beans (EJB) that are active or have been moved to secondary storage, the number of time that a container can and cannot find an EJB in the cache, as well as the total number of EJBs in the cache.

This report returns information for:

- *Stateful EJBs*, which hold data for a client between calls to the EJB. Stateful EJBs can use considerable amount of server resources.
- *Stateless EJBs*, which hold data for only one call to the EJB, and then deletes that data. Stateless EJBs use fewer system resources than stateful EJBs
- **JVM Runtime**
The report charts the heap size (in kilobytes) of the Java Virtual Machine (JVM) on the WebLogic server, as well as amount memory (in kilobytes) available to the JVM.
- **Transaction Manager**
The report charts the number of transactions that were committed or completed successfully, as well as total number of transactions that are rolled back.
- **Servlets**
The report charts the number of requests that were made to the HTTP servlets that are running on the WebLogic server.

Optionally, click *Select All Options* to use all of the options that are listed above.


4. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
5. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
6. If you are generating reports for specific systems in your environment, select them from the *List of Systems and Nodes*.
7. Select a report generation option. See [Report Generation Options](#) for details
8. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the WebLogic Report

Since WebLogic is large and complex, it can be difficult to pinpoint the source of a problem with the server or an application running on the server. This is especially true when that problem is intermittent. Watching for problems in real time only gives you a snapshot of the problem. The *up.time* WebLogic report, on the other hand, gives you a detailed historical perspective of the problem. Using the information report, you can find the source of the problem.

For example, users have trouble logging into an application that is running on the WebLogic server. Checking the *Connection Pool charts* section of a WebLogic report, you might see that the size of the connection pool has reached its maximum, and that there are a large number of connections that are waiting in the pool. From there, you can then adjust the size of the connection pool to allow more connections.

Or, if a WebLogic application is using a large amount of memory you could check the *JVM charts* section of the report.



If there are increases or sudden spikes in the heap size or memory usage of the JVM, then you can tune the JVM to ensure that it is working at optimal levels.

Reports for Virtual Environments

Virtualization platforms such as VMware vSphere enable you to consolidate servers and applications in a virtual environment. Using virtual machine managers such as VMware ESX, you can run multiple servers or applications on a single system, but without using as much hardware. Each server or application runs in its own VMware instance. You can use VMware vSphere to manage and monitor ESX servers, as well as allocate resources among virtual machines.

up.time's VMware and pSeries reports enable you to visualize the performance of systems that are consolidated on virtual machines, whether using VMware or IBM pSeries Logical Partitions (LPARs):

VM Sprawl Report

The VM Sprawl report helps you assess the extent of sprawl across your virtual infrastructure, and provides you with the information needed to reduce it. Using the report, you can perform the following types of tasks:

- see which VMs are underused, and which VMs' tasks should be consolidated elsewhere
- create an inventory of VMs that are always off or suspended, thus were unnecessarily created
- see which VMs have not been powered on recently
- audit or verify VM creation and destruction balance
- plot VM population trends to see if sprawl is growing or being reduced

Depending on how it is configured the VM Sprawl report consists of at least one of the following components:

1. monthly VM population trends for up to the last year
For each month, stacked bars show the following VM totals:
 - created during that month and are currently running
 - created during that month and are currently suspended or powered off
 - destroyed during that month

- created and destroyed during that same month
2. a list of VMs created during the time period
 3. a list of VMs destroyed during the time period
 4. lists summarizing various VM power state scenarios:
 - existing VMs created that were powered off or suspended for 20% or more of the time, during the time period
 - existing VMs that were not powered on during the time period
 - existing VMs that were in a suspended state for the time period

Creating a VM Sprawl Report

To create a VMware vSphere Workload report, do the following:

1. In the Reports Tree panel, click *VM Sprawl*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. In the *Report Options* section, select which types of VM population totals you would like to include in the report.
4. Select the Elements whose child-object VMs you would like to include in the report totals.
In most cases, for the greatest visibility with a high-level view, the selected Element will most likely be a VMware vCenter server or datacenter Element.
5. Select a report generation option. See [Report Generation Options](#) for details.
6. Do one of the following:
 - Click the *Generate Report* button.
 - Enter a name for the report in the *Save to My Portal As* field, and optionally enter text in the *Report Description* field. Then, click *Save Report*.

The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
7. To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Scheduling Reports](#) for more information on configuring a scheduled report.

vSphere Workload Report

The VMware vSphere Workload report provides a broad view of workloads across your entire virtual infrastructure. For selected Elements, it includes a resource usage overview (CPU, memory, disk, and network). It also includes detailed resource usage charts for selected Elements' child objects in the vSphere hierarchy.



For reporting periods, total resource usage is reported regardless of how the VMware vSphere object's child objects change during that time period. For example, a four-week overview chart for an ESX server will include performance of VMs that have since migrated.

Using this report, you can visualize resource usage at the datacenter, cluster or ESX server level, as well as with resource pools, vApps, and virtual machines. You can also understand how these Elements' respective child objects contribute to their usage levels.

vSphere Workload Report Metrics

The following metrics can be displayed in a vSphere Workload report:

CPU Usage	the following is shown for the interval: <ol style="list-style-type: none"> 1. for ESX servers, the total CPU capacity, in GHz 2. the current amount of CPU resources, in GHz, used by ESX hosts or VMs 3. the total CPU capacity of the Element
Memory Usage	the following is shown for the interval: <ol style="list-style-type: none"> 1. for ESX servers, total memory capacity, in GB 2. the amount of memory, in GB, used by ESX hosts or VMs 3. the total memory capacity of the Element
Disk I/O	the rate, in MBps, at which ESX hosts or VMs are reading and writing data to and from disk
Network I/O	the rate, in Mbps, at which ESX hosts or VMs are receiving or transmitting data over the network
% 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

Creating a vSphere Workload Report

To create a VMware vSphere Workload report, do the following:

1. In the Reports Tree panel, click *vSphere Workload*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#)
3. Select the Elements you would like to include in the vSphere Workload report from the following sections where applicable:
 - List of vCenter Servers


- List of Datacenters
 - List of Clusters
 - List of ESX Servers
 - List of Resource Pools
 - List of vApps
 - List of Virtual Machines
4. Select which *vSphere Element Summaries to Display* .
For each of the Elements you selected in the previous step, you can display resource usage summaries of their VMware vSphere child objects in the detailed sections of the report.
For example, if you select a VMware vCenter server Element to be in the report, selecting the *Datacenter* and *Cluster* summary options will include summaries for datacenters and clusters that are child objects of the vCenter server. Note that selecting a VMware vSphere child object will not automatically include its own child objects (e.g., ESX servers that are part of a cluster); a selection exclusively includes that specific object. Only VMware vSphere child objects that are relevant to the parent will be shown:
 - if you select the *Datacenter* summary option, VMware vCenter server Elements included in the report will include resource usage summaries for its child datacenter objects
 - if you select the *Cluster* summary option, VMware vCenter server and datacenter Elements in the report will include resource usage summaries for their child cluster objects
 - if you select the *ESX Server* summary option, datacenter and cluster Elements in the report will include resource usage summaries for their child ESX server objects
 - if you select the *Resource Pool* or *vApp* summary options, any datacenter, cluster, or ESX server Elements included in this report will include resource usage summaries of its top-level resource pools and vApps
 - if you select the *Resource Pool* or *vApp* summary options, any parent resource pool or vApp Elements included in this report will include resource usage summaries for their child resource pools and vApps
 - if you select the *Virtual Machine* summary option, ESX servers, resource pools, and vApps included in this report will include resource usage summaries for their child VMs
 - if you select the *Virtual Machine* summary option, datacenters and clusters included in this report will include percent-ready and percent-wait summaries for its child VMs
 5. In the *Report Options* section, select whether to *Show Overview Charts* in the report.
The overview charts consist of CPU, memory, disk, and network usage overviews for the defined time range.
 6. Select a report generation option. See [Report Generation Options](#) for details.
 7. Do one of the following:
 - Click the *Generate Report* button.
 - Enter a name for the report in the *Save to My Portal As* field, and optionally enter text in the *Report Description* field. Then, click *Save Report* .
The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
 8. To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Scheduling Reports](#) for more information on configuring a scheduled report.

VMware Workload Report

A VMware server often slows down because an instance on the server is consuming large amounts of such system resources as CPU, disk I/O, and memory. The problem could lie with an instance that is currently slow or another instance on the same server.


The VMware Workload report charts the workload of both the server on which VI3 or vSphere 4 is running, and the ESX servers that it is managing. It does this by graphing the key performance counters the *up.time* collects from VI3 or vSphere 4.

You can also use the VMware Workload report to determine whether or not you are using a particular VMware server to its optimal capacity. The VMware Workload report can be a useful tool for determining whether or not a VMware server is being used to its optimal capacity. Consider the following example, in which the VMware Workload report returns the following information about the top ten CPU loads on the VMware server:



This graph indicates that, on average, the ten most CPU-intensive instances use only 20% of the server's CPU capacity. The CPU on the server can handle up to three to four times its current load.

The memory usage section of the report indicates that the instances are using roughly the same amount of memory:



The server appears to have an ample amount of memory available.

The report indicates that you can add more instances to the VMware server.

Creating a VMware Workload Report

To create a VMware Workload report, do the following:

1. In the Reports Tree panel, click *VMware Workload*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
3. In the *Report Options* section, select one of the following:
 - **Workload Profile - CPU**
The percentage of CPU time that is being used by a VMware instance. This is a percentage of the available maximum amount of CPU time. This ensures that all of the CPU usage figures add up to the overall CPU usage of the server.
 - **Workload Profile - Memory**
The amount of physical memory, in kilobytes, that is being used by a VMware instance.
 - **Workload Profile - Disk IO**
The amount of the disk I/O capacity, in kilobytes per second, that is being used by a VMware instance.

- Workload Profile - Network IO
The amount of the network I/O capacity, in kilobits per second, that is being used by a VMware instance.
 - Workload Profile - % Ready
The amount of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server.
 - Workload Profile - % Used
The percentage of CPU time that an instance running on an ESX server is using.
4. If you want to generate reports for systems in specific groups, select the groups from the *List of Groups* area.
 5. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
 6. If you are generating reports for specific systems in your environment, select them from the *List of Elements*.
 7. Select a report generation option. See [Report Generation Options](#) for details.
 8. Do one of the following:
 - Click the *Generate Report* button.
 - Enter a name for the report in the *Save to My Portal As* field, and optionally enter text in the *Report Description* field. Then, click *Save Report*.
The report parameters are saved to the *My Portal* panel. Doing this does not generate the report.
 9. To schedule the saved report to run at a specific time or interval, click the *Scheduled* checkbox.
See [Scheduling Reports](#) for more information on configuring a scheduled report.

VMware Infrastructure Density Report

The VMware Infrastructure Density report enables you to assess the carrying capacity and workload distribution of your ESX infrastructure. To accomplish this, virtual machine counts are tracked and reported on a daily basis, where the peak VM count for a given day is used as that day's tally. The information available in the report includes the following:

1. Virtual Infrastructure Density: The total number of virtual machines in relation to the total number of ESX servers over a given time period. A trend line is mapped onto the totals, indicating whether VM counts, and corresponding workloads, are increasing or decreasing in relation to available ESX server capacity.
2. Total Virtual Machine Count: The total number of virtual machines running on all, or a group of, ESX servers. The VM totals are separated into individual ESX server totals.
3. ESX Server Virtual Machine Count: The total number of virtual machines running on a specific ESX server.

Using this report, you can have a better understanding of virtualized workloads by seeing ESX server use and trends, and quantifying VM creation overall, and on a server-by-server basis.

Creating a VMware Infrastructure Density Report

To create a VMware Infrastructure Density report, do the following:

1. In the Reports Tree panel, click *VMware Infrastructure Density*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
3. In the *Report Options* section, indicate whether you want to *Include Charts for Individual ESX Servers* by selecting or clearing the check box. When this option is enabled, a separate chart with VM counts will be created for each ESX server that is included in the report.
4. In the *Report Options* section, select the level of granularity at which the virtual infrastructure density information will be presented (i.e., daily, weekly, or monthly).
5. If you want to generate reports for groups of systems, select the groups from the *List of Groups* area.
6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
7. If you are generating reports for specific systems in your environment, select them from the *List of Systems and Nodes*.
8. Select a report generation option. See [Report Generation Options](#) for details.
9. To save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

LPAR Workload Report

The LPAR Workload report charts the workload of the individual logical partitions (LPARs) on an IBM pSeries server. It does this by graphing the following workload data:

- CPU
- Memory
- Network I/O
- Disk I/O

Using the information in the report, you can gain insight into the overall workload on an IBM pSeries server. This enables you to accurately adjust the CPU entitlements of the LPARs and keep track of the overall workload over time.

Creating an LPAR Workload Report

To create an LPAR Workload report, do the following:

1. In the Reports Tree panel, click *LPAR Workload*.
2. In the *Date and Time Range* area, select the dates and times on which to report.
For more information, see [Understanding Dates and Times](#).
3. Select one or more of the following report options:
 - CPU Workload
The CPU entitlements of the LPARs, and their use of the entitlements.
 - Memory Workload
The amount of memory, in kilobytes, that is being used by the LPARs on the system.

- Disk IO Workload
The amount of data, measured in kilobytes per second, that is being read from and written to the disk by the LPARs on the system.
 - Network IO Workload
The amount of data, measured in kilobytes per second, that is being sent and received over the network interface by the LPARs on the system.
Optionally, click *Select All* to generate a report on all of the options that are listed above.
4. If you selected more than one report option and plan to report on more than one system, you can optionally click the *Group report options by system* checkbox.
Selecting this option combines the metrics for each system for which you are generating the report.
 5. To generate reports for systems in specific groups, select the groups from the *List of Groups* area.
 6. To generate reports for one or more views, select the groups from the *List of Views* area.
See [Working with Views](#) for more information about views.
 7. If you are generating reports for specific systems, select the systems from the *List of Systems* .
 8. Select a report generation option. See [Report Generation Options](#) for details.
 9. If you want to save the report or schedule it to run at a specific time or interval, complete the settings in the *Save Reports* section of the subpanel.
See [Saving Reports](#) and [Scheduling Reports](#) for more information.

Using the LPAR Workload Report

The LPAR Workload report takes the guesswork out of determining CPU entitlements for the LPARs on a pSeries server. The entitlements indicate the amount of CPU power that is assigned to an LPAR.

For example, you have an LPAR with hard entitlement (one that cannot use spare processing power from another CPU on the server) and its CPU usage is constantly at or near the maximum. In this case, you can either increase the CPU entitlement of the LPAR, or change it to a soft entitlement.

If, on the other hand, the LPAR has a soft entitlement (one which can use spare processing power from another CPU on the server) and its CPU usage is consistently at or greater than the entitlement, you can increase it.