up.time User Guide
version 5.5

up.time software
because downtime is not an option

Contacting uptime software

By mail:
uptime software inc.
555 Richmond Street West,
PO Box 110
Toronto, Ontario
Canada
M5V 3B1

Telephone: 416-868-0152
Fax: 416-868-4867

Contacting Sales

To contact sales, use the main telephone line: +1-416-868-0152, and follow the prompts.

Please have the following information available so we may serve you better:

- Operating systems
- Key applications and databases
- Deployment Timeframe
- Project to deploy
- Key problems
- Present tools

Contacting Support

uptime software delivers responsive customer support. Customer support is available to licensed and demonstration users.

uptime software offers user support through the following:

- Documentation
- Application
- Telephone
- E-mail
- Internet site

Before contacting support, consult the up.time User Guide, up.time Release Notes, or the help system from the Help button in the application.

To contact sales, use the main telephone line: +1-416-868-0152, and select option #2.

TABLE OF CONTENTS

Welcome to up.time

Understanding up.time

Installing the up.time Enterprise Monitoring Station

Getting Started

Using My Portal

Defining and Managing Your Enterprise

Overseeing Your Enterprise

Working with Service Level Agreements

Configuring Users

Alerts and Actions

Understanding Graphing

Using Graphs

Configuring and Managing up.time

Reference

End User License Agreement

Index

CHAPTER 1

Welcome to up.time

This chapter introduces up.time in the following sections:

Introducing up.time

up.time monitors, manages, and reports on systems, network devices, and applications in single datacenters, and across multiple datacenters in a real-time, centralized view.

At the datacenter level, up.time continuously monitors your servers, applications, databases and IT resources, and alerts you to problems. Using the information that up.time gathers, you can solve problems before they impact your business.

For example, a service monitor detects that a large volume of email messages are going back and forth between a particular email address in your organization and an external domain. This could indicate that a high number of legitimate emails are being sent, or it could indicate that a virus or a trojan is active on a system in your environment.

You can also generate reports and graphs to visualize the information that up.time gathers. By analyzing the information, reports, and graphs you can do the following:

- identify and isolate performance bottlenecks
- monitor and report on the availability of services
- determine the specific causes of a problem in your network
- perform capacity planning
- consolidate servers where necessary
- develop more precise management reports

At the enterprise level, up.time enables the same monitoring, alerting, and reporting across multiple datacenters, accommodating organizations whose essential IT services are globally distributed. While the aforementioned tasks are still performed on specific datacenters' infrastructure—allowing datacenter-specific system management—managers of global IT services can perform similar, but larger scoped tasks using aggregated datacenter metrics in a central view.

As IT assets continue to be globalized while the management of these assets continues to be centralized, up.time allows managers to perform vital, enterprise-scale, systems management tasks:

- monitor global IT assets and identify datacenter-level bottlenecks

- perform capacity planning and server consolidation from a global perspective
- maintain visibility on hardware assets that are outsourced to managed service providers (MSPs)
- selectively view, monitor, and report on either deep or broad portions of the entire enterprise, from perspectives that differ from manager to manager
    - local datacenter views
    - global, enterprise views
    - logical views in a global deployment (e.g., all database management servers in the organization)
    - logical groupings of global assets (e.g., all applications, servers, and network devices associated with an enterprise application)

Who Should Read This Guide

The up.time User Guide is intended for various types of users:

- system administrators who want to use up.time to monitor a single system or multiple systems in a distributed environment
- owners of global enterprise applications whose associated servers and network devices—regardless of location—need to be selectively monitored
- users who gather information about their systems to perform analysis and make key business decisions
- IT managers who will determine the availability of resources, applications, and data for their user community

up.time Architecture

up.time consists of an Enterprise Monitoring Station that retrieves information from Local Datacenters, whose own Monitoring Stations are collecting data from . The Enterprise Monitoring Station aggregates data from each Local Datacenter by linking with each LDC database, retrieving relevant metrics, and recording them to its own database. The following diagram illustrates the general architecture of up.time:

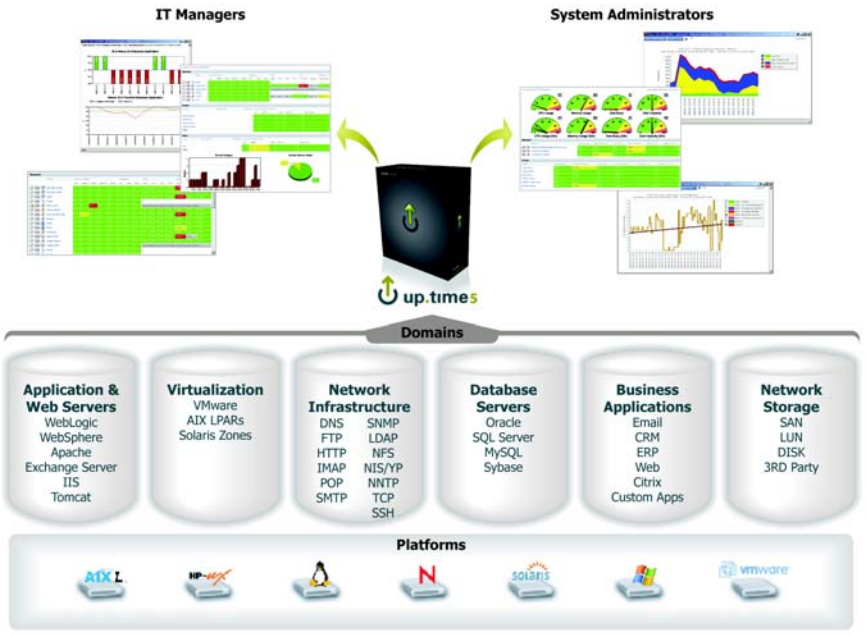

up.time Service Monitoring Concepts

Before you start using up.time, you should first understand the underlying service monitoring concepts.

- Alert Profiles

  Templates that tell up.time exactly how to react to various alerts – issuing alert notifications and performing recovery options – generated by your service checks.

- Monitoring Periods

  Specific windows during which you want to have up.time generate and send alert notifications. For example, you can specify that alerts only be sent between 9 a.m. and 5 p.m. on weekdays.

- Monitor Escalations

  The exact definitions of when and how up.time should escalate service alerts if they have not been acknowledged by specific users within pre-defined time limits.

CHAPTER 2

Understanding up.time

This chapter explains underlying concepts in the following sections:

Understanding the up.time Interface

The up.time Web interface consists of six or seven main sections, depending on whether you are logged in to the Enterprise Monitoring Station or a Local Datacenter. The following image displays the up.time application screen for the EMS. The panels change according to the task area that is selected from the tool bar.

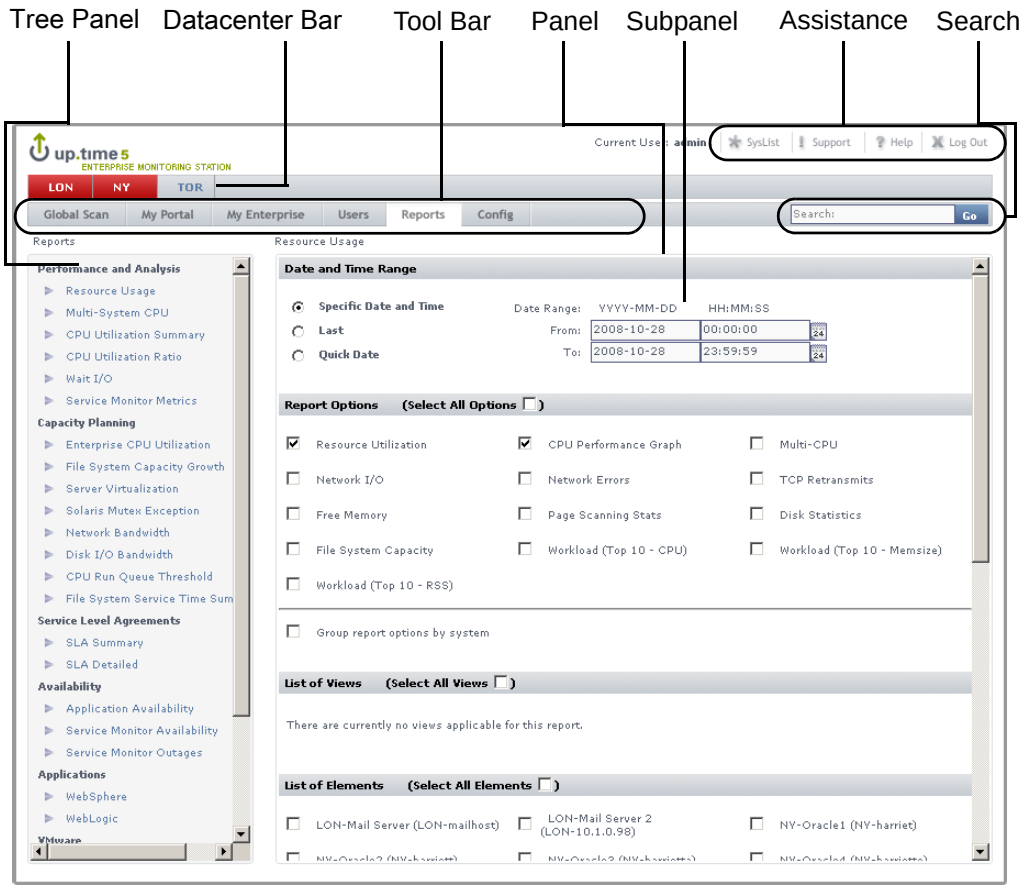

up.time Tool Bar

The up.time tool bar provides quick access to the following panels:

- Global Scan
- My Portal
- My Enterprise
- Users
- Reports
- Config

Global Scan

The **Global Scan** panel provides information about the status of your Datacenters, and EMS-based Applications and SLAs. You can drill down by Datacenter, Replication Group, system group, replicated system, or alert status to manage the resources in your enterprise.

For more information about using the **Global Scan** panel, see “Overseeing Your Enterprise” on page 75.

My Portal

When you log into up.time, the first screen you see is the **My Portal** panel. The **My Portal** panel gives quick access to basic up.time functions and to saved reports. The **My Portal** panel is divided into the following sections:

- Assistance
- My Preferences
- Latest News
- My Reports

For more information about using the **My Portal** panel, see “Using My Portal” on page 47.

My Enterprise

The **My Enterprise** panel provides an inventory of your network resources. You can view information about systems and their monitoring status. From the **My Enterprise** panel, you can add and view:

- Datacenters
- Groups
- Applications
- Service Level Agreements
- Views

For more information about using the **My Enterprise** panel, see “Defining and Managing Your Enterprise” on page 51.

Users

The **Users** panel enables you manage all users, user groups, Notification Groups and their associated permissions. You can view, create, edit, and delete the following:

- users
- user groups
- Notification Groups
- user roles

For more information about using the **Users** panel, see “Configuring Users” on page 117.

Reports

The **Reports** panel enables you to manage and create detailed, custom reports on the performance and availability of the resources in your enterprise.

Using the **Reports** panel, you can:

- generate a report and schedule when you want it to be generated
- select how and where you would like the report delivered

For more information about using the **Reports** panel, see “Using Reports” on page 175.

Config

The **Config** panel enables you to configure the following:

- up.time license information and the license key
- archive policies
- mail servers
- Monitoring Periods
- remote reporting instances
- user authentication

You can also generate problem reports and edit the `uptime.conf` file from the **Config** panel. For more information about using the **Config** panel, see “Configuring and Managing up.time” on page 289.

Icons

Entries in various panels have icons beside them. These icons enable you to perform the following tasks:

-  Edit

  Opens a window in which you can modify any entry in a panel.

-  View

  Displays the properties of any entry in a panel.

-  Delete

  Deletes any entry in a panel. You will need administrator privileges to delete certain entries.

These icons do not appear in the up.time Web interface if users do not have permissions to access the functions represented by the icons.

Understanding Reports and Graphs

up.time includes a powerful set of reporting and graphing tools that enable you to visualize performance data. You can use the reports and graphs as the starting point when analyzing problems in your environment.

Understanding Reports

Reports enable you to visually analyze how individual critical resources—such as memory, CPU, and disk resources—are being consumed over specific period of time.

For detailed information about reports, see “Using Reports” on page 175.

If you need to regularly run certain reports, you can save them to the **My Portal** panel. See “Scheduling Reports” on page 169 for more information.

Understanding Graphs

You can graph performance information when you need to view the most common or pertinent performance information for servers in your environment. For example, you can use a graph to determine CPU usage or the available capacity on a file system. Graphs give you a fine level of performance detail.

You can view graphs in two ways:

- With Internet Explorer in Microsoft Windows. Graphs are rendered using an ActiveX graphing control. You can edit and manipulate a graph once it has been displayed, and you can create trend lines.
- Using the Java graphing tool on any platform (e.g., in Firefox, running on Linux).

For more information on graphing, see “Understanding Graphing” on page 241 and “Using Graphs” on page 249

Understanding the up.time DataStore

The DataStore is a database in which up.time stores different types of information:

- configuration information for up.time
- configuration information for the Datacenters that you are monitoring
- the performance data replicated from Datacenters, which is used for generating graphs and reports
- user information, including user names and passwords (encrypted if it is sensitive information)
- reports that Enterprise Monitoring Station users have saved, and are scheduled to run at specific intervals.

Like any other database, the DataStore consists of a number of tables. Data that you enter and save, or which up.time collects from hosts, is written to specific tables in the DataStore.

Access to the DataStore is determined by one of the three installed user accounts: root, uptime, and reports. Each account gives users varying levels of access to the contents of the DataStore.

Understanding the Status of Services

up.time monitors can return the following statuses for a service:

- `0` – OK

  The services are functioning properly.

- `1` – Warning

  There is a potential problem with one of more of the services.

- `2` – Critical

  There is a critical problem with one or more services.

- `3` – Unknown

  This status is returned when:

  - The host on which the service sits is offline.
  - The host on which the service sits is in a scheduled maintenance or downtime period.
  - The Monitoring Station could not execute the service monitor.

Each status reflects the state of the service that has been assigned to the system that you are currently viewing. up.time picks up these error codes and triggers an alert or an action. This status reporting, and subsequent alerting can be acknowledged if it's tied to an EMS-based element (i.e., Datacenters, SLAs, or Applications); otherwise, alerting for Local-Datacenter-based services will be acknowledged and managed by LDC-based administrators.

The status of the services associated with a system are displayed in the **Global Scan** panel, as shown below:

Name	Service Status					Outages			CPU				Disk		Memory
	OK	WARN	CRIT	MAINT	UNKN	1hr	12hr	24hr	USR	SYS	WIO	TOT	% Used	% Busy	Swap Used
LON-Mail Server	2	0	0	0	0	0	0	0	0%	1%	0%	1%	96%	1%	2%
LON-Mail Server 2	2	0	0	0	0	0	0	0	1%	0%	0%	1%	90%	0%	28%
TOR-ESX4	1	0	0	0	0	0	0	0	20%	0%	0%	20%	100%	0%	0%
TOR-ESX7	1	0	0	0	0	0	0	1	26%	0%	0%	26%	100%	0%	0%
TOR-WebSphere	0	0	1	0	3	0	0	0	The availability check has failed.						

The figures in each column in the Global Scan panel indicate the number of services for that particular machine that are in each state. Click a number to

view the **System Status** screen for a particular system. See “Viewing the Status of a System” on page 251 for more information.

Understanding Dates and Times

When you are configuring graphs or reports, you must specify a range of dates and times over which the graph or report will chart information. up.time will only collect information for the periods that you specify.

You specify data and time ranges in the **Date Range** area of the **Reports** and **Graphing** subpanels, as shown below:

To set dates and times for a graph or report, do one the following:

- Click the **Specific Date and Time** option. Then, in the **Date Range** area, select the start date and time of the report by:
  - entering the start and end times (HH:MM:SS) in the **From** and **To** text boxes
  - entering the start and end dates (YYYY-MM-DD) in the **From** and **To** text boxes

  You can also click the calendar icon ( ) to select dates.

- Click the **Last** option, then do the following:
  - select a number from 1 to 10 from the first dropdown list
  - select `Days`, `Weeks`, or `Months` from the second dropdown list

    The end date for any of these options is the current date and time. For example, if you select `1` and `Days`, then the graph or report will cover the 24 hour period from the previous day until the date and time on which you created the report.

- Click the **Quick Date** option, and then select one of the following options from the dropdown list:
  - Today
  - Yesterday
  - This Week
  - Last Week (Sun-Sat)
  - This Month
  - Last Month

The **This Month** option collects information from the first day of the current month to the day on which the report or graph is being generated. The **Last Month** option collects information from the beginning to the end of the previous month.

Understanding Retained Data

up.time enables you to save some or all of the metrics that its monitors collect to the DataStore. You can use the retained data to generate a Service Metrics report (see “Service Monitor Metrics Report” on page 187) or a Service Metrics graph (see “Viewing System and Service Information” on page 38).

The data that you can retain varies from monitor to monitor. For example, with the Windows Service Check monitor you can save the Service Status and Response Time metrics. With the Exchange monitor you can save all Web Mail and SMTP metrics.

You can save data to the DataStore by clicking the **Save for Graphing** checkbox on a monitor template, as shown below:

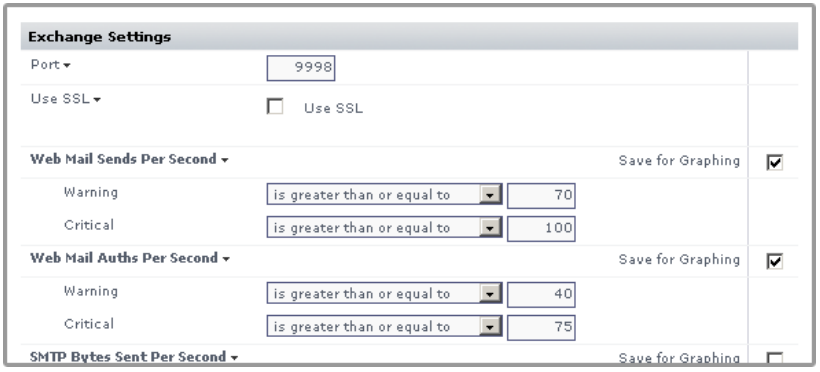

CHAPTER 3

Installing the up.time Enterprise Monitoring Station

This chapter explains how to install up.time in the following sections:

Installation Plan

Before installing up.time you must:

- identify the system that will act as a central Enterprise Monitoring Station
- ensure that all Local Datacenters from which you will aggregate data have been set up, and are accessible over the network
- ensure that all systems on which Monitoring Stations are running (including the EMS system) are synchronized with Internet time servers to avoid time drift across the multi-datacenter deployment

If you purchased the boxed version of up.time, the Enterprise Monitoring Station system must have a CD-ROM drive from which to load the server software. A CD-ROM drive is not required if you have downloaded the up.time software from the Internet.

The installation procedure creates the user ID `uptime` on the Enterprise Monitoring Station. The `uptime` user ID should also exist on all of the clients, as using this ID will minimize any security risks by not running the agents as a privileged process.

Wherever possible, do not use the `root` account to run the Enterprise Monitoring Station.

You can use other existing user accounts for the agent, such as `nobody`, `bin`, or `adm`. However, using these accounts may pose security risks depending on other system processes that run under these accounts.

Installation Requirements

This section describes the system requirements for the up.time Enterprise Monitoring Station. Before installation, it is recommended that you check the uptime software Web site (http://www.uptimesoftware.com) for the most up-to-date list of hardware and software requirements.

up.time Enterprise Monitoring Station

EMSThe up.time Enterprise Monitoring Station is a computer running the core up.time software that retrieves information from other up.time Monitoring Stations. The EMS has a self-contained Web server and database that enables easy access to the application and data.

The Enterprise Monitoring Station can run on the operating systems listed below. You should refer to the uptime software Client Care Web site for the most up-to-date list of supported platforms.

Operating System	Version(s)
Microsoft Windows Server 2008	Standard or Enterprise R2 (with 32-bit execution)
Microsoft Windows Server 2008	
Microsoft Windows Server 2003	Standard or Enterprise R2
Microsoft Windows 7	
Microsoft Windows Vista	
Microsoft Windows XP	Professional
Red Hat Enterprise Linux	4.7; 5.4–6
Solaris SPARC	10
SUSE Linux Enterprise Server	11–11.1

Note – Suse Linux systems may require additional SSL libraries.

Supported Web Browsers

You can use the following Web browsers with up.time:

- Internet Explorer 7 or higher
- Firefox 3.6 or higher
- Chrome 10 or higher

Minimum Hardware Configuration

The hardware configurations for an Enterprise Monitoring Station can change depending on the number of Local Datacenters that you want to monitor, the reports that you want to generate, and the amount of data that in the up.time DataStore.

The following is the recommended minimum hardware:

- 2.4 GHz dual-core processor
- 4 GB of memory
- 80 GB of disk storage (does not include Oracle storage requirement)
- 100 Mbps network interface (1 Gbps preferred)

Database Requirements

Although up.time Local Datacenters include MySQL as the DataSource, the Enterprise Monitoring Station requires integration with Oracle 11g R1 or R2.

up.time Local Datacenters

Before the Enterprise Monitoring Station (EMS) is installed, you should already be using up.time Local Datacenters (LDCs) to collect performance data from respective groups of servers, network devices, Applications, or SLAs. After the EMS has been installed, you will need to define Datacenters from which data will be aggregated.

Adding Datacenters is covered later in this guide, in “Working with Datacenters” on page 53.

Installing the up.time Enterprise Monitoring Station

The Enterprise Monitoring Station is installed a single directory:

- `/usr/local/uptime` on Linux
- `/opt/uptime` on Solaris
- `C:\Program Files\uptime software\uptime` on Windows

On Windows, the up.time Enterprise Monitoring Station is installed using a graphical installer that guides you through the steps of the installation process. On Solaris or Linux, the installer is a console application.

Before installing up.time, you must be logged in as a local (i.e., non-domain) administrator (in Windows) or as root (in Solaris or Linux).

Before You Begin

There are two ways in which to install the up.time Enterprise Monitoring Station:

1. **From an archive downloaded from the uptime software Web site.**

   If you have downloaded the up.time distribution from the uptime software Web site, copy the archive to a temporary directory on the system that will host the Enterprise Monitoring Station. For the Windows installer, extract the contents of the archive using a utility like WinZip.

2. **From the distribution CD.**

   If you are installing up.time from the distribution CD, do the following:

   - Insert the CD in the CD-ROM drive.
   - If you are installing up.time on Solaris or Linux, mount the CD-ROM drive if you are not using automount.
   - Change to the following directory on the CD:

     `up.time_MonitoringStation`

Once preparations have been made, refer to the procedures in the "Installing the Enterprise Monitoring Station on Windows" on page 24, or "Installing the Enterprise Monitoring Station on Solaris or Linux" on page 26 for details on completing the installation for your platform.

Installing the Enterprise Monitoring Station on Windows

To install the up.time Enterprise Monitoring Station on Windows, do the following:

1. **If you are upgrading, ensure you have logged out of the up.time Web application by clicking the Logout button.**
2. **Ensure you are logged in to the Enterprise Monitoring Station system as the local administrator.**

   up.time may not function properly if the Enterprise Monitoring Station is installed when you are logged in as a domain or non-local administrator.
3. **Double click the following file:**

   `up.time-5.0.<build#>-win32-x86.exe`

   Where `<build#>` is the number of the up.time build that you are installing. For example:

   `up.time-5.0.455-win32-x86.exe`
4. **On the Introduction screen, click Next.**
5. **On the License Agreement screen, carefully read the up.time end user license agreement, and then click the I accept the terms of the license agreement option.**
6. **Click Next.**
7. **Do one of the following to set the location where up.time will be installed:**

   - Click **Next** to accept the default location (`C:\Program Files\uptime software\uptime`).
   - In the **Please Choose a Folder** field, type the name of the directory where you want to install the application and then click **Next**.

- Click **Choose** and select a directory from the **Browse for Folder** window.
- To recover the default directory, click **Restore Default Folder**.

8 **Do one of the following to specify the basic up.time configuration information:**

- Click **Next** to accept the defaults.
- Enter information in the following fields:
  - Email address

    The email address from which the Enterprise Monitoring Station will send alerts and reports to users.
  - DataStore Port

    The number of the port on which the DataStore (the up.time database) will listen for requests. The port number is written to the file `uptime.conf`. Leave this at the default value.
  - Web Server Name

    The name of the computer that is hosting the Web server. This name is written to the file `httpd.conf`, which contains configuration information for the Web server used by up.time.
  - Web Server Port

    The number of the port on which the Web server for the Enterprise Monitoring Station will listen for requests. The port number is written to the file `httpd.conf`.

9 **Select an option for setting up icons in the Windows Start menu and then click Next.**

10 **On the Install Summary screen, review the installation options that you selected and then do one of the following:**

- Click **Previous** to change the settings.
- Click **Install** to begin the installation process.

The installation process will take several minutes.

11 **When the software is installed, click Next.**

The following occurs:

- The Web server, DataStore and Data Collector are installed.
- The Web server and DataStore are started.
- The DataStore is populated with default data.
- The Data Collector is started.

12 **On the Install Complete screen, click Next.**

13 **Click Finish.**

Installing the Enterprise Monitoring Station on Solaris or Linux

Installation on Solaris or Linux is done at the command line. In addition to installing the up.time application, the installation process attempts to create the `uptime` user ID (which run applications in non-privileged mode). If it already exists, then the installer will use that account.

Installing the Enterprise Monitoring Station

To install the up.time Enterprise Monitoring Station on Solaris or Linux, do the following:

1 **If you are upgrading, ensure you have logged out of the up.time Web application by clicking the Logout button.**

2 **Ensure you have logged in to the Entperprise Monitoring Station system as root.**

up.time may not function properly if the Enterprise Monitoring Station is installed when you are logged in as a domain or non-local administrator.

3 **Type the following command:**

```
sh up.time-5.0.<build#>-<platform>.bin
```

where `<build#>` is the number of the up.time build that you are installing, and `<platform>` is the operating system on which you are installing up.time. For example:

- Linux: `up.time-5.0.455-rhes4-x86.bin` or `up.time-5.0.455-sles9-x86-upgrade.bin`

- Solaris: `up.time-5.0.455-solaris-sparc.bin`

It can take up to several minutes for the components of the installer to be extracted from the `.bin` file. Wait while this process completes.

4 **On the Introduction page, press Enter to continue.**

5 **On the License Agreement page, carefully read the up.time end user license agreement. Press Enter to scroll through the agreement.**

6 **At the `DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N)` prompt, type `y` and press Enter.**

7 **Do one of the following to set the directory in which up.time will be installed:**

- Press Enter to accept the default location (`/opt/uptime` on Solaris, and `/usr/local/uptime/` on Red Hat and SLES)
- Type a new location at the command prompt (for example, `/opt/uptime` on Solaris), then press Enter.

The uptime user account must be able to access the directory that you specify.

8 **Do one of the following to specify the basic up.time configuration information:**

- Press Enter to accept the default for each option that is listed below.
- Type new information for each of the following options:
  - Web Server Name

    The name of the computer that is hosting the Web server. This name is written to the file `httpd.conf`, which contains configuration information for the Web server used by up.time.
  - Web Server Port

    The number of the port on which the Web server for the Enterprise Monitoring Station will listen for requests. The port number is written to the file `httpd.conf`.
  - up.time email address

The email address from which the Enterprise Monitoring Station will send alerts and reports to users.

- DataStore Port

  The number of the port on which the DataStore (the up.time database) will listen for requests. The port number is written to the file `uptime.conf`. Leave this at the default value.

9 **On the Install Summary page, review the installation options and then do one of the following:**

- Type `back` and then press Enter to change any of the settings.
- Press Enter begin the installation process.

The installation process will take several minutes.

10 **When the software is installed, press Enter.**

The following occurs:

- the Web server, DataStore and Data Collector are installed
- the Web server and DataStore are started
- the DataStore is populated with default data
- the Data Collector is started

11 **On the Install Complete page, press Enter.**

It can take up to a minute for the up.time services to start. Wait before attempting to log into the Enterprise Monitoring Station.

12

Post-Installation Tasks

After installing up.time, you will need to do the following:

- configure the newly installed Monitoring Station system for its role as the Enterprise Monitoring Station
  - switch the Monitoring Station to be the Enterprise Monitoring Station in the `uptime.conf` file
  - modify the `uptime.conf` file to switch the up.time database from the bundled MySQL implementation to your organization's Oracle database server
  - create database links between the EMS and the LDCs' Oracle instances
- log in to, and customize, the Enterprise Monitoring Station
  - set up the administrator account when you first log in (see "Setting Up the Administrator Account" on page 36)
  - provide the host name of the SMTP server when you first log in (see "SMTP Server" on page 294)
  - install the license for the up.time Enterprise Monitoring Server (see "License Information" on page 317)
  - add users (see "Configuring Users" on page 117)

Enabling the Enterprise Monitoring Station

To enable the Enterprise Monitoring Station, edit the `uptime.conf` file, which is located at the root of the up.time installation directory:

- `/opt/uptime` on Solaris
- `/usr/local/uptime/` on Red Hat and SLES
- `C:\Program Files\uptime software\uptime` on Windows

In the `uptime.conf` file, add the following line: `ems.enabled=true`. This ensures all EMS functionality exists when you log in to the Monitoring Station.

Configuring the Monitoring Station to Use Oracle

In a multi-datacenter deployment, all up.time Monitoring Stations must write to an Oracle database instance instead of the default bundled MySQL implementation. To switch the database used by the Enterprise Monitoring Station, edit the `uptime.conf` file.

To edit the `uptime.conf` file to use an Oracle database instance instead of MySQL, do the following:

1. **Remove or comment out the default MySQL settings, as shown below:**

```
# dbDriver=com.mysql.jdbc.Driver
# dbType=mysql
# dbHostname=localhost
# dbPort=3308
# dbName=uptime
# dbUsername=uptime
# dbPassword=uptime
```

2. **Show (i.e., uncomment) the Oracle database settings.**
3. **For the dbHostname and dbPort settings, enter the address and port for your Oracle database server.**
4. **For the dbName setting, provide a name for the Enterprise Monitoring Station's Oracle database instance.**
5. **In the dbUsername and dbPassword fields, enter the authentication details to access and write to the database.**
6. **Save your changes.**
7. **Use the `resetdb` utility with the `really` option to delete, then recreate the database structure that is used by up.time by running the appropriate command:**

   - Linux: `/usr/local/uptime/resetdb really`
   - Solaris: `/opt/uptime/resetdb really`
   - Windows: `C:\Program Files\uptime software\uptime\resetdb really`

Linking the EMS and LDC Databases

In order to aggregate data from Local Datacenters, the Enterprise Monitoring Station must be able to access the Oracle instances used by each LDC. You can establish these database links by making two key configuration changes on the EMS for each LDC that is part of the multi-datacenter deployment: configure the EMS to communicate with an LDC database instance, then actually establish a database link.

By default, Oracle can link to up to four databases, which means the EMS can aggregate data from up to four Datacenters. For assistance on linking to more database instances, contact uptime software Client Care.

Establishing Communication Through TNS

You can establish communication with LDC databases by modifying the EMS' Oracle Transparent Network Substrate file (`tnsnames.ora`). This file is found in the `/ORACLE_HOME/network/admin/` directory on UNIX-based systems, and `%ORACLE_HOME%\network\admin` directory on Windows-based sytems.

For each LDC, add the following configuration entry to the `tnsnames.ora` file. For the placeholder names displayed in blue, provide values that reflect the LDC's database instance details:

```
$ALIAS$=
 (DESCRIPTION=
  (ADDRESS=(PROTOCOL=TCP)(HOST=$HOSTNAME$)(PORT=$PORT$))
   (CONNECT_DATA=
     (SERVER=DEDICATED)
     (SERVICE_NAME=$SERVICE_NAME$)
   )
 )
```

The service name, host name, and port (which identify the name and location of the LDC database), were set when the LDC's administrator first

configured it to use an Oracle database instance, and can be found in the LDC's `uptime.conf` file as `dbName`, `dbHostname`, and `dbPort`, respectively.

For example:

```
ORASERV_UK=
 (DESCRIPTION=
   (ADDRESS=(PROTOCOL=TCP)(HOST=ldc_EU_2)(PORT=1541))
   (CONNECT_DATA=
     (SERVER=DEDICATED)
     (SERVICE_NAME=ora10)
   )
 )
```

You will determine the alias used to refer to the LDC. This alias will be required in the next section.

Linking the Databases

Using a utility such as SQL*Plus, for each LDC Oracle instance, use the `CREATE DATABASE LINK` statement to link the LDC instance to the EMS database. For the placeholder values displayed in blue, provide values that reflect the LDC's database instance details:

```
CREATE database link "DB_LINK_NAME" connect to
"USERNAME" identified by "PASSWORD" using 'ALIAS';
```

The user name and password used to access the LDC database instance were set when the LDC's administrator first configured it to use an Oracle database instance, and can be found in the LDC's `uptime.conf` file as `dbUsername` and `dbPassword`. The alias was assigned to the `tnsnames.ora` file entry for this LDC in the previous section.

For example:

```
CREATE database link "LDC-UK" connect to "oradmin"
identified by "orapassowrd" using 'ORASERV_UK';
```

You will determine the name of the database link for the LDC. This link name will be required when you are adding LDCs to the EMS (see "Adding

Replication Groups to a Datacenter" on page 56 for more information).

CHAPTER 4

Getting Started

This chapter introduces you to the basic features of up.time in the following sections:

Accessing and Exiting up.time

Before logging into up.time, you will need a user name and password from your system administrator. Your system administrator will provide assistance if this is your first time logging into the application.

Setting Up the Administrator Account

The first user to log into up.time should be the system administrator. While the administrator account has the default user name `admin`, you will have to set the password and email address for the administrator account. You will only need to do this the first time that you log into up.time.

To set up the administrator account, do the following:

1. **Enter the following in the address bar of a Web browser:**

   `http://<uptime_hostname>:<port>`

   Where `<uptime_hostname>` is the name or IP address of the server that is hosting the Enterprise Monitoring Station. For example:

   `http://localhost:9999`

   The up.time log in window opens in a Web browser.

2. **Enter the password for the administrator in the Password field.**
3. **Re-enter the password in the Confirm Password field.**
4. **Enter your email address in the Administrator's Email field.**
5. **Click the Login button.**

Accessing up.time

Once an administrator sets up your up.time account, you can navigate and log in to the Enterprise Monitoring Station.

To start up.time, do the following:

1. **Start a Web browser.**
2. **Enter the following in the address bar of the Web browser:**

   `http://<uptime_hostname>:<port>`

   Where `<uptime_hostname>` is the name or IP address of the server that is hosting the Enterprise Monitoring Station.

   The up.time log in window opens in the Web browser.
3. **Enter your assigned user name in the User Name field.**
4. **Enter your assigned password in the Password field.**
5. **Click the Login button.**

Exiting up.time

To exit up.time, click the **Logout** button () in the top right corner of the screen.

Viewing System and Service Information

You can view information about the following:

- basic configuration of systems in your environment
- services and service groups assigned to the system
- user groups assigned to the system

Viewing System Information

To view system information, do the following:

1. **In the Global Scan or My Enterprise panels, click the name of a system.**

   The general information for the system appears in the sub panel.

2. **Click the Info tab, and then click one of the following options in the Tree panel:**

   - Info & Rescan

     Lists the basic information about the system, including the following:

     - the display name of the system in up.time
     - the host name
     - the number of processes the monitors will retrieve
     - whether or not the system is being monitored
     - the name of the domain on which the system resides (e.g., `uptimesoftware.com`)
     - the name and version of the operating system that is running on the system
     - the number of CPUs on the system
     - the amount of memory, in megabytes, on the system
     - the size of the paging file, in megabytes, on the system

3 **Click the Rescan Configuration button to refresh the configuration information for an agent or a Net-SNMP host. You would do this, for example, if a disk was added to the system. A progress window appears.**

When the message `Configuration Rescanning Completed` appears, click **Close Window**. Information about the configuration changes, if any, appears in **Configuration Changes** section of the subpanel.

If the system that you selected in step 1 is a node, then only the following information appears: the display name and host name of the node, its parent group, and whether or not the node is monitored.

- CPU Information

  Lists the speed (in MHz) of all of the CPUs on the system.

- Network

  Lists the network interfaces on the system, as well as the IP addresses of those interfaces.

- Disks/File System

  Lists the disks that are on Solaris and Linux systems and the names of the file systems that up.time is monitoring.

- Poll Agent

  Displays the output from an up.time agent that you suspect may have a problem. You can forward the output to uptime software Client Care when you encounter problems with up.time.

- Services

  Lists the services assigned to the system, as well as the interval (in minutes) at which the services are checked.

- User Groups

  Lists the user groups that are associated with the system.

Viewing Service Information

To view system information, do the following:

1. **In the Global Scan or My Enterprise panels, click the name of a system.**
2. **Click the Services tab in the Tree panel.**
3. **Click one of the following options in the Tree panel:**

   - Status

     Lists the status of each service assigned to the system, for example:

     `up.time agent running on subway [up.time agent running on subway, up.time agent 4.0 solaris]`

An arrow (➔) at the end of a status message indicates that there is more text. Hold your mouse over the arrow to view the full message.

When up.time issues an alert, you can acknowledge the alert in the **Status** subpanel. For more information, see “Acknowledging Alerts” on page 73.

   - Trends

     Displays one or more graphs that chart the status of the services associated with a host, as shown below:

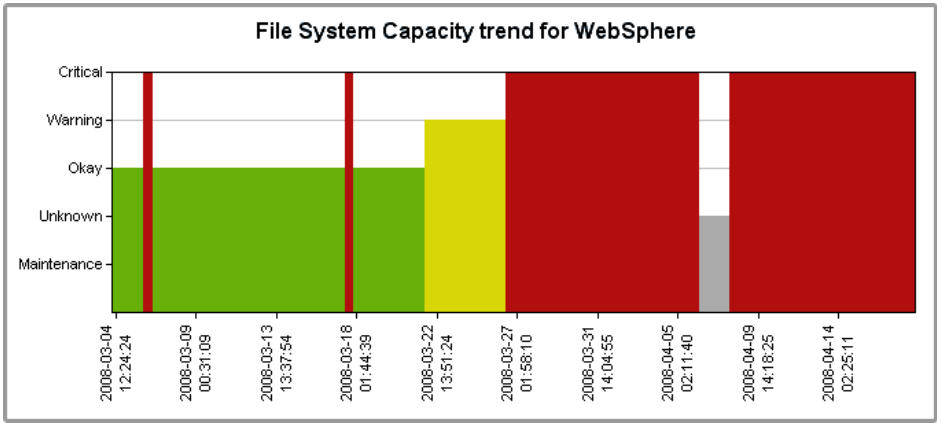

For more information on what each status means, see “Understanding the Status of Services” on page 14.

- Outages

  Lists, in tabular format, the services that have suffered outages, along with the time at which the outage occurred. The Outages table is shown below:

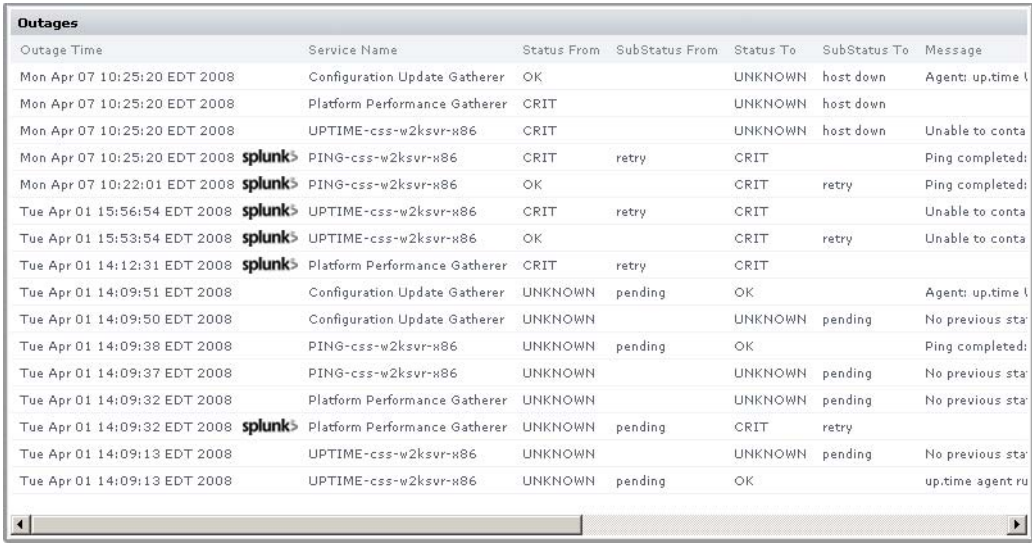

Outages

Outage Time	Service Name	Status From	SubStatus From	Status To	SubStatus To	Message
Mon Apr 07 10:25:20 EDT 2008	Configuration Update Gatherer	OK		UNKNOWN	host down	Agent: up.time \
Mon Apr 07 10:25:20 EDT 2008	Platform Performance Gatherer	CRIT		UNKNOWN	host down	
Mon Apr 07 10:25:20 EDT 2008	UPTIME-css-w2ksvr-x86	CRIT		UNKNOWN	host down	Unable to conta
Mon Apr 07 10:25:20 EDT 2008 splunk>	PING-css-w2ksvr-x86	CRIT	retry	CRIT		Ping completed:
Mon Apr 07 10:22:01 EDT 2008 splunk>	PING-css-w2ksvr-x86	OK		CRIT	retry	Ping completed:
Tue Apr 01 15:56:54 EDT 2008 splunk>	UPTIME-css-w2ksvr-x86	CRIT	retry	CRIT		Unable to conta
Tue Apr 01 15:53:54 EDT 2008 splunk>	UPTIME-css-w2ksvr-x86	OK		CRIT	retry	Unable to conta
Tue Apr 01 14:12:31 EDT 2008 splunk>	Platform Performance Gatherer	CRIT	retry	CRIT		
Tue Apr 01 14:09:51 EDT 2008	Configuration Update Gatherer	UNKNOWN	pending	OK		Agent: up.time \
Tue Apr 01 14:09:50 EDT 2008	Configuration Update Gatherer	UNKNOWN		UNKNOWN	pending	No previous sta
Tue Apr 01 14:09:38 EDT 2008	PING-css-w2ksvr-x86	UNKNOWN	pending	OK		Ping completed:
Tue Apr 01 14:09:37 EDT 2008	PING-css-w2ksvr-x86	UNKNOWN		UNKNOWN	pending	No previous sta
Tue Apr 01 14:09:32 EDT 2008	Platform Performance Gatherer	UNKNOWN		UNKNOWN	pending	No previous sta
Tue Apr 01 14:09:32 EDT 2008 splunk>	Platform Performance Gatherer	UNKNOWN	pending	CRIT	retry	
Tue Apr 01 14:09:13 EDT 2008	UPTIME-css-w2ksvr-x86	UNKNOWN		UNKNOWN	pending	No previous sta
Tue Apr 01 14:09:13 EDT 2008	UPTIME-css-w2ksvr-x86	UNKNOWN	pending	OK		up.time agent ru

  The Outages table also lists all changes to the states and substates for services and host checks – for example, from `OK` to `CRIT` and then from `CRIT` to `OK`.

  As well, up.time displays a message describing the outage – for example:

```
Socket error has occurred connecting to elinux
Error text: Connection timed out: connect
```

  If you are using the Splunk IT search engine with up.time, the Splunk icon (splunk>) appears beside the names of services that are in WARN or CRIT states. You can click the icon to check the Splunk logs for information about the outage.

- Availability

  Lists the state – `OK`, `WARN`, `CRIT`, `MAINT`, `UNKNOWN` – of the monitors that are associated with a specific host or device, as well as:

- the amount of time that the services have been in each state and the total of all times
- the percentage of time each service has been in each state

The Availability table is shown below:

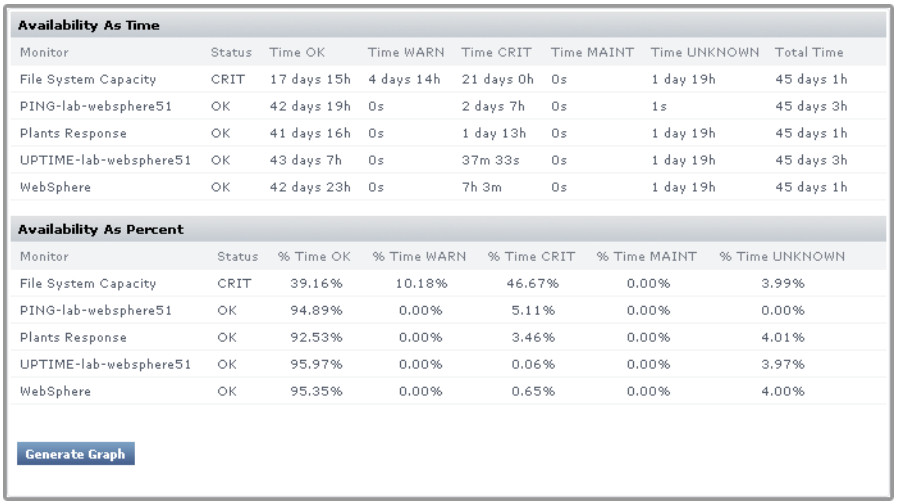

Availability As Time

Monitor	Status	Time OK	Time WARN	Time CRIT	Time MAINT	Time UNKNOWN	Total Time
File System Capacity	CRIT	17 days 15h	4 days 14h	21 days 0h	0s	1 day 19h	45 days 1h
PING-lab-websphere51	OK	42 days 19h	0s	2 days 7h	0s	1s	45 days 3h
Plants Response	OK	41 days 16h	0s	1 day 13h	0s	1 day 19h	45 days 1h
UPTIME-lab-websphere51	OK	43 days 7h	0s	37m 33s	0s	1 day 19h	45 days 3h
WebSphere	OK	42 days 23h	0s	7h 3m	0s	1 day 19h	45 days 1h

Availability As Percent

Monitor	Status	% Time OK	% Time WARN	% Time CRIT	% Time MAINT	% Time UNKNOWN
File System Capacity	CRIT	39.16%	10.18%	46.67%	0.00%	3.99%
PING-lab-websphere51	OK	94.89%	0.00%	5.11%	0.00%	0.00%
Plants Response	OK	92.53%	0.00%	3.46%	0.00%	4.01%
UPTIME-lab-websphere51	OK	95.97%	0.00%	0.06%	0.00%	3.97%
WebSphere	OK	95.35%	0.00%	0.65%	0.00%	4.00%

Generate Graph

Optionally, click the **Generate Graph** button to display pie charts that graph the status of each service, as shown below:

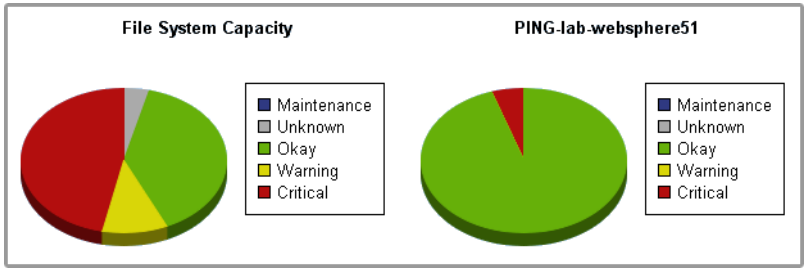

4 **Optionally, click Service Metrics to generate a graph that visualizes retained data over a given period of time. For more information about retained data, see “Understanding Retained**

Data" on page 18.

To generate a graph, do the following:

- Select the date range for the graph from the **Date Range** area. For more information, see "Understanding Dates and Times" on page 16.
- In the **Current Retained Service Metrics** area, select the retained data variables that you want to graph, as shown below:

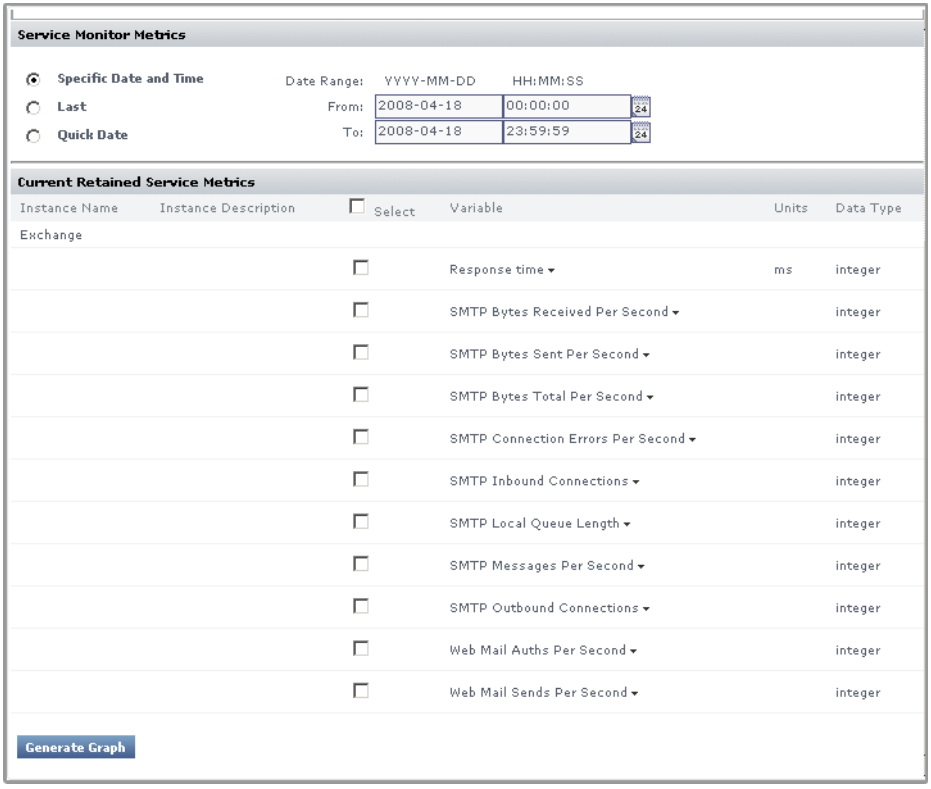

- Click **Generate Graph**.

Searching and Filtering

If you have a large number of hosts on your system, you can use the search and filtering functions in the up.time Web interface to quickly display and view information about specific hosts.

Using the Search Box

You can use the search box at the top of the up.time Web interface to display the basic information about a particular host.

To use the search box, do the following:

1 **From anywhere in the up.time Web interface, enter any of the following information in the Search box:**

- The name of the system for which you want to search.

You can enter a partial name in the **Search** box. For example, if you want to display all systems whose names start with `Web`, enter `Web` in the **Search** box.

- Details about the architecture of the servers. For example, to use an operating system as the search criteria enter `Linux` in this field.
- Any information that may appear in the Custom fields in the profile for the system.

2 **Click Go.**

The following information is displayed:

- name of the host
- description of the host (if any)
- the operating system and type of hardware on which the host is running

- any information in the four custom fields in the system profile (e.g., the job being done by the system, and its physical location)

Audit Logging

up.time can record changes to the application's configuration in an audit log. The details of the configuration changes are saved in the file `audit.log`, found in the `logs` directory.

Windows Vista users can find the audit log in the Virtual Store instead of the default location (i.e., C:\Users\uptime\AppData\Local\VirtualStore\Program Files\<uptime-install-directory>

There are many uses for the audit log. For example, you can use the audit log track changes to your up.time environment for compliance with your security or local policies. You can also use the audit log to debug problems that may have been introduced into your up.time installation by a specific configuration change; the audit log enables you to determine who made the change and when it took effect.

The following is an example of an audit log entry:

```
2006-02-23 12:28:20,082 - dchiang: ADDSYSTEM [cfgcheck=true,
port=9998, number=1, use-ssl=false, systemType=1,
hostname=10.1.1.241, displayName=MailMain,
systemSystemGroup=1, serviceGroup=, description=,
systemSubtype=1]
```

Enabling the Audit Log

By default, the audit log is disabled. To enable it, edit the `uptime.conf` file, which is located at the root of the up.time installation directory:

- `/opt/uptime` on Solaris
- `/usr/local/uptime/` on Red Hat and SLES
- `C:\Program Files\uptime software\uptime` on Windows

In the `uptime.conf` file, locate the "`auditEnabled=`" entry and modify it to be "`auditEnabled=yes`". If the entry does not exist, add the entry to the file.

CHAPTER 5

Using My Portal

This chapter explains the **My Portal** panel.

Overview

When you log into up.time, the first screen you see is the **My Portal** panel. The **My Portal** panel gives quick access to basic up.time functions and to saved reports. The **My Portal** panel is divided into several sections:

- Assistance
- My Preferences
- Latest up.time Articles
- up.time Information
- My Alerts
- Saved Reports
- Custom Dashboards

Assistance

The top portion of the **My Portal** panel gives you quick access to:

- tutorials that demonstrate how to perform basic tasks in up.time
- up.time's online help
- the uptime software community support forums

There is also a search engine with which you can find information in the Client Care Web site Knowledge Base and support forums.

The following image illustrates the top portion of the **My Portal** panel:

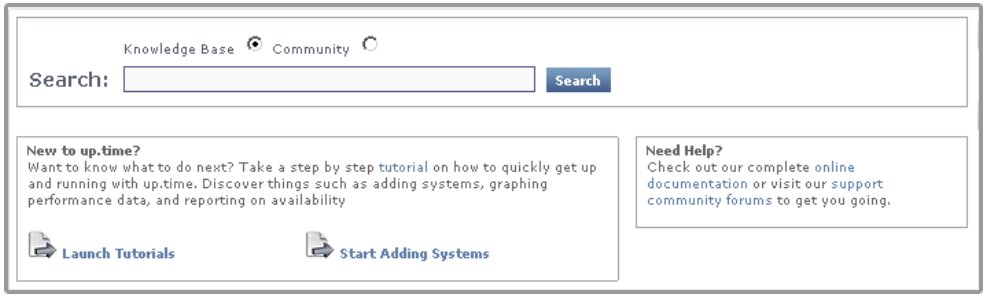

My Preferences

The **My Preferences** section enables you to:

- View your user account settings. Click the **View** icon () or your user name to open your account settings in the subpanel. You can also edit your user information by clicking **Edit User**.
- Change your user account settings. Click the **Edit** icon (). The Edit User window appears. See “Editing User Information” on page 124 for details on editing your user account settings.

Latest up.time Articles

The **Latest up.time Articles** section contains a list of recent Knowledge Base articles. This list is fed to the **My Portal** panel via RSS (Really Simple Syndication, a method for delivering summaries of and links to Web content). You simply click the title of the article to open it in your Web browser.

up.time Information

The **up.time Information** section contains the following information about your Enterprise Monitoring Station:

- Whether or not updates are available. If an update is available, there will be a link to the uptime software Client Care Portal where you can download the update.
- The status of your license, including the type of license and the numbers remaining before the license expires.

My Alerts

The **Current Issues** section contains a list of systems that are in a warning or critical state.

Saved Reports

The **Saved Reports** tab lists the reports that you have scheduled and saved. For more information on scheduling reports, see “Scheduling Reports” on page 169.

This section contains the following information about the reports:

- the name of the report
- an optional description of the report
- whether or not the report is scheduled to run at a specific time
- whether or not the report will be saved to a directory on the Enterprise Monitoring Station or on another server
- the time at which the report will next be run, in the following format:

  ```
  Wed Oct 12 14:30:00 EDT 2005
  ```

The **My Portal** panel only displays the reports and graphs that you have defined. However, a system administrator or a user with administrator privileges can view all saved reports.

Custom Dashboards

A custom dashboard tab displays the contents of an external Web page that is referenced by URL. Creating one or more custom tabs allows up.time users to view customized content through **My Portal**.

Custom dashboards are visible to members of specific, dashboard-related User Groups. For information on configuring a custom dashboard, see “Custom Dashboard Tabs” on page 316.

up.time

CHAPTER 6

Defining and Managing Your Enterprise

This chapter explains the **My Enterprise** panel in the following sections:

Overview

The **My Enterprise** panel is your starting point for monitoring the systems in your environment. From the **My Enterprise** panel, you can add:

- Applications, which provide the overall status for one or more services
- service level agreements, which measure compliance to infrastructure performance goals
- groups, which are sets of Applications or SLAs that have been combined in a meaningful way
- views, which enable non-administrative users to view only the systems in which they are interested
- Agentless WMI

A Windows-based system whose metrics collection is managed by WMI (Windows Management Instrumentation), and does not have an up.time agent installed on it. Adding a WMI System to up.time

1

1

Working with Datacenters

up.time's monitoring, alerting, and reporting capabilities for a distributed, but single-site environment can be scaled to a multi-datacenter (MDC) environment. In an up.time MDC environment, each datacenter is independently monitored by an up.time instance. Data collected by Local Datacenters (LDCs) are aggregated at the central Enterprise Monitoring Station (EMS).

Monitoring and reporting on aggregated LDC data allows enterprise-wide systems management from a single up.time Monitoring Station. Additionally, the health of the LDCs themselves are also monitored, providing EMS-level administrators and managers a high-level view of the enterprise infrastructure through which system performance monitoring is possible. The enterprise-wide systems management view also allows capacity planning and server consolidation beyond the realm of a single datacenter.

In addition to system and device monitoring at the infrastructure and enterprise levels, centralized monitoring also enables owners of specific domains within an organization to selectively monitor the health of Elements that are most relevant to the domain owner's role (e.g., database server management, overseeing of global applications).

Representing Datacenters on the EMS

On the Enterprise Monitoring Station, each LDC is monitored and managed as an individual entity, or Datacenter, and the Elements that are a part of that LDC (and are monitored by that up.time instance) form one or more Replication Groups.

Representing each LDC as an individual Element on the EMS provides administrators a big-picture view of all Datacenters in the **Global Scan** panel. Administrators with proper Datacenter access can click through the EMS to the LDC Monitoring Station to further analyze performance data at the Datacenter level.

Representing groups of an LDC's Elements in Replication Groups provides administrative flexibility, allowing administrators to create groups that match the structure of responsibility in their organization.

Consider, for example, an IT group whose responsibilities are separated into the areas of infrastructure, application servers, and database servers. The EMS administrator could create three Replication Groups for each Datacenter, each comprised of Elements related to those areas; clicking a Datacenter entry in the **Global Scan** panel would then display these Replication Groups, allowing other administrators to easily monitor the status of particular Datacenter Elements. In another scenario, an IT group that only needs to treat a datacenter as a whole would create a single Replication Group, placing all of the Datacenter's Elements in it.

Note that each Datacenter Element can only be a part of one Replication Group, and every Datacenter Element does not have to be represented in a Replication Group.

To provide a complete picture of multi-datacenter performance, all relevant Datacenter information is replicated on the EMS:

- Element configuration information
- knowledge of the service monitors that are attached to each Element
- service monitor configuration information
- SLA and Application status monitors
- all performance and retained data for Elements and service monitors that are viewed in up.time graphs and reports (see "Using Graphs" on page 249 and "Using Reports" on page 175 for more information):
    - disk statistics such as Top 10 Disks, File System Capacity, and Disk Peformance Statistics
    - network statistics such as TCP Retransmits, I/O, and Errors
    - virtual environment reports: VMware Workload and LPAR Workload
    - process-related information such as Number of Processes,
    - workload information such as Workload by User, Workload by Process Name, and Workload Top 10

About Unreplicated Configuration Information

Datacenter configuration information that is not replicated include the following:

- Alert Profiles and Action Profiles
- Datacenter-level user information (i.e., users, user groups, user roles)
- Element groups and Element views

Configuration information pertaining to the Local Datacenter itself is not replicated on the EMS. The responsibilities of administrators (and resulting configuration of Elements, users, and alerting policies) at each respective Local Datacenter can differ from EMS-level administrators.

For example, since an LDC's Element group or view structures are not replicated, a flat list of all Elements are available to EMS administrators for group and view creation. This flat list allows enterprise-focused users (e.g., a global application owner), to assemble a custom view on the EMS that includes Elements from not only the aforementioned LDC, but relevant Elements from other monitored LDCs as well. For more information, see "Working with Groups" on page 66 and "Working with Views" on page 69.

Configuring Datacenters

To add a Datacenter to the EMS, do the following:

1. **In the My Enterprise panel, click Add Datacenter.**

   The **Add Datacenter** window appears.

2. **Enter a descriptive Name for the Datacenter.**

   This name will appear in both the **My Enterprise** and **Global Scan** panels.

3. **Optionally enter a Description of the Datacenter.**

4. **Enter a Short Name for the Datacenter.**

   This short name will appear as the click-through link in the Datacenter Bar, which is visible throughout the EMS to administrator-level users. The short name will also be used as a display-name prefix for all Elements from that Datacenter.

5. **Enter the address for the LDC in the Local Datacenter URL field.**

   This address is used as the click-through link in the Datacenter bar, allowing direct access to the LDC's Monitoring Station. For information on which login information is used, see "Understanding Cross-Datacenter Access" on page 59.

6 **In the Oracle Database Link Name field, enter the name of the database link to the LDC's Oracle instance.**

This link must already exist on the EMS configuration, and was established when the EMS was first installed. See "Linking the EMS and LDC Databases" on page 31 for more information.

7 **In the Time Zone field, begin entering the name of the time zone, city, country, or continent in which the LDC exists, then select the correct entry from the dropdown list.**

8 **Click Save.**

The Datacenter short name appears on the Datacenter bar, and in the main view, the Datacenter's **General Information** subpanel is displayed. From this page you can add Replication Groups.

Adding Replication Groups to a Datacenter

To add a Replication Group to a Datacenter, do the following:

1 **In the My Enterprise panel, click the name of the Datacenter that you want to edit.**

The Datacenter's **General Information** subpanel is displayed.

2 **Click Add Replication Group.**

The **Add Replication Group** window appears, indicating the Datacenter to which this Replication Group will be attached:

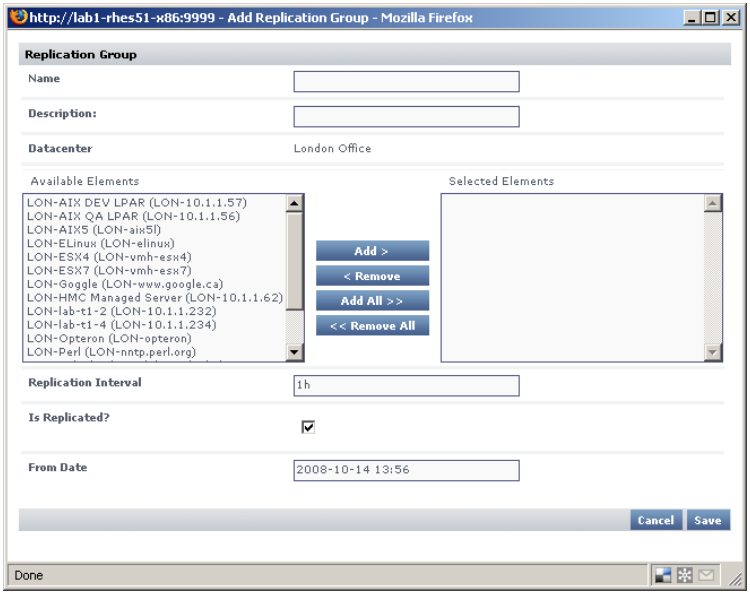

3 **Enter a descriptive Name for the Replication Group.**

This name will appear in the **Global Scan** panel.

4 **Optionally, enter a description for the Replication Group in the Description field.**

5 **Select the systems you want to add to this Replication Group from the Available Elements list, then click Add.**

The listed Elements are all those that are a part of the Local Datacenter, as indicated by the display-name prefix, which was configured when the Datacenter was first created.

6 **In the Replication Interval field, indicate how often the EMS will replicate performance data from the Datacenter for the selected Elements.**

7 **If desired, disable the Is Replicated checkbox if all you do not yet want to replicate Element data, and only want to configure the Replication Group.**

8 **In the From Date field, enter the starting date from which the Datacenter's Element data will be replicated.**

The date must be in `yyyy-mm-dd` format (e.g., `2007-10-01`). For greater precision, you can also include a precise time of day as a starting point in an `hh:mm` format using a 24-hour clock (e.g., `2007-10-01 15:00`).

9 **Click Save to create the Replication Group and begin the data replication process, based on the Datacenter Elements and starting date selected.**

Adding Dependent Nodes to a Datacenter

To improve diagnosis of a Datacenter that reaches a critical state, you can supplement its profile by defining any nodes on which it is dependent. The Enterprise Monitoring Station will check the status of a Datacenter's Dependent Nodes if an EMS attempts to retrieve a configuration update from the Datacenter, but fails.

In most cases, a Datacenter's dependent nodes would be network hardware (e.g., routers or switches) that manage data communication within the Datacenter's infrastructure.

To add a Dependent Node to a Datacenter's profile, do the following:

1 **In the My Enterprise panel, click the name of the Datacenter that you want to edit.**

The Datacenter's **General Information** subpanel is displayed.

2 **Click Add Dependent Node to display the configuration window.**

3 **In the Enter a descriptive Name for the Dependent Node.**

This name appears on the Datacenter's **General Information** subpanel in **My Enterprise**, as well as well as its **Current Status** page (see "Viewing the Status of a Datacenter" on page 60 for more information).

4 **Optionally enter a Description.**

5 **In the Hostname / IP field, enter the network address for the node.**

6 **Click Save.**

Once saved, the new node appears in the **Dependent Nodes** section of the Datacenter's **General Information** subpanel.

Understanding Cross-Datacenter Access

Typically, Local Datacenter administrators will only be concerned with the infrastructure encompassed by the datacenter in which they are located; however, EMS administrators are more likely to have enterprise-wide responsibilities and concerns, and may need be able to perform analysis beyond the EMS level—in other words, access a Local Datacenter's Monitoring Station. Since administrator-level users on the EMS have the ability to "click through" (i.e., log in) to any Datacenter's Monitoring Station, it is important that your organization's up.time usage and account policies allow co-ordination from an administrative perpsective.

The following are paths an Enterprise Monitoring Station administrator can take to access a Local Datacenter:

- clicking a Local Datacenter "short name" in the Datacenter bar, visible throughout the EMS
- clicking the **Link to LDC** icon, located beside any Datacenter name anywhere on the Monitoring Station (e.g., the **Global Scan** or **My Enterprise** panels)
- trying to view more details about an LDC-based SLA or Application (only their status is displayed at the EMS level)

When a user on the EMS attempts to directly accesses a Datacenter, their user name and password used to access the EMS are again used to access the LDC Monitoring Station. If the LDC does not have a user account with that particular EMS user's information, the login attempt will fail.

It is essential that a multi-datacenter deployment, especially one that includes numerous system administrators with different access needs, accommodates cross-datacenter access. In most cases, this means ensuring EMS administrator accounts are created on every LDC Monitoring Station. Depending on how your organization is structured, this could also mean that EMS-level administrator accounts on LDC Monitoring Stations are given a more limited User Role, allowing LDC administrators full control over the infrastructure for which they are exclusively responsible.

Viewing the Status of a Datacenter

If a Datacenter reaches a warning- or critical-level state, it is reported in both the **Global Scan** panel and the Datacenter Bar: its name is highlighted in yellow or red.

There are a number of connectivity requirements between a Datacenter and Enterprise Monitoring Station that must be fulfilled in order for a multi-datacenter monitoring deployment to function correctly; as such, a Datacenter can reach a critical state for any one of these reasons. Use the Current Status view for a Datacenter to see a breakdown of the status checks that have been performed on it.

You can display the Current Status for a Datacenter by clicking its Datacenter Current Status icon, visible in both the **Global Scan** and **My Enterprise** panels.

The Health Check Status indicator reflects all of the status checks. If a single critical-level status check fails, the overall Health Check Status will become critical. The following status checks contribute to the overall health of a Datacenter:

Status Check	Description	Failure Result
Configuration Replication	Indicates whether the EMS was able to connect to, and replicate, an LDC's configuration profile. If it cannot, independent checks on the LDC Web server and Dependent Nodes are performed.	critical
Status of Web Server	Indicates whether the EMS can connect to the Datacenter's Monitoring Station. This check is only performed when a Configuration Replication test fails, and is meant to help you diagnose the problem.	critical
Dependent Node Ping Status	Indicates whether the EMS could ping any defined Dependent Nodes. This check is only performed when a Configuration Replication test fails, and is meant to help you diagnose the problem.	critical

Status Check	Description	Failure Result
Recent Data Check	Ensures the Datacenter's aggregated performance data are current (i.e., not older than five minutes).	critical
Server Time Check	Ensures the LDC and EMS times are in sync (i.e., not greater than 60 seconds apart).	warning
Database Link Check	Not reported as an individual status check, the database link check ensures the LDC's Oracle instance exists, and can be aggregated on the EMS. If a link is not found, the Health Check Status automatically moves to a critical state, and all other status checks cease until this issue has been resolved.	critical

Working with Applications

An Application provides the overall status for one or more services. You can, for example, add an Application that checks the status of a system's Web services, database, and file system capacity.

When creating an Application, you must specify the following:

- master service monitor(s)

  One or more monitors can be used to determine the status of the Application as a whole.

- regular service monitors

  Other service monitors that are associated with a master service monitor, but are not used to determine the status of the Application as a whole.

For more information on services, see "Using Service Monitors" on page 135. For information on viewing information about Applications, see "Viewing Details About Applications" on page 64.

Adding Applications

To add an Application, do the following:

1. **In the My Enterprise panel, click Add Application.**
2. **In the Add Application window, enter a descriptive name for the Application in the Name of Application field.**

   This name will appear in both the **My Enterprise** and **Global Scan** panels.
3. **Optionally, enter a description for the Application in Description of Application field.**
4. **Optionally, select the group of systems in your** up.time **environment with which this system will be associated from the Parent Group dropdown list.**

   By default, the Application is added to the My Enterprise group.

   For more information on groups, see "Working with Groups" on page 66.

5 **Select one of the following options from the dropdown list above the Available Master Service Monitors list:**

- the name of a specific system, which displays all its service monitors
- **All**, which displays all service monitors for every system in your environment

6 **Select one or more of the service monitors from the Available Master Service Monitors list, and then click Add.**

7 **Select one of the following options from the dropdown list above the Available Regular Service Monitors list:**

- the name of a specific system, which displays all its service monitors
- **All**, which displays all service monitors for every system in your environment

8 **Select one or more of the service monitors from the Available Regular Service Monitors list and then click Add.**

9 **Click Save.**

After closing the **Add Application** window, the name of the newly created Application appears in the **My Enterprise** panel as a link that can be clicked to view the Application's details.

10 **If required, associate Alert Profiles with the Application by clicking Edit Alert Profiles when viewing the Application's details.**

11 **In the Alert Profile Selector pop-up window, select one or more of the Available Alert Profiles from the list, then click Save.**

12 **If required, associate Action Profiles with the Application by clicking Edit Action Profiles when viewing the Application's details.**

13 **In the Action Profile Selector pop-up window, select one or more of the Available Action Profiles from the list, then click Save.**

Viewing Details About Applications

After you have added an Application to up.time, the name of the Application appears in the **My Enterprise** panel. The name of the Application is a hyperlink.

You can view detailed information about that Application by clicking the name of the Application, which opens the **Application General Information** subpanel.

The **Application Profile** section of the subpanel displays the following information about the Application:

- the name of the Application
- the description, if available
- the group of systems to which the Application belongs
- whether or not the Application is being monitored

The **Application Member Services** section of the subpanel contains the following information about the service monitors that are part of the Application:

- the name of the service that is being monitored
- whether or not the service is a master service monitor

The **Alert Profiles** section of the subpanel displays which Alert Profiles have been associated with the Application.

For information about viewing more details about Applications, see “Viewing System and Service Information” on page 38.

Editing Applications

To edit an Application, do the following:

1. **In the My Enterprise panel, right-click the name of the Application that you want to modify, then click Edit.**

   The **Edit Application** window appears.

2. **Edit the Application setting as described in “Adding Applications” on page 62.**

Working with SLAs

In up.time, a service level agreement (SLA) measures your organization's ability to meet pre-defined performance goals. These goals focus on various aspects of your IT infrastructure, and each can include any number of monitored systems.

From the **My Enterprise** panel, you can view your existing SLA details by clicking the SLA name (see "Viewing SLA Details" on page 100 for more information).

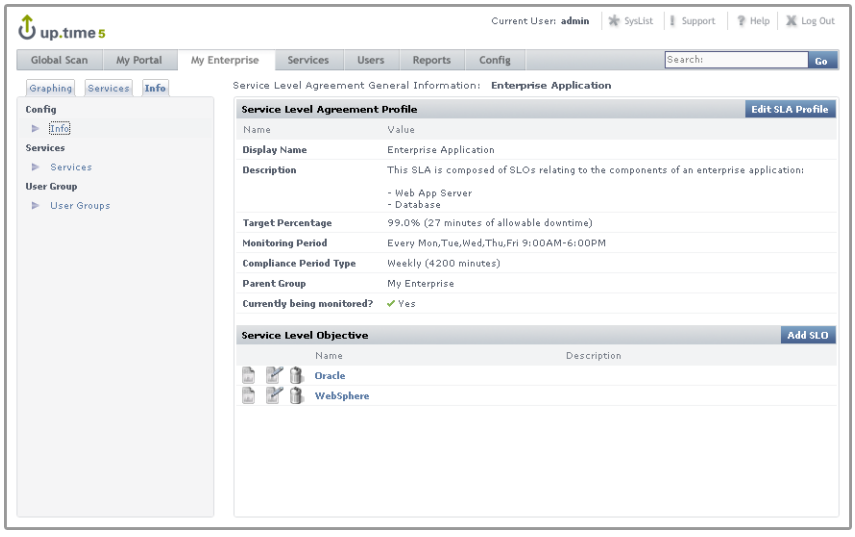

For information about creating and using SLAs, see "Adding and Editing SLA Definitions" on page 111.

Working with Groups

At sites with multiple systems to monitor, searching through a large list of systems is time consuming. To avoid this problem, you can define *groups* of systems. Groups are sets of systems that have been combined in a meaningful way.

You can group systems by their geographical location or by their function. The name of the group should describe the servers or they way in which they have been grouped. For example, you can create a group called `Database Servers` that contains all of the database servers in your environment.

You can assign the following to groups:

- SLAs or Applications
- the user groups that are allowed to view the systems or Elements in a group (see "Working with User Groups" on page 125 for more information on user groups)

If you plan to group your systems, you should first map out what groups you need and which systems will be part of those groups.

Adding Groups

To add a group, do the following:

1. **On the My Enterprise panel, click Add Group.**
2. **Enter a descriptive name for the group in the Group Name field.**
3. **Optionally, enter a description of the group in the Group Description field.**
4. **To make this group a subgroup, select the name of the existing group to which it will be subordinate in the Parent Groups list, then click Add.**

If this is the first group that you have defined, only **My Enterprise** will appear in the dropdown list.

5 **To give this group its own subgroups, select one or more entries from the Available Groups list, then click Add.**

6 **Select the Elements that you want to add to this group from the Available Elements list, then click Add.**

7 **Select one or more sets of users who can view this group from the Available User Groups list, then click Add.**

8 **Click Save.**

Adding Nested Groups

You can also create *nested groups*. Nested groups enable you to further group your systems. For example, you can create a parent group called Datacenters, and then add two nested groups called Production and Disaster Recovery.

You can assign the following to nested groups:

- groups of Elements
- individual Elements
- the up.time user groups that are allowed to view the systems or Elements in a group

Note that you cannot assign a parent group to a subgroup or to any other ancestor.

Before you begin, ensure that you have at least one parent group defined. For more information, see “Adding Groups” on page 66.

Adding a Nested Group

To add a nested group, do the following:

1 **In the My Enterprise panel, click Add Group.**

2 **Enter a descriptive name for the group in the Group Name field.**

3 **Optionally, enter a description of the group in the Group Description field.**

4. **Select the group with which the new one will be associated from the Parent Group dropdown list.**
5. **To give this nested group its own subgroups, select one or more entries from the Available Groups list, then click Add.**
6. **Select the Elements that you want to add to this group from the Available Elements list, and then click Add.**
7. **Select one or more sets of users who can view this group from the Available User Groups list, and then click Add.**
8. **Click Save.**

Editing Groups

To edit groups, do the following:

1. **In the My Enterprise panel, right-click the group you want to modify, then click Edit.**

   The **Edit Element Group** window appears.
2. **Edit the group as described in** "Adding Groups" on page 66.
3. **Click Save.**

To delete a group, right-click it then click **Delete**, but note that only empty groups can be deleted from the My Infrastructure panel.

Working with Views

Not every administrator (i.e., users assigned the “superadmin” user role) that accesses the Enterprise Monitoring Station needs to view all Elements that are a part of your enterprise. In a multi-datacenter environment, different EMS-level administrators are typically interested in monitoring and reporting on the enterprise from different perspectives. By limiting the Elements that one or more administrators can see, you can logically group Elements into *views* that facilitate efficient analysis and monitoring.

For example, given a global deployment of IT systems, one EMS user may be primarily concerned with the enterprise’s back-end performance and work with a view of all database management servers across several Datacenters. In another example, an enterprise application owner will need a view of all up.time Applications, servers, and network devices associated with that application. By creating views, it comes easier for users to not only monitor different aspects of your enterprise, but to also browse and compare historical data.

Due to the scope of view offered at the Enterprise Monitoring Station, non-administrator users (i.e., those who are not assigned the “superadmin” user role), by default, cannot see any replicated Elements. However, EMS administrators can create views that include replicated Elements; these views can be made available for non-administrator users.

Views appear in the Views section on the **My Enterprise** panel, as well as the the **Global Scan** panel.

Adding Views

To add a view, do the following:

1. **In the My Enterprise panel, click Add View.**
2. **In the Add View window, enter a descriptive name in the View Name field.**

   This name will appear when listing views in the **My Enterprise** panel.
3. **Optionally, enter a description in View Description field.**

4. **To make this view a child of an existing one, select it from the Parent View dropdown list.**

   If this is the first group that you have defined, only **My Enterprise** will appear in the dropdown list.

5. **To give this view its own child views, select one or more entries from the Available Element Views list, then click Add.**
6. **Select one or more Elements from the Available Elements list, then click Add.**

   If you have combined your Elements or Datacenters into groups, select a group from the dropdown at the top of the list. Or, select **All** from the dropdown to display all of the Elements in your environment
7. **Select one or more users from the Available Users for View list, then click Add.**
8. **To add previously defined groups of users, select one or more entries from the Available User Groups list, then click Add.**
9. **Click Save.**

Adding Nested Views

You can also create nested views in order to categorize and better manage a larger set of existing views. The following can be assigned to nested views:

- existing Element views
- individual Elements
- individual users who have view access to the Elements in a view
- up.time user groups with similar privileges

You cannot assign a parent view to a child view or to any other ancestor.

Before you begin, ensure that you have at least one parent view defined. For more information, see “Adding Views” on page 69.

Adding a Nested View

To add a nested view, do the following:

1. **In the My Enterprise panel, click Add View.**
2. **In the Add View window, enter a descriptive name in the View Name field.**

   This name will appear when listing views in the **My Enterprise** panel.
3. **Optionally, enter a description in View Description field.**
4. **In the Parent View dropdown list, select the view to which this nested view will be subordinate.**
5. **To give this nested view its own child views, select one or more entries from the Available Element Views list, then click Add.**
6. **Select one or more users who can view this group from the Available Users list, then click Add.**
7. **To add previously defined groups of users, select one or more entries from the Available User Groups list, then click Add.**
8. **Click Save.**

Editing Views

To view and edit views, do the following:

1. **In the My Enterprise panel, right-click the View you want to modify, then click Edit.**

   The **Edit View** window, which contains system and user information, appears.
2. **Edit the view as described in** “Adding Views” on page 69**.**
3. **Click Save.**

Deleting Elements, Applications, and Views

If you have administrator privileges, you can delete an Datacenter, Element or view in the **My Enterprise** panel.

To delete a system or network device, do the following:

1. **Locate the system or network device, Application, or view that you want to delete in the My Enterprise panel.**
2. **Right-click the Element, then click Delete.**
3. **On the dialog box that appears, click OK.**

Acknowledging Alerts

For Datacenters as well as EMS-based SLAs and Applications (i.e., EMS-specific, non-replicated Elements), alerts notify assigned users of status problems. Qualified users can also acknowledge an alert.

When you acknowledge an alert, up.time:

- records the acknowledgement, which can be viewed in the Service Monitor Outages report
- sends an acknowledgement message to any up.time user who received the last alert
- turns off alert escalation, but continues monitoring the problem, and only sends an alert when the status of the system or Application returns to OK

To acknowledge alerts, do the following:

1 **In the My Enterprise panel, click the name of the EMS-based, non-replicated Element that generated the alert.**

The **System General Information** subpanel appears.

2 **In the Tree panel, click the Services tab and then click Status.**

Status information for the monitors associated with the Element appears in the subpanel, as shown below:

3 **Click the Acknowledge icon ( ) in the Ack column.**

The acknowledgement message window appears.

4 **Type a comment relating to the alert or why it has been acknowledged, and then click Submit.**

An email containing the following information is sent to any up.time user who received the last alert:

- the user name and email address of the person who acknowledged the alert
- the name of the Element and service monitor involved
- a comment relating to the alert or reason for acknowledgement

The following is a sample alert acknowledgement message:

```
up.time Administrator (jsmith@myDomain.com)
acknowledged the WARN status of File System Capacity (Web
Server 2) with comment:
Initial check of problem. More information to come.
```

In the up.time Web interface, the acknowledge icon changes to  .

up.time

CHAPTER 7

Overseeing Your Enterprise

This chapter explains the **Global Scan** panel in the following sections:

Overview

The **Global Scan** panel enables you to view the current status of all of the Datacenters, Replication Groups, and Elements (servers and devices, Applications, and SLAs) in your environment. When initially viewed, the **Global Scan** panel typically contains a list of all the Datacenters and Elements that are being monitored by up.time, as shown below:

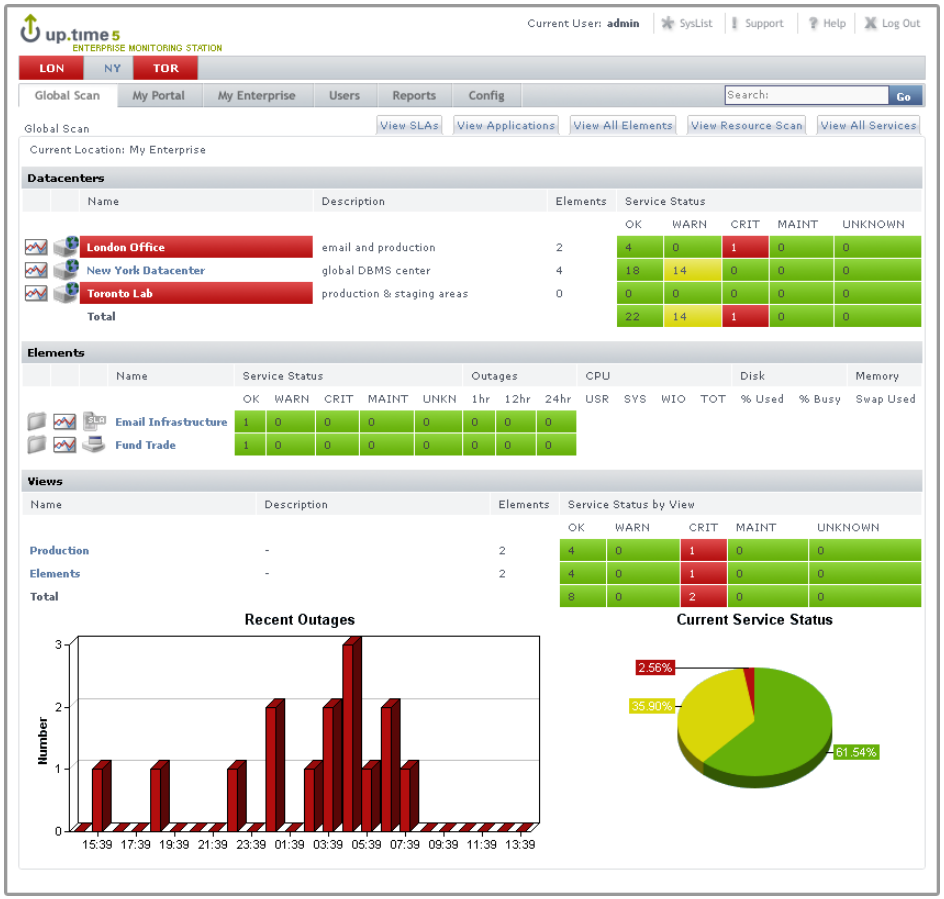

The **Datacenters** table displays the status and number of services that are associated with each Datacenter.

The **Elements** table displays the status and number of outages for SLAs and Applications

Users who are not assigned to the "superadmin" user role will not see any replicated Elements or Datacenters; they will only see Elements that were created on the Enterprise Monitoring Station. EMS administrators can allow regular users to see replicated Elements with custom views. See "Working with Views" on page 69 for more information.

Service status indicators range from normal (green), to Warning (yellow), to Critical (red), and also include an Unknown state (gray). An Unknown state indicates that no performance data for the last 10 minutes exists for the Element. To avoid false positives, note that recently added Elements will have this status until 10 minutes' worth of performance data has been collected; also, in cases where the up.time Data Collector service is down for more than 10 minutes, all Elements will have this status until the service has been restarted and enough data has been collected.

The thresholds for the service status indicators are typically 70% for a warning state, and 90% for a critical state. These thresholds can be customized (see "Changing Reporting Thresholds" on page 95).

The bar chart at the bottom left of the panel displays the number of service monitors that have moved from a normal (OK) to critical (CRIT) status over the past 24 hours. up.time takes a data sample from the database for any *new* critical-status services every 15 minutes, and charts it on the bar chart. The number of services in each state appears in the graph.

The pie chart at the bottom right of the panel visualizes the current availability of systems or devices. The services for unmonitored systems in groups are not shown in the pie chart.

Reporting Datacenter Status

Datacenters, as single entities, are continuously checked to ensure they are functioning correctly for a multi-datacenter environment. Since data aggregation at the central Enterprise Monitoring Station requires precise synchronization as well as Datacenter infrastructures that are not experiencing downtime, any of several Datacenter health issues can result in a critical status. In the **Global Scan** panel, Datacenters that are in a critical state are highlighted in red. Additionally, critical-state Datacenters

are also highlighted in red in the Datacenter Bar, which is visible to any EMS administrator.

You can also view the status of a Datacenter in greater detail by clicking its graph icon, which takes you to its **Current Status** page. See “Viewing the Status of a Datacenter” on page 60 for information on the types of checks that are made against a Datacenter to ensure it is functioning.

Datacenters reach a warning-level state, and are highlighted in yellow, when time drift is detected between Local Datacenters and the Enterprise Monitoring Station.

Datacenter states are shown in the following Global Scan panel image:

Viewing More Information

You can view the list of Replication Groups that are a part of a Datacenter by clicking the latter’s name. Clicking a Replication Group lists all of its Elements.You can view detailed information about an Element by clicking its name. To view the details of each metric (for example, CPU usage) click the number in the column for that variable to go to its Graphing page, where you will be able to generate a graph.

Groups and Views in the Global Scan Panel

When you create groups or views (see “Working with Groups” on page 66 and “Working with Views” on page 69), they appear in their own sections in the **Global Scan** panel. The following information is displayed:

- the names and descriptions of the groups
- the number of Elements in each group
- the status of the hosts that make up the group
- the number of alerts per group

When you click a group or view in the **Global Scan** panel, the systems that make up the group or view and details about their status are displayed.

Viewing All SLAs

Service level agreements in the **Global Scan** panel indicate whether performance targets are being met. For EMS-native SLAs, although the main summary displays the status of the SLA definition as a whole, you can also expand the view to verify how well component service level objectives (SLOs) are meeting targets. (SLOs are made up of monitored services that, as a group, are used to measure a specific performance goal.) For SLAs that are being replicated from Local Datacenters, you are provided with basic status information, and need to log in to the actual LDC Monitoring Station to view more detailed information such as SLO targets.

In the **Service Level Agreements** subpanel (accessed by clicking the **View SLAs** tab), the following SLA information is provided in the default view:

- the list of SLAs, and whether any are in a critical or warning-level state
- headway into the time period during which compliance is measured (for SLAs that were created on the EMS)
- the percentage of allowable downtime used, after which the SLA's status becomes critical (for SLAs that were created on the EMS)

SLA Status Indicators

The color coding used in the **Service Level Agreements** subpanel indicates, at a glance, whether the SLAs' respective limits are in danger of or have already been exceeded:

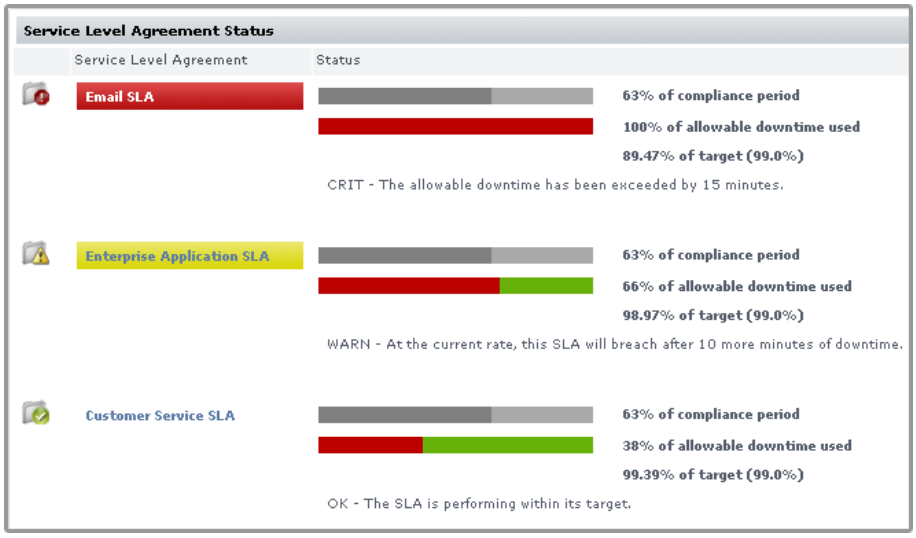

The **Downtime** progress bar allows you to gauge how close the SLA is to reaching a critical state:

- an SLA whose allowable downtime exceeds 100% reaches a critical state, is highlighted with red, and is accompanied by the critical state icon ( )
- an SLA whose allowable downtime, at the current rate of use, will be depleted before the compliance period has ended enters a warning-level state, is highlighted with yellow, and is accompanied by the warning state icon ( )
- an SLA whose graphed allowable downtime does not exceed the graphed progress through the compliance period is in a compliant state

Note that once an SLA reaches a critical state, it will remain in that state until the compliance period has restarted the following week or month; an SLA that enters a warning-level state can be downgraded to a normal state if the rate at which allowable downtime is used decreases to a "safer" value.

Generating an SLA Detailed Report

Clicking an SLA's corresponding **Detailed Report** button instantly generates an SLA Detailed report for the last 24 hours.

See "Reports for Service Level Agreements" on page 215 for more information.

SLA View Types

The **Service Level Agreements** subpanel provides two types of views: Condensed View and Detailed View. The latter view is suitable if you have one or two defined SLAs.

Condensed View

The following image illustrates the Condensed View of the **View SLAs** subpanel:

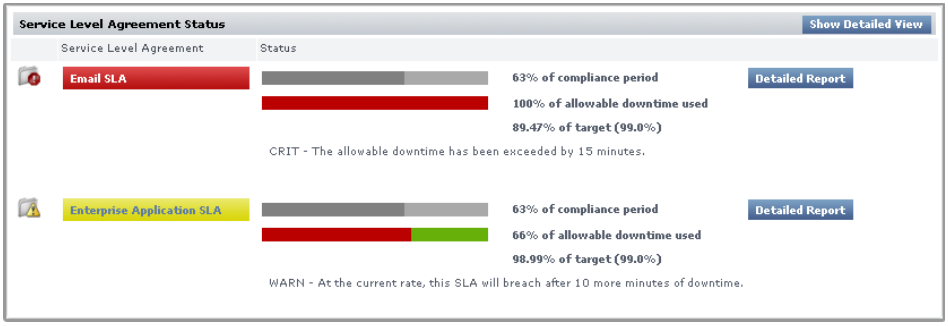

The Condensed View is the default view of this subpanel and displays the following information:

- the name of the SLA
- a status breakdown of the SLA for the current time period:
  - time period elapsed
  - available downtime used for the current time period
  - how close the SLA is to its performance target

- status message

Detailed View

Click the **Show Detailed View** button to expand each SLA to include SLOs:

An SLA's compliance is based on the downtime of its component SLOs: when one or more of the SLOs experience downtime, it counts towards overall SLA non-compliance.

Clicking an SLO name displays the status of the SLO, and all of the services that make up the SLO:

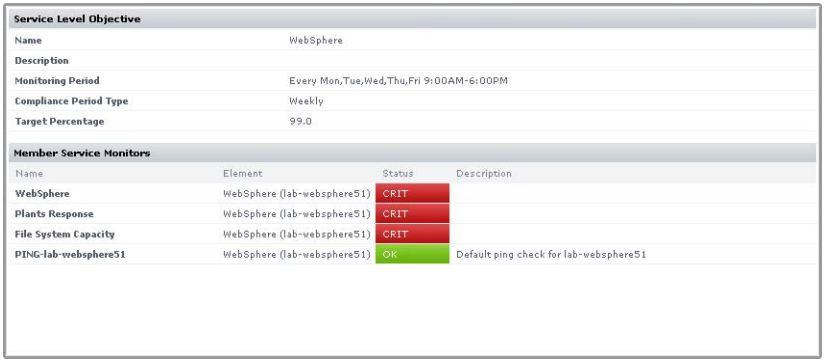

Using the Detailed View allows you to pinpoint which SLO is causing SLA non-compliance, and in turn which monitors are causing the SLO to experience downtime.

For more information about viewing SLA details, and defining SLOs that help you accurately gauge the performance of your IT infrastructure, see “Working with Service Level Agreements” on page 97.

Viewing All Applications

Applications provide the overall status for one or more services that up.time monitors. Applications group services, such as ping checks and checks for the status of the up.time agents that are installed on a system. An Application can contain many services, and enable you to better analyze component outages versus true Application outages.

An Application consists of:

- master service monitors

  One or more monitors can be used to determine the status of the Application as a whole.

- regular service monitors

  Other service monitors that are associated with a master service monitor, but are not used to determine the status of the Application as a whole.

The status of each Application is color coded:

- Applications highlighted in green are functioning normally
- Applications highlighted in yellow are in a warning state
- Applications that are in a critical state (when one or more master service monitors reaches a critical state) are highlighted in red and include the critical icon ()

The color coding also indicates whether an Application is offline or is in scheduled maintenance:

- an Application that is offline is highlighted in red and marked by the offline icon, and a message indicating that the Application is offline appears in the **Applications** subpanel
- an Application that is in scheduled maintenance is grayed out, the message `System is in scheduled maintenance` is displayed in the **Applications** subpanel, and the Application is marked with the scheduled maintenance icon ()

The **Applications** subpanel displays the status of each Application that you have added to, or are replicating on, the EMS. For replicated Applications, you are presented with basic status information, and need to log in to the actual LDC Monitoring Station to view more detailed information.

This subpanel has two views: Condensed View and Detailed View.

Condensed View

The following image illustrates the Condensed view of the **View Applications** subpanel:

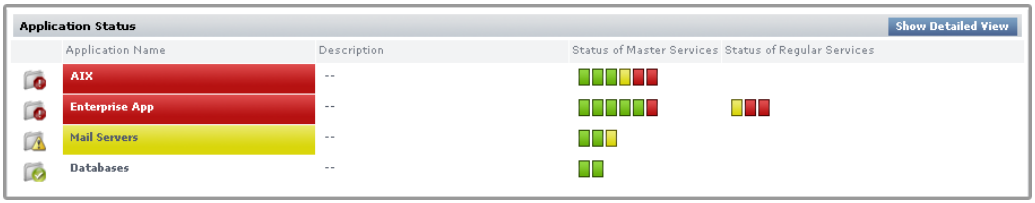

The Condensed view is the default view for this subpanel and displays the following information:

- the name of the Application
- a description of the Application, if one was added when the Application was defined
- the overall status of the Application if it is being replicated from another Datacenter
- the status of each service in the Application if it was created on the Enterprise Monitoring Server

  The status of the service is denoted by a colored bar in the **Status of Master Services** and **Status of Regular Services** columns. For example, if there are three services associated with the Application and their status is OK then three green bars appear in this column.

Detailed View

Click the Show Detailed View button to change to the Detailed view of the **View Applications** subpanel, as illustrated below:

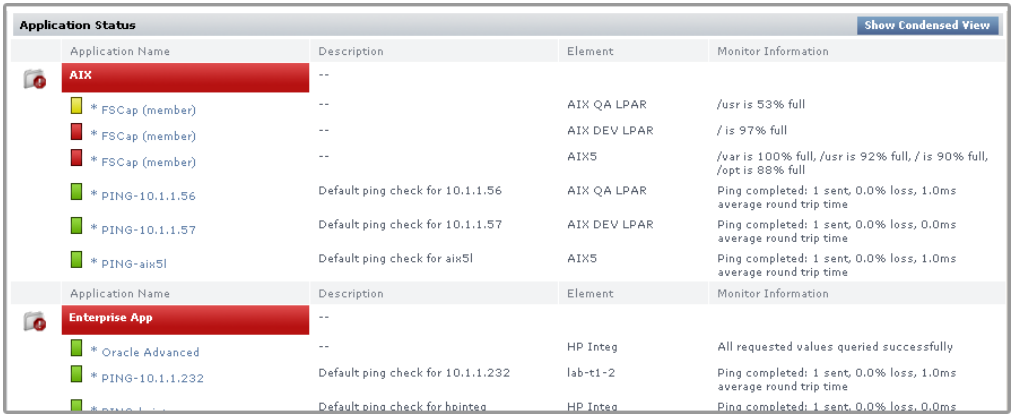

The name of the master Application group is in the far left column – for example, `Databases` in the image above. The names of the individual Applications are in the columns on the right – for example, `PING-mckay` and `UPTIME-mckay` in the image above. Master service monitors in an Application are marked with an asterisk (*).

The status of a service is denoted by a colored bar beside the name of the service – green for services that are functioning normally; yellow for services that are in a warning state; and red for services that are in a critical state.

The name of each Application is a hyperlink. Click a link to view detailed information about an Application. For details about the Application information that is displayed, see “Viewing System and Service Information” on page 38.

Viewing All Elements

Elements are the replicated systems, network devices, Applications, and SLAs that up.time is currently monitoring. In the **Global Scan** panel, you can view the status of all monitored Elements in the **All Elements** subpanel. This can be accessed by clicking the **View All Elements** tab.

The following image illustrates the **View All Elements** subpanel:

Elements

Name	Service Status						Outages			CPU				Disk		Memory
	OK	WARN	CRIT	MAINT	UNKN	ACK	1hr	12hr	24hr	USR	SYS	WIO	TOT	% Used	% Busy	Swap Used
AIX DEV LPAR	2	0	0	0	0	0	2	9	2	20%	15%	0%	35%	97%	1%	1%
AIX QA LPAR	2	0	0	0	0	0	0	5	8	18%	3%	0%	21%	53%	7%	9%
AIX5	2	0	0	0	0	0	0	2	4	14%	75%	1%	90%	100%	2%	15%
bogus	0	0	2	0	0	0	0	1	1							
Dev's SLES	2	0	0	0	0	0	2	6	5	0%	3%	0%	3%	33%	0%	0%
Dev1-w2k3se	0	0	1	0	1	0	0	0	0	The availability check has failed.						
DNS and The Fogz	3	0	2	0	0	0	0	12	5	No performance data available. Is the System online?						
ELinux	2	0	0	0	0	0	0	10	8	0%	0%	0%	0%	20%	0%	5%
ESX4	1	0	0	0	0	0	0	8	6	4%	1%	0%	5%	82%	0%	0%
ESX7	1	0	0	0	0	0	1	5	6	21%	41%	0%	62%	82%	0%	0%
Exchange	3	0	0	0	0	0	0	0	0	0%	0%	0%	0%	63%	1%	2%
FilterSNMP	1	0	0	0	0	0	0	7	14	3%	4%	1%	8%	91%	n/a	0%
Ginger Agent	3	0	0	0	0	0	1	4	9	0%	1%	0%	1%	64%	0%	0%
GingerSNMP	0	0	0	0	0	0	0	0	0	No performance data available. Is the System online?						
HMC Managed Server	0	0	0	0	0	0	0	0	0							
HP Integ	4	0	0	0	0	0	0	13	12	15%	39%	1%	55%	69%	51%	32%
Mail Server	2	0	0	0	0	0	0	3	13	1%	3%	0%	4%	63%	1%	2%
McKay	1	0	0	0	0	0	0	5	2							
my app	2	0	0	0	0	0	0	0	0							
MyMachine	3	0	0	0	0	0	1	10	10	2%	0%	0%	2%	42%	13%	28%
Novell	1	0	0	0	0	0	0	6	6	n/a	n/a	n/a	1%	0%	n/a	83%
Opteron	3	1	0	0	0	0	0	0	0	2%	3%	0%	5%	84%	0%	43%
Perl	0	0	0	0	0	0	0	0	0							
QA RedHat Instance	2	0	0	0	0	0	1	6	5	1%	21%	0%	22%	45%	8%	2%

The **View All Elements** subpanel lists the following information:

- the names of the Elements in your enterprise that are replicated on the EMS (including the source Local Datacenters' prefix names)
- the status of the services that are assigned to each Element
- the number of outages over the last hour, 12 hours, and 24 hours
- the percentage of CPU resources being consumed by users, the system, and by disk I/O
- the percentage of the system disk that is being used and the percentage that is busy

- the amount of memory swap space that is being used

If up.time cannot contact an Element, then the following message is displayed:

```
The availability check has failed
```

The values in each column are hyperlinks. Click one of the links to display the following information in the system information or graphing subpanels:

- Click any value in the **OK**, **WARN**, **CRIT**, **MAINT**, or **UNKNOWN** columns to open the **Status** subpanel. See “Status” on page 40 for more information.
- Click any value in the **Outages** column to open the **Outages** subpanel. See “Outages” on page 41 for more information.
- Click any value in the **USR**, **SYS**, **WIO**, or **TOT** columns to open the **Usage% Busy** report subpanel. For more information, see “Usage (% busy)” on page 253 for more information.
- Click any value in the **% Used** column to open the **File System Capacity** report subpanel. See “File System Capacity Graph” on page 280 for more information.
- Click any value in the **% Busy** column to open the **Disk Performance Statistics** report subpanel. See “Disk Performance Statistics Graph” on page 276 for more information.

Viewing All Services

Services are specific tasks, or sets of tasks, performed by an application in the up.time environment. up.time service monitors continually check the condition of services to ensure that they are providing the required functions to support your business.

You can view the services assigned to each system in your environment by clicking on the **View All Services** tab. This tab contains the following information:

- the name of the service
- the monitor that is associated with the service
- the status of the service
- the date and time on which the last check was performed
- the number of days, hours, and minutes since the last check
- a human-readable text message that was returned by the monitor (e.g., "`up.time agent running on MailServer, up.time agent 3.7.2 linux`")

Viewing the Resource Scan Report

Resource Scan is a dynamically-updated report that charts the percentage of various resources that are being used by the systems in your environment. You can view this report by clicking the **View Resource Scan** tab.

Resource Scan is divided into three sections – a set of performance gauges, 24-hour performance graphs, and an Elements chart.

As you click through lists in the Resource Scan report, the status reported in the gauges and charts reflects your current view, whether it is focused on parent groups, nested groups, or individual Elements.

Performance Gauges

There are two sets of gauges that are updated every 15 minutes with new data. The top row of gauges displays an average of the most recent 15-minute time frame; the bottom row of gauges displays a minimum, maximum and average value for the last 24-hour period, up to the most recent 15-minute time frame. The gauges show the following information:

- CPU Usage

  The percentage of the system's CPU resources that are being used.

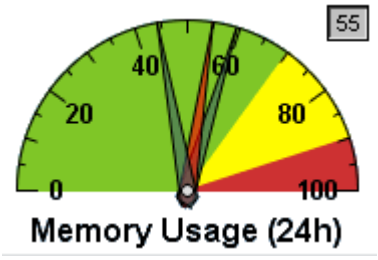

- Memory Usage

  The amount of memory, expressed as a percentage of total available memory, being consumed by a process.

- Disk Busy

  The percentage of time that the disk is handling transactions in progress.

- Disk Capacity

  The percentage of space on the system disk that is being used.

24-Hour Performance Graphs

The 24-hour gauges display a minimum, maximum, and average value; the full 24-hour performance history is displayed in the graphs below:

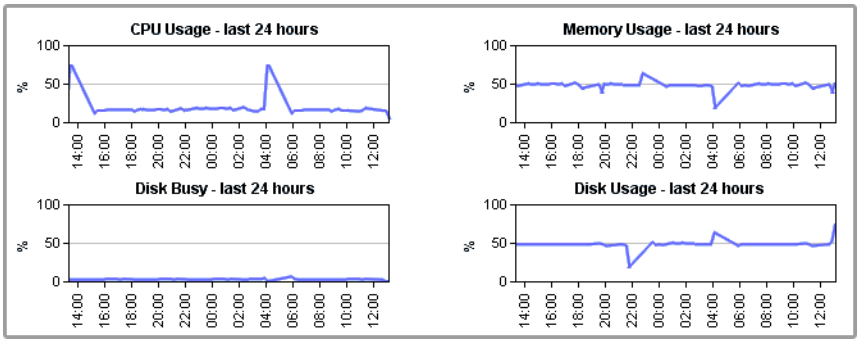

Elements Chart

The Resource Scan chart displays the following information for all of the Elements in your environment:

- CPU Usage

  The percentage of CPU resources that are being used.

- Memory Usage

  The amount of memory, expressed as a percentage of total available memory, that is being consumed by a process.

- Disk Capacity

  The percentage of storage space on the system disk that is being used.

- Network In

  The average amount of traffic coming in over the network interface.

- Network Out

  The average amount of traffic going out over the network interface.

The following image illustrates the Resource Scan chart:

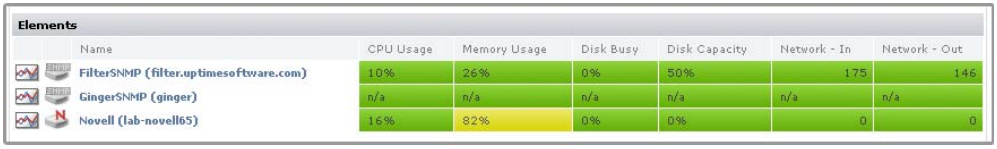

Elements						
Name	CPU Usage	Memory Usage	Disk Busy	Disk Capacity	Network - In	Network - Out
FilterSNMP (filter.uptimesoftware.com)	10%	26%	0%	50%	175	146
GingerSNMP (ginger)	n/a	n/a	n/a	n/a	n/a	n/a
Novell (lab-novell65)	16%	82%	0%	0%	0	0

You can view the Resource Scan gauges for a particular server by clicking the name of the server in the chart.

If you have grouped your servers, the names of individual servers do not appear in the Resource Scan chart. Instead, the names of the groups are displayed. To view a list of Elements in a group, click the name of the group.

When viewing a Resource Scan for a system, you can navigate to other groups by selecting the name of the group from the **Current Location** dropdown list at the top of the **Resource Scan** panel, as shown below:

Viewing Scrutinizer Status

Scrutinizer is a NetFlow analyzer that takes advantage of communications standards for Cisco IOS networking devices, as well as other compatible switches and routers, to retrieve and store network traffic information for users, systems, and applications. It allows administrators to monitor, graph, and report on network usage patterns, and locate the heaviest traffic creators.

Scrutinizer can be integrated with up.time. Doing so allows you to add node-type Elements that are exporting NetFlow data to Scrutinizer, as well as call a Scrutinizer instance from a commonly-monitored Element's status page (whether the Element is a NetFlow-exporting node, or a non-node Element).

You can also access all of Scrutinizer's features, such as the MyView status panel, from within **Global Scan** by clicking the **NetFlow** tab:

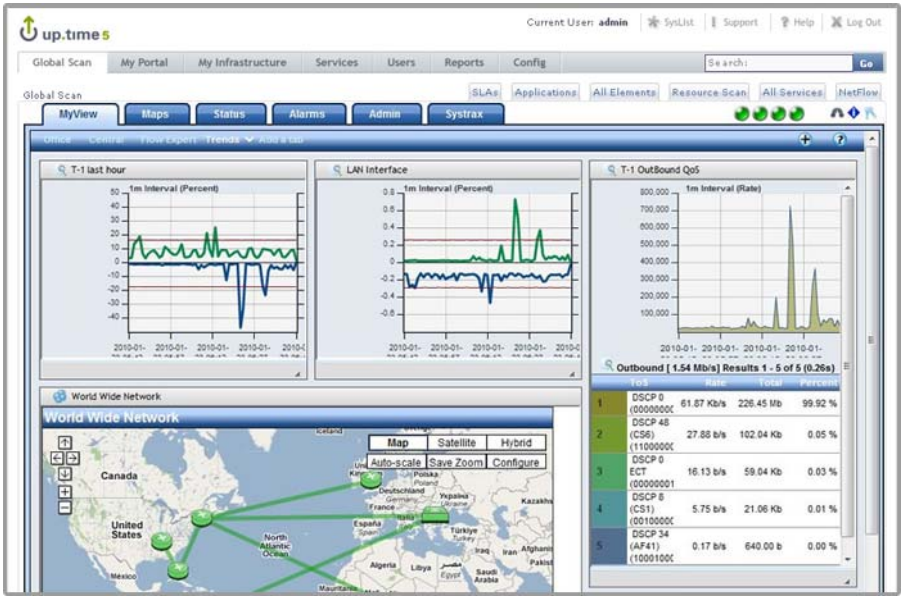

Changing Reporting Thresholds

The thresholds that determine when an Element's reported status changes between normal, Warning, and Critical (i.e., green, yellow, and red) can be modified for both **Global Scan** and the Resource Scan.

Global Scan and the Resource Scan thresholds are configured by separate sets of attributes that can be changed in the **up.time Configuration** panel. By changing these attributes, you can set how large the color ranges are on resource gauges, and at what point table cells change color. See "Status Thresholds" on page 309 for more information.

Note that when you change **Global Scan** threshold values, the changes are not retroactively applied to all existing Elements monitored by up.time; changes only apply to Elements added to up.time after the threshold changes are made. Conversely, the Resource Scan gauge ranges are updated immediately.

up.time

CHAPTER 8

Working with Service Level Agreements

This chapter explains how to configure up.time to monitor for compliance with Service Level Agreements (SLAs) in the following sections:

Overview

In up.time, a service level agreement (SLA) measures your IT infrastructure's ability to meet performance goals, particularly from the end-user perspective. Different goals can focus on different aspects of your infrastructure from underlying network performance, to back-end database availability, to user-facing application server response time. Given this broad coverage, a performance goal encompasses anything from a handful of monitored systems to an entire production center.

Defining and working toward fulfilling SLAs provides you with more insight into the performance and planning of your infrastructure:

- measure the performance of your infrastructure from the end-user perspective

  An SLA can measure the success of your IT infrastructure by using end-user-focused service monitors such as the Web Application Transaction monitor and the Email Delivery monitor.

- translate IT infrastructure demands into quantifiable and reportable goals

  Use SLAs to methodically set expectations on all or the most critical aspects of your infrastructure. SLAs provide you with metrics with which you can gauge the success of your network administration.

- use trends to anticipate new infrastructure requirements

  Trend lines in SLA reports can give you an estimate for when your current hardware deployment will require augmentation.

- generate SLA reports that demonstrate compliance and break down objectives

  Compliance reports quantify the value of the IT department's efforts, and objective-based reports exist to identify recurring problems that affect business outcomes.

SLAs, Service Monitors, and SLOs

Like other up.time Elements (i.e., systems, network devices, and Applications) an SLA definition consists of service monitors that you have previously created. Depending on its use, an SLA can consist of a single service level objective (SLO) that in turn consists of a single service monitor.

In other cases, an SLA's coverage can be broad enough to include an ungainly list of service monitors; in this case the SLA can be refined to consist of multiple SLOs that focus on different aspects of the SLA. Creating multiple objectives helps you further refine your performance targeting and reporting.

For example, consider an SLA called "Web Application" that focuses on IT performance for end users. The SLA's objectives could be broken down by performance:

- SLO 1, application availability: the application is available 99% of the time (e.g., using an HTTP monitor)
- SLO 2, application speed: the application's Web transactions always complete in fewer than 10 seconds (e.g., using the Web Application Transaction monitor)

Consider another example: an SLA called "Customer Service Group" that focuses on the operational readiness of a support team. The SLA's objectives could be broken down by application:

- SLO 1: helpdesk application
- SLO 2: bug-tracking application
- SLO 3: email service

Viewing Service Level Agreements

Service level agreements, and the type of information displayed, are viewed in the **Global Scan** panel from a monitoring perspective, and in **My Enterprise** from a configuration perspective.

Viewing SLA Status

You can view the status of all your SLAs in the **Service Level Agreements** subpanel, which can be accessed by clicking the **View SLAs** tab when you are in the **Global Scan** panel.

For more information about what kind of SLA information you can view in the **Global Scan** panel, see “Viewing All SLAs” on page 80.

Viewing SLA Details

The details of an SLA definition can be viewed in the **Service Level Agreement General Information** subpanel. This can be accessed from the **My Enterprise** panel by clicking the SLA name listed among the

Elements, or from the **Global Scan** panel by clicking the **Info** tab in the Tree panel, then clicking **Info**:

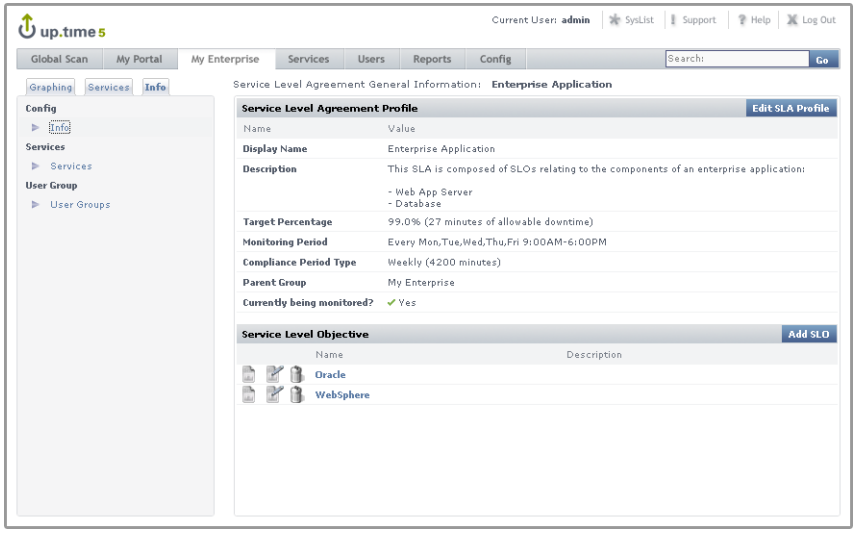

The **General Information** subpanel displays a summary for the SLA that includes the following:

- Target Percentage: the targeted percentage of up time of the SLA's component services over the Monitoring Period
- Monitoring Period: the days and time frames during which uptime is measured
- Compliance Period Type: the compliance period intervals over which SLA compliance is measured (i.e., weekly or monthly)
- Service Level Objectives: a listing of the SLOs into which the SLAs services have been organized

For more information about system information in general, see "Viewing System Information" on page 38.

You can view information about the services that make up the SLA by clicking the **Services** tab in the Tree panel. The options available in the Tree panel are summarized in "Viewing Service Information" on page 40.

Clicking the **Graphing** tab in the Tree panel, then clicking **Current Status** displays a verbose status summary of the SLA that includes the following:

- Trend Analysis: SLA status indicator for the current compliance period

- Compliance Period and Allowable Downtime Used: the current progress through the compliance period, and how close the SLA is getting to reaching a critical state
- Achieving (SLA): how close the SLA is to its performance target; how recoverable a failing SLA is, based on how far it is from its target
- Achieving (SLOs): an SLO-level breakdown of how well or poorly each SLO is meeting its performance target; how recoverable failing SLOs are, based on how far it is from its target

See “A Note About SLOs and Compliance” on page 105 for more information about SLOs and the Achieving statistic.

SLA Compliance Calculation

SLA downtime occurs when any of the SLA's services are in a critical state. An SLA is compliant if its downtime has not exceeded a maximum number of minutes over a one-week or one-month Monitoring Period.

For example, consider an SLA whose compliance period type is weekly and its Monitoring Period is Monday through Friday, 9 p.m. to 5 p.m. The Monitoring Period consists of five eight-hour days—in other words, 40 hours, or 2400 minutes. If the SLA's target is 95%, it has 120 minutes of allowable downtime for any of its services.

Reporting SLA Status

An SLA's reported status in the **Global Scan** panel includes the following in the form of progress bars: the percentage of the Monitoring Period that has expired, and the percentage of allowable downtime consumed during the Monitoring Period. (See "Viewing All SLAs" on page 80 for information about SLA information in the **Global Scan** panel.)

An SLA will reach a critical state when its allowable downtime has been depleted. An SLA will reach a warning-level state when its allowable downtime, at the current rate of use, will be depleted before the compliance

period has ended. These states, and their conditions under which they happen, are shown in the **Global Scan** status display:

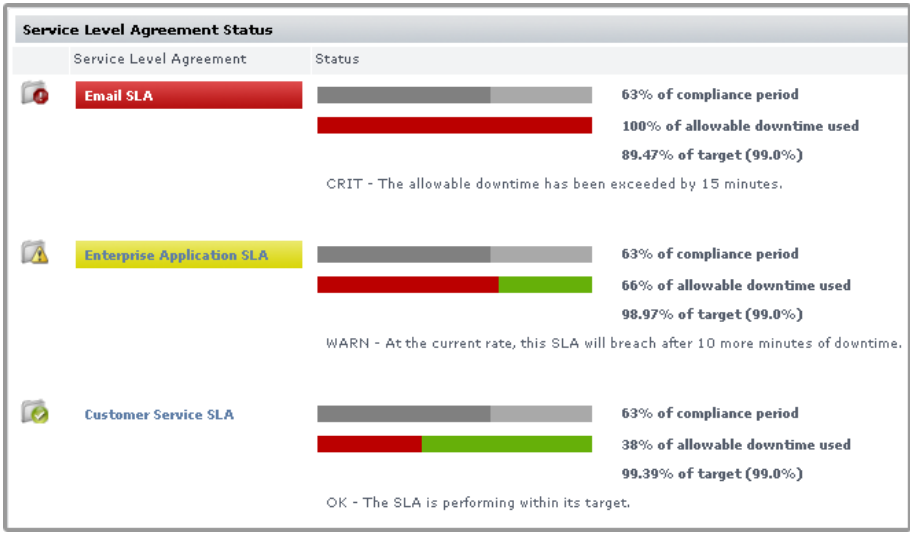

Handling Simultaneous Service Downtime

The simultaneous downtime of multiple services does not cumulatively impact an SLA's remaining allowable downtime; the term "allowable downtime" can be expanded to mean the amount of time during which there can be any service downtimes (until the compliance period has ended, after which the counters are reset).

In the following outage graph for an SLO, note that any time an outage is experienced—whether by one or four services—the SLO is deemed to have experienced an outage, which is reflected in the top red line:

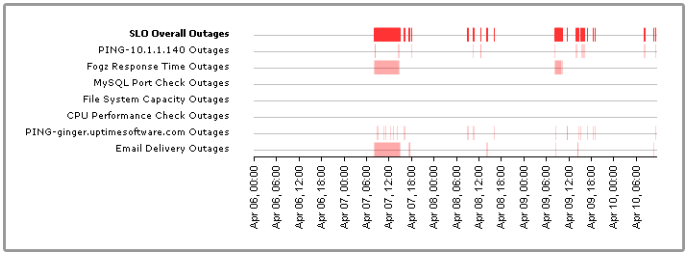

A Note About SLOs and Compliance

It is important to note the role an SLO plays regarding SLA compliance: SLOs exist to help you conceptually separate services into logical groups that make it easier for you to monitor, diagnose, and set performance goals for them. Although the descriptions of "allowable downtime" in the previous section implied that service downtime affects SLA downtime, it is more accurate to say that service downtime affects SLO performance—which in turn, affects SLA downtime.

SLO outages affect reported SLA compliance in the same way service outages affect SLO compliance: allowable downtime is reduced when any outage is experienced. This is also pertinent if you are scanning the "Achieving" statistic for an SLA Summary. (This statistic can be viewed in the **Service Level Agreement** subpanel of **My Enterprise**, by clicking the **Graphing** tab, then clicking **Current Status**.)

You can verify how well or poorly an SLA is achieving its target, but you can also view how the component SLOs are performing for the time period. In the following example, the email server performance SLO is achieving 90.03% of its 99.0% target. Although the email server availability SLO is achieving its target (99.43% vs. 99%), both SLOs' downtime affects SLA downtime. In thise case, combined SLO downtime results in the SLA only achieving 89.47% of its target—resulting in a critical status.

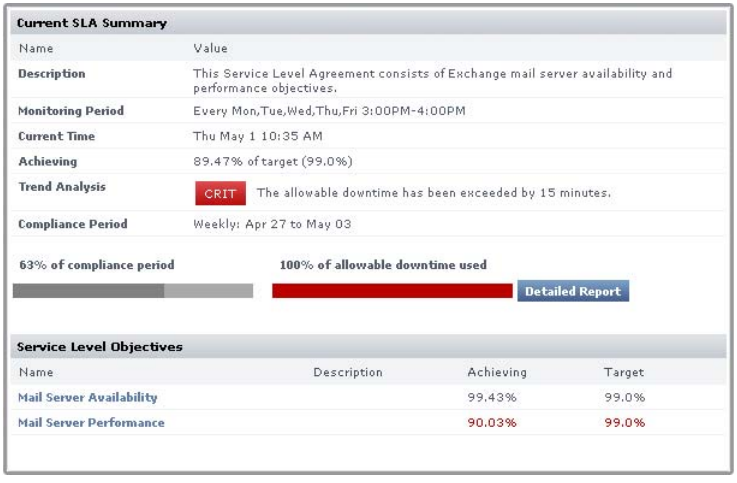

See "Viewing SLA Details" on page 100 for information on how to find information such as the Achieving statistic in an SLA summary.

SLA-Creation Strategies

The key to an effective SLA is defining a service level that satisfies end users, yet is also attainable by IT staff and their systems configurations. This section covers the suggested steps to pinpointing this target service level:

- ensure service monitors exist for all SLA-related Elements (if you are a new up.time user, all of these will need to be created)
- define an SLA and its objectives
- use the SLA Detailed report to identify and resolve outages or underperforming Elements
- use the SLA Summary report to develop a baseline

Setting Up and Gathering Data for Monitors

Determine which service monitors will best reflect the end-user experience, based on the aspect of your infrastructure that your SLA will cover. See “SLAs, Service Monitors, and SLOs” on page 99 for some sample SLAs and objectives.

up.time users who do not have existing service monitors should create them and allow them to accumulate data for at least one week. Having historical data is essential to determining what level of service you should target.

Identifying Outages and Improvable Performance

When added to an SLA, service monitors that have been collecting data will immediately contribute to the SLA’s reported status. For example, if all of an SLA’s service monitors have a year’s worth of historical data, creating a trial SLA will allow you to see how it would have performed over that last year. Having this historical data in SLA reports helps you analyze each component service monitor in the context of the SLA.

Consider a sample SLA called System Performance that is meant to ensure your application servers are not experiencing excessive loads; this can be indicated by CPU usage and disk space. The first service level objective is

based on the Performance Check monitor for the application servers. A critical state occurs when CPU usage exceeds 90%. The second service level objective is based on the File System Capacity monitor. A critical state occurs when remaining disk space falls under 10%.

After creating an SLA based on these objectives, the SLA is immediately shown to be in a critical state—for the current Monitoring Period, one or both of the objectives have already failed to meet the defined service level:

You can investigate outages using the SLA Detailed report. In this example, you determine that the cause the SLA failure was a prolonged disk-space-related outage that, based on the outage graph, appears to have been resolved:

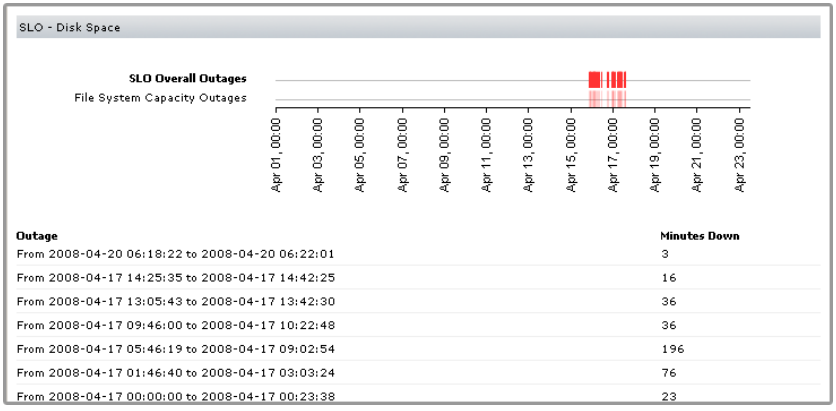

Outage	Minutes Down
From 2008-04-20 06:18:22 to 2008-04-20 06:22:01	3
From 2008-04-17 14:25:35 to 2008-04-17 14:42:25	16
From 2008-04-17 13:05:43 to 2008-04-17 13:42:30	36
From 2008-04-17 09:46:00 to 2008-04-17 10:22:48	36
From 2008-04-17 05:46:19 to 2008-04-17 09:02:54	196
From 2008-04-17 01:46:40 to 2008-04-17 03:03:24	76
From 2008-04-17 00:00:00 to 2008-04-17 00:23:38	23

However, there may be cases were analyzing the SLA Detailed report will show intermittent outages that have not caused your trial SLA to fail, but represent underperforming services that should be optimized:

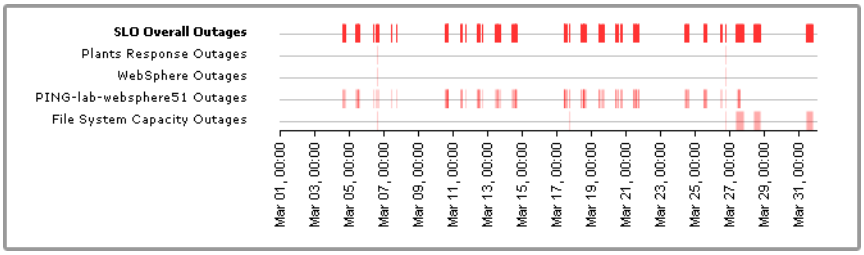

Developing Baselines

After outages and underperforming systems have been addressed, use the SLA Summary report to compare test service levels to historical data.

Find a service level that is attainable. For example, in the SLA graph below, a 95% service level would be more realistic than the default 99% level, given the historical data. In the bottom SLA graph, although the 90% service level is compliant based on historical data, the performance history

shows that a 95% service level is attainable if the IT department is able to isolate and improve key underperforming systems.

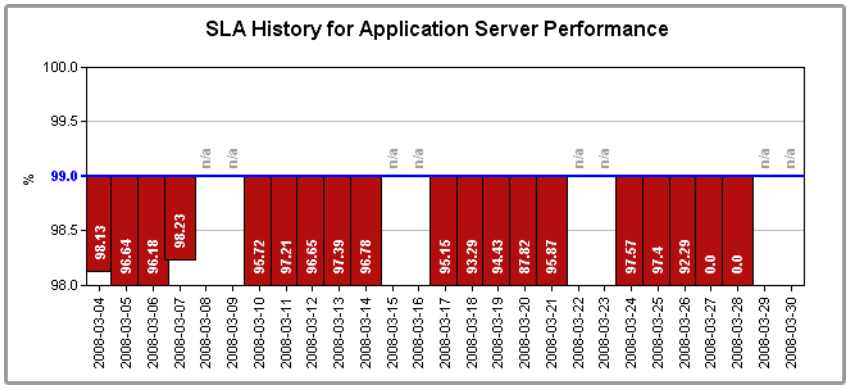

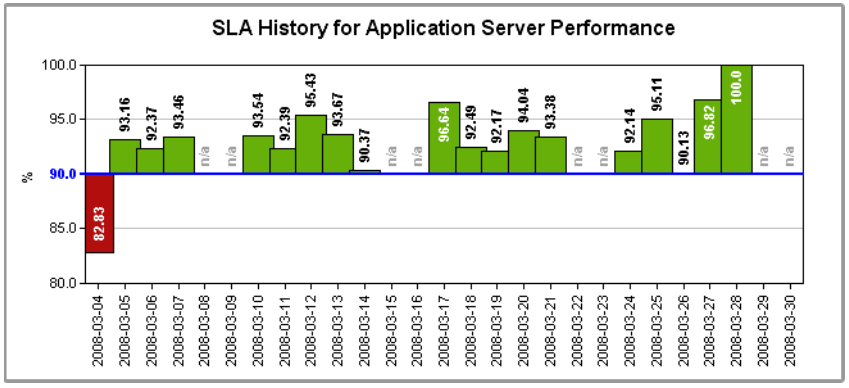

Working with SLA Reports

up.time provides two types of SLA reports. The SLA Summary report provides high-level SLA compliance information, and the SLA Detailed report provides SLO- and service-level compliance information for system administrators.

See “Reports for Service Level Agreements” on page 215 for more information.

Adding and Editing SLA Definitions

Adding and using an SLA requires that you first define the SLA, then add one or more SLOs to it.

When you create an SLA, it will be inserted into the current compliance period. For example, a newly created SLA that reports over a monthly compliance period will, if created on the 15th of the month, already be around 50% through the period.

Adding a Service Level Agreement

To add a service level agreement to up.time, do the following:

1 **In the My Enterprise panel, click Add Service Level Agreement.**

The Add Service Level Agreement window appears:

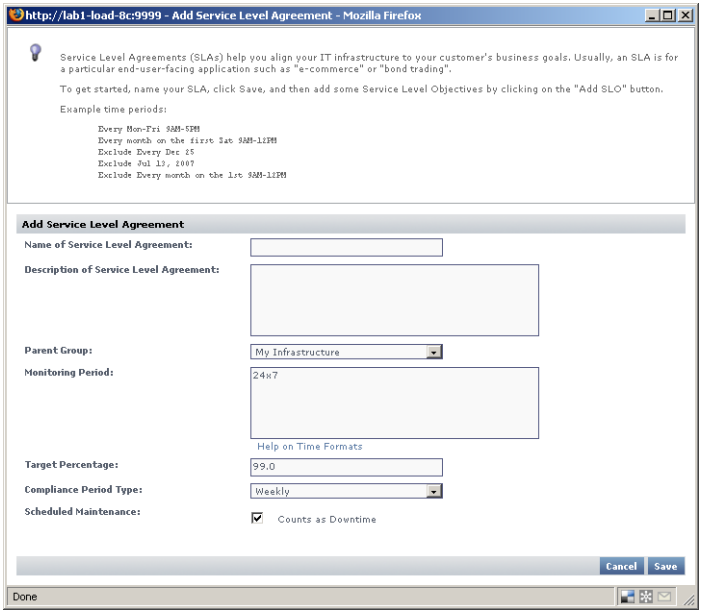

2 **Enter a descriptive name for the SLA in the Name of Service Level Agreement field.**

This name will appear in both the **My Enterprise** and **Global Scan** panels.

3 **Optionally enter a description for the SLA in Description of Service Level Agreement field.**

Although this step is optional, this description will appear in generated SLA reports; therefore, it is recommended that you provide a detailed description of the SLA including what it is meant to accomplish and of which SLOs it consists.

4 **Optionally select the group of systems in your up.time environment with which this system will be associated from the Parent Group dropdown list.**

By default, the SLA is added to the **My Enterprise** group.

For more information on groups, see "Working with Groups" on page 66.

5 **If it is not continuous (i.e., "24x7"), enter a Monitoring Period during which the SLA's compliance will be measured.**

You will need to create a time period definition (e.g., "`Every Mon-Sat 8AM-6PM`"). See "Monitoring Periods" on page 160 and "Time Period Definitions" on page 321 for more information.

6 **If it is not the default 99.0%, enter a Target Percentage against which the SLA's compliance will be measured.**

7 **Ensure you have selected the correct Compliance Period Type from the dropdown list.**

8 **Indicate whether scheduled system maintenance will count as downtime.**

For Elements that are part of the SLA definition, Maintenance Windows are created and assigned at the LDC level.

9 **Click Save.**

Once saved, the SLA's **Service Level Agreement General Information** subpanel is displayed (see "Viewing SLA Details" on page 100 for more information). From this page, you can add SLOs, as well as associate Alert Profiles and Action Profiles to the SLA.

Adding Service Level Objectives to an SLA

To add a service level objective to an SLA, do the following:

1. **In the My Enterprise panel, click the name of the Service Level Agreement that you want to edit.**

   The **Service Level Agreement General Information** subpanel appears.

2. **Click Add SLO.**

   The Add Service Level Objective window appears:

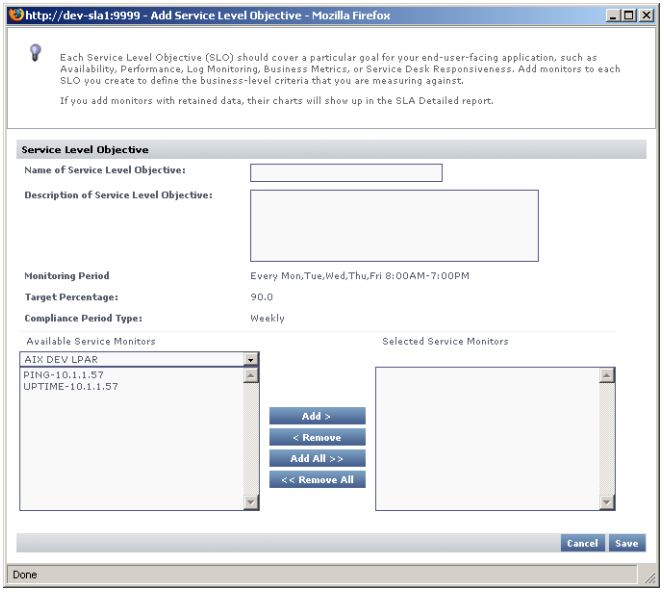

3. **Enter a descriptive name for the SLO in the Name of Service Level Objective field.**

   This name will appear anywhere in **My Enterprise** and **Global Scan**.

4. **Enter a description for the SLO in Description of Service Level Objective field.**

   Although this step is optional, this description will appear in SLA Detailed reports; therefore, it is recommended that you provide a detailed description of the SLO including what goal is being accomplished, and of which service monitors it consists.

5. **Add a service monitor that will be associated with the SLO by first selecting its host from the dropdown list, then adding the service monitor.**
6. **Continue to add service monitors to the SLO as required.**
7. **Click Save.**

Associating Alert and Action Profiles to an SLA

To add a service level objective to an SLA, do the following:

1. **In the My Enterprise panel, click the name of the Service Level Agreement that you want to edit.**

   The **Service Level Agreement General Information** subpanel appears.
2. **Associate Alert Profiles with the SLA by clicking Edit Alert Profiles.**
3. **In the Alert Profile Selector pop-up window, select one or more of the Available Alert Profiles from the list, then click Save.**
4. **If required, associate Action Profiles with the SLA by clicking Edit Action Profiles.**
5. **In the Action Profile Selector pop-up window, select one or more of the Available Action Profiles from the list, then click Save.**

Editing SLA and SLO Definitions

To edit a service level agreement, do the following:

1. **In the My Enterprise panel, right-click the name of the Service Level Agreement that you want to modify, then click Edit.**

   The **Edit Service Level Agreement** window appears.
2. **Edit the SLA as described in the previous section.**

   See “Adding a Service Level Agreement” on page 111 for information.

Since SLA reporting and monitoring is based on weekly or monthly compliance periods, changing any of the following on an existing SLA affects the reported SLA status and generated reports:

- Monitoring Period
- target percentage
- compliance period type

Any changes made are immediately reflected in any SLA reporting.

To edit a service level objective, do the following:

1. **In the My Enterprise panel, click the name of the Service Level Agreement that you want to modify, then click Edit.**

   The **Service Level Agreement General Information** subpanel appears.

2. **Click the SLO's corresponding Edit icon ( ).**

3. **Edit the SLO as described in the previous sections.**

   See "Adding Service Level Objectives to an SLA" on page 113 for information.

   Since SLA reporting and monitoring is based on weekly or monthly compliance periods, changing the service monitors that make up an SLO definition will affect the reported SLA status and generated reports.

Any changes made are immediately reflected in any SLA reporting.

up.time

CHAPTER 9

Configuring Users

This chapter describes the up.time user management functions in the following sections:

Working with User Roles

User roles define the following:

- what a user will see when they log in to the up.time Monitoring Station
- the items that a user can add, view, edit, or delete when using the Monitoring Station

The user roles that you create should reflect that needs of the users to whom the roles will apply. For example, a user who only needs to generate graphs and reports does not need to be able to view or add accounts for other up.time users.

Adding User Roles

To add user roles, do the following:

1. **On the up.time tool bar, click Users.**
2. **In the Tree panel, click Add New User Role.**

   The **Add User Role** window appears.
3. **Type a name for this role in the Name of User Role field.**

   This name will appear in the up.time Web interface.
4. **Optionally, type a short description in the Description of User Role field.**
5. **In the first Permissions area of the Add User Role window, you assign the user permissions to View, Add, Edit, or Delete the following items by clicking the checkbox beside each item:**

   - Users
   - Elements
   - Services
   - Element Groups
   - Action Profiles
   - Alert Profiles

- Time Periods
- Service Level Agreements
- Element Views

6 **Optionally, in the second Permissions area enable one or more of the following options by clicking the Allowed checkbox:**

- Administrator

  The user can perform all up.time administration tasks.

- Acknowledge Alerts

  The user can acknowledge an alert. See “Understanding Alerts” on page 142 for more information.

- Save Reports

  The user can save reports. Links to the saved reports will appear in the **My Portal** panel, or the user can save reports to a local or network drive. “Saving Reports” on page 166 for more information.

7 **Click Save.**

Viewing User Roles

You can view a user role to ensure that the permissions for the role are properly configured.

To view user roles, do the following:

1 **In the Tree panel, click View User Roles.**

A list of the user roles appears in the **Users** subpanel. Clicking a user role displays a table that summarizes the role’s configured permissions; those

which have been granted as denoted by a green check mark (✔), as shown below:

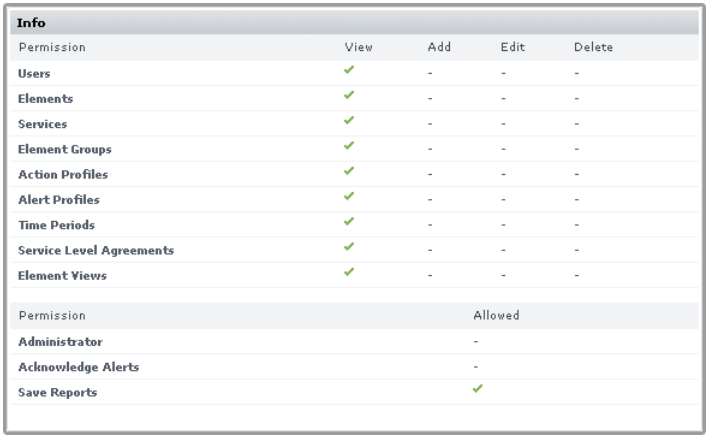

Info

Permission	View	Add	Edit	Delete
Users	✔	-	-	-
Elements	✔	-	-	-
Services	✔	-	-	-
Element Groups	✔	-	-	-
Action Profiles	✔	-	-	-
Alert Profiles	✔	-	-	-
Time Periods	✔	-	-	-
Service Level Agreements	✔	-	-	-
Element Views	✔	-	-	-

Permission	Allowed
Administrator	-
Acknowledge Alerts	-
Save Reports	✔

Editing User Roles

To edit user roles, do the following:

1. **In the Tree panel, click View User Roles.**
2. **Click the name of the user role that you want to edit, and then click Edit User Role in the Users subpanel.**

   The **Edit User Roles** window appears.
3. **Edit the user role information as described in the section “Adding User Roles” on page 118.**

Working with Users

Users are the individuals who have access to up.time and its various functions. You can grant permissions to users to do any or all of the following:

- view information about specific systems in your environment
- generate and save reports about specific systems
- receive alerts

Adding Users

To add users, do the following:

1. **In the Tree panel, click Add New User.**

   The **Add User** window appears.

2. **Type a name for the user, which will be used to log into up.time, in the Username field.**

   If you are using Active Directory or an LDAP directory to authenticate up.time users, the user name you input should be identical to the user's name in the central directory.

3. **If AD/LDAP is enabled for user authentication, leave the Password field blank; otherwise enter a password that will be stored in the up.time DataStore.**

   If using an AD or LDAP directory to authenicate users, up.time will refer to the directory for password information during user login. For more information, see "Changing How Users Are Authenticated" on page 133.

4. **If you have set a user password, re-enter it in the Confirm Password field.**

5. **Enter the full name of the user in the First Name and Last Name fields.**

6. **Optionally, enter the user's geographical location or department in the Location field.**

7. **If the user will be receiving alerts via email, enter the user's email address in the Email Address field.**

8 **Select one of the following options from the Time Period for Emailing dropdown list:**

- 24x7
- 9am to 5pm weekdays
- another Monitoring Period that you have previously created

9 **If the user will receive alerts on their cell phone or pager, enter the email address of the user's cell phone or pager in the Pager/Cellphone Address field.**

The email address takes the following format:

`<number>@mobile_provider_domain`

Where `<number>` is the user's cell phone number, and `mobile_provider_domain` is the Internet domain of the user's mobile phone service. For example, `1234567890@mymobile.com`.

10 **Select an option from the Time Period for Pager/Cellphone Messages dropdown list.**

The options are the same as the ones listed in Step 8.

11 **If the user will receive alerts via the Window messaging service, enter the name of the user's computer in User's Windows Desktop Hostname field.**

To receive popup alerts, you must enable the Windows messaging service on the user's computer. See "Enabling the Windows Messaging Service" on page 145 for information.

12 **Enter the workgroup or domain to which the user's computer belongs in the User's Windows Desktop Workgroup field.**

13 **Select an option from the Time Period for Windows Popups dropdown list**

The options are the same as the ones listed in Step 8.

14 **If the user will receive alerts, select the Should the user receive alerts? option.**

If you select this option, you must also enter information in the **Email Address** or **Pager/Cellphone Address** fields.

15 **If you selected the Should the user receive alerts? option in step 14, select one of the following options:**

- Alert on Critical

  The user receives an alert when up.time detects a critical problem with one or more of monitored services.

- Alert on Warning

  The user receives an alert when up.time detects a potential problem with one or more monitored services.

- Alert on Unknown

  The user receives an alert when up.time detects an error in the configuration of the monitor, or if up.time cannot execute the service check.

- Alert on Recovery

  The user receives an alert when the service recovers from an error – for example, an application, process or service restarts, or a server reboots.

16 **Click the Disable ActiveX Graphs option to display graphs using a Java applet instead of in 3D.**

ActiveX graphs are only available to users accessing up.time with Internet Explorer.

Do not select this option if the user is working with Internet Explorer.

17 **Click the Show Tips option to disable graphical tool tips on pages like View Notification Groups.**

18 **Select a role for the user from the User Role dropdown list.**

For more information on user roles, see the section “Working with User Roles” on page 118.

19 **In the Available User Groups field, select the user group to which this user will belong and then click Add.**

For more information on user groups, see the section "Working with User Groups" on page 125.

20 **Click Save.**

Viewing Users

To view users, do the following:

1 **In the Tree panel, click View Users.**

A list of users appears in the **Users** subpanel.

Editing User Information

To edit user information, do the following:

1 **Do one of the following:**

- Click the **Edit** icon () beside the name of the user.
- Click the name of the user whose information you want to edit, and then click **Edit User** on the **User Information page**.

The **Edit User** window appears.

2 **Edit the information as described in the section "Adding Users" on page 121.**

Working with User Groups

User groups are sets of up.time users who have been assigned similar privileges. These privileges enable the members of a group to do the following:

- work with specific systems or network devices
- receive up.time alerts from those systems and devices
- participate in any number of defined service alert monitoring escalation paths

A member of a user group can view either individual systems or multiple systems in a system group. The following diagram illustrates how user groups work in up.time:

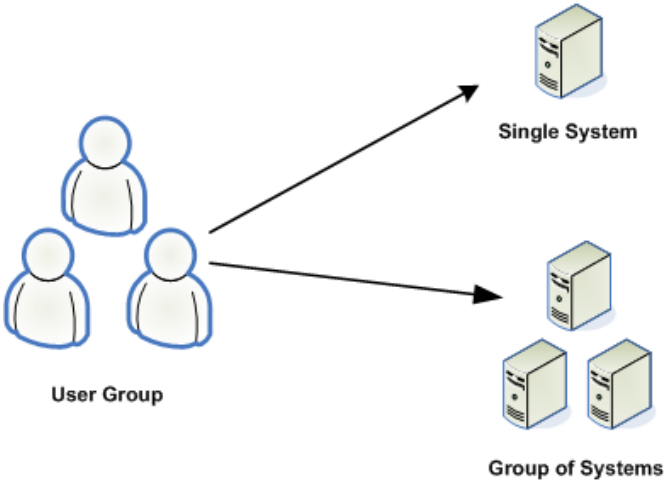

Each up.time user must belong to at least one user group. In a small installation of up.time there may only be one user and one user group. In larger installations, you can set up such user groups as Operators, Help Desk, System Administrators, Network Administrators, DBAs, Development, QA, Operations Management, and the like.

Adding User Groups

To add user groups, do the following:

1. **In the Navigation pane, click Add New User Group.**
2. **Enter a name for this group in the User Group Name field.**
3. **Optionally, type a short description in the User Group Description field.**
4. **Select the users to add to the group in the Available Users list, then click Add.**
5. **Optionally, select one of the systems or Elements from the Available Elements list, then click Add.**
6. **Optionally, select one of the groups from the Available Element Groups list, then click Add.**
7. **Optionally, select one of the views from the Available Entity Views list, then click Add.**
8. **Click Save.**

Viewing User Groups

To view user groups, do the following:

1. **In the Tree panel, click View User Groups.**

   A list of user groups appears in the **User Groups** subpanel.

Editing User Groups

To edit user groups, do the following:

1. **In the Tree panel, click View User Groups.**
1. **Do one of the following:**

   - Click the **Edit** icon () beside the name of the user group.
   - Click the name of the user group whose information you want to edit, and then click **Edit User Group** in the **User Group** subpanel.

The **Edit User Group** window appears.

2 **Edit the information as described in the section "Adding User Groups" on page 126.**

Deleting User Groups

To delete user groups, do the following:

1 **In the Tree panel, click View User Groups.**

2 **Click the Delete icon ( ) beside the name of the user group that you want to delete.**

You cannot delete the SysAdmin user group.

3 **On the warning dialog box that appears, click OK.**

Managing Distribution Lists

A Distribution List allows you to use an email alias to send alerts to end users who, aside from wanting to be informed of status alerts, have no other reason to use up.time. Using a Distribution List is an easy way to broadcast to a large group of users without having to create and manage individual up.time user profiles for each member.

Distribution Lists, like individual user profiles, are associated with Notification Groups, and can be configured to broadcast specific types of status alerts (e.g., only Critical-level and Recovery alerts).

Adding Distribution Lists

To add Distribution Lists, do the following:

1. **Click Users on the up.time tool bar.**
2. **In the Tree panel, click Add New Distribution List.**
3. **Type a descriptive name in the Display Name field.**

   You will select this name when defining a Notification Group.
4. **Select a Monitoring Period from the Time Period for Emailing list:**

   - 24x7
   - 9am to 5pm weekdays
   - another Monitoring Period that you have previously created
5. **Select the Should the Distribution List receive alerts? check box.**
6. **Configure the type of alerts those on the Distribution List will receive by selecting one or more of the following check boxes:**

   - Alert on Critical

     The user receives an alert when up.time detects a critical problem with one or more monitored services.

- Alert on Warning

  The user receives an alert when up.time detects a potential problem with one or more monitored services.

- Alert on Unknown

  The user receives an alert when up.time detects an error in the configuration of the monitor, or if up.time cannot execute the service check.

- Alert on Recovery

  The user receives an alert when the service recovers from an error – for example, an application, process or service restarts, or a server reboots.

7 **Click Save.**

Viewing Distribution Lists

You can view the details of a Distribution List to ensure is properly configured. The details of a Distribution List include an email address, and the conditions under which alerts will be sent.

To view Disbrituion Lists, do the following:

1 **Click Users on the up.time tool bar.**

2 **In the Tree panel, click View Distribution Lists.**

A list of Distribution Lists appears in the **Distribution Lists** subpanel.

3 **Click the name of the Distribution List that you want to view.**

The details of the group appear in the **Distribution Lists** subpanel.

Editing Distribution Lists

If you find that a Distribution List is not properly configured, you can edit that list.

To edit Distribution Lists, do the following:

1 **Do one of the following:**

- Click the **Edit** icon () beside the name of the Distribution List.

- Click the name of the Distribution List you want to edit, then click **Edit Distribution List** on the **Distribution List Information page**.

The **Edit Distribution List** window appears.

2 **Edit the group as described in "Adding Distribution Lists" on page 128.**

Working with Notification Groups

When up.time detects a problem with a system or service in your environment, it can issue alerts to specific users. If a group of users in your enterprise should receive certain notifications, you can ensure that they do by defining *Notification Groups* and adding those users to the group.

A Notification Group specifies the users who will receive the notifications, as well as the Alert Profile that will be used to react to the problems. See the section “Alert Profiles” on page 145 for more information.

Users can only view the Notification Groups to which they are members. While users can see the members of Notification Groups to which they belong, they can only view detailed user information for users that belong to the same user groups.

Adding Notification Groups

To add Notification Groups, do the following:

1. **Click Users on the up.time tool bar.**
2. **In the Tree panel, click Add New Notification Group.**
3. **Type a descriptive name in the Name of Notification Group field.**

   You will select this name when defining Alert Profiles. For more information on Alert Profiles, see “Alert Profiles” on page 145.
4. **Optionally, type a description of the group in the Description of Notification Group field.**
5. **Select one or more Alert Profiles to apply to the group from the Available Alert Profiles list, then click Add.**
6. **Select one or more users to add to the group from the Available Users list, then click Add.**
7. **Select one or more Distribution Lists to add to the group from the Available Distribution Lists, then click Add.**
8. **Click Save.**

Viewing Notification Groups

You can view the details of a Notification Group to ensure that the group is properly configured. The details of a Notification Group include:

- the Alert Profiles assigned to the group
- the users in the group
- whether or not the users are configured to receive alerts
- the conditions on which alerts are sent to the users

To view Notification Groups, do the following:

1. **Click Users on the** up.time **tool bar.**
2. **In the Tree panel, click View Notification Groups.**

   A list of Notification Groups appears in the **Notification Groups** subpanel.
3. **Click the name of the Notification Group that you want to view.**

   The details of the group appear in the **Notification Groups** subpanel.
4. **To view the details of an Alert Profile, click the name of the profile.**

Editing Notification Groups

If you find that a Notification Group is not properly configured, you can edit that group.

To edit Notification Groups, do the following:

1. **Do one of the following:**

   - Click the **Edit** icon () beside the Notification Group.
   - Click the name of the notification whose information you want to edit, and then click **Edit Notification Group** on the **Notification Group Information page**.

   The **Edit Notification Group** window appears.
2. **Edit the group as described in "Adding Notification Groups" on page 131.**

Changing How Users Are Authenticated

By default, user management and authentication is based entirely in up.time: a profile for a User is created in up.time, and all profile information is kept in the DataStore. up.time user lists exist, and are maintained, separately from any other user management framework your organization may be using. In light of this, you can elect to use Active Directory or an LDAP-based service for authentication and user detail synchronization.

If you configure up.time to authenticate users against a central AD or LDAP directory, password entry on login will refer to that directory instead of the DataStore. Additionally, if you choose to synchronize specific user attributes (e.g., email address), the up.time user profiles will draw all information from the central directory instead of the DataStore. Both measures ensure up.time access is automatically kept in sync with the current access levels in your organization: up.time administrators do not have to manually update user access to match staffing changes.

If user detail synchronization with Active Directory or LDAP is enabled, you will no longer be able to manually add users from within up.time: the **Add New User** option on the **Users** panel will not be available.

Regardless of which authentication and synchronization method is selected, the up.time "admin" user profile will always be stored, and authenticated against the password found in, the DataStore.

Active Directory Authentication

To use Active Directory for user management, you need to provide up.time with your organization's AD information. You can also define whether, and how much, user information is synchronized between AD and up.time's user list.

Enabling Active Directory for Authentication

To configure up.time to check an Active Directory listing for user passwords, do the following:

1 **On the up.time tool bar, click Config.**

2. **In the Tree panel, click User Authentication.**
3. **Click Edit Configuration.**
4. **Select Active Directory as the authentication method.**

   You will next need to provide access details for the Active Directory server.
5. **In the Primary Domain Controller field, enter the host name of the server acting as the domain controller, most likely enabled as the global catalog.**
6. **If applicable, in the Backup Domain Controller field, enter the name of the server acting as an additional domain controller on the same domain.**
7. **Enter the Port through which communication to the domain controller occurs.**
8. **If communication to the domain controller is secure, select the SSL check box.**
9. **In the Domain Name field, enter the domain that contains the domain controller.**
10. **Continue to the next section to enable and configure synchronization from the Active Directory listing to up.time user profiles. If you do not wish to synchronize users, click Save.**

    Clicking **Save** switches the authentication source to Active Directory. Administrators still need to create profiles for all up.time users, but will not need to set a password for each one. See “Adding Users” on page 121 for more information.

Defining Active Directory Synchronization Mapping

Before synchronizing user details, a populated “uptime” group must already exist in the Active Directory listing; you will also need to know its distinguished group name, as it will be required during configuration.

All DataStore-based user profiles will be deleted when you switch to Active Directory for synchronization—a list of affected users will be displayed during configuration. Before continuing, you should ensure your up.time users are also in the AD listing.

To configure user detail synchronization from the Active Directory list, do the following:

1. **Click Edit Configuration to open the User Authentication Configuration pop-up window.**
2. **Select the Synchronization Enabled check box.**

   All user synchronization configuration options appear.
3. **In the Synchronize Users field, enter the frequency at which up.time user information will be synchronized with the Active Directory listing.**

   By default, synchronization occurs every hour.
4. **In the AD Group Distinguished Name field, enter the name of the AD group of up.time users (e.g., CN=uptime users, CN=Groups, DC=yourdomain, DC=com).**
5. **If required, enter an appropriate administrative AD Username and AD Password required to access the directory.**
6. **In the User Name field, provide the name attribute used to retrieve the user name (e.g., sAMAcountName).**

   For AD synchronization, a user name is the minimum amount of directory information up.time needs to map to a user profile.
7. **For the remaining Field Mappings, provide attibutes for other user details you would like to synchronize with the up.time user profile:**
   - i **First Name (e.g., givenName)**
   - ii **Last Name (e.g., sn)**
   - iii **Location (e.g., physicalDeliveryOfficeName)**
   - iv **Email Address (e.g., userPrincipalName)**
   - v **Pager/Cellphone**
   - vi **User's Windows Desktop Host Name**
   - vii **User's Windows Desktop Workgroup**

Any user attributes chosen to be synchronized with the directory will not be editable in up.time.

8. **Select a User Role to which any newly detected users will be assigned.**
9. **Select a User Group to which any newly detected users will be assigned.**
10. **Click Save.**

    Once saved, up.time will synchronize its list of users with the up.time group in Active Directory at the specified interval.

LDAP Authentication

To use LDAP for user management, you need to provide up.time with your organization's LDAP information. You can also define whether, and how much, user information is synchronized between LDAP and up.time's user list.

Enabling LDAP for User Authentication

To configure up.time to check an LDAP listing for user passwords, do the following:

1. **On the up.time tool bar, click Config.**
2. **In the Tree panel, click User Authentication.**
3. **Click Edit Configuration.**
4. **Select LDAP as the authentication method.**

   You will next need to provide access details for the Active Directory server.
5. **In the LDAP URL field, enter the address for the LDAP server.**

   If directory communication occurs through secure channels, such as TLS or SSL, ensure this is reflected in the server address (e.g., "`ldaps://`" instead of "`ldap://`").
6. **Enter the LDAP Query that up.time will use on the LDAP server to look up a user's name.**
7. **Continue to the next section to enable and configure synchronization from the Active Directory listing to up.time user profiles. If you do not wish to synchronize users, click Save.**

Clicking **Save** switches the authentication source to the LDAP directory. Administrators still need to create profiles for all up.time users, but will not need to set a password for each one. See “Adding Users” on page 121 for more information.

Defining LDAP Synchronization Mapping

Before synchronizing user details, a populated “uptime” group must already exist in the LDAP directory; you will also need to know its distinguished group name, as it will be required during configuration.

Note that all DataStore-based user profiles will be deleted when you switch to an LDAP directory for synchronization—a list of affected users will be displayed during configuration. Before continuing, you should ensure your up.time users are also in the LDAP directory.

To configure user detail synchronization from the Active Directory list, do the following:

1. **Click Edit Configuration to open the User Authentication Configuration pop-up window.**
2. **Select the Synchronization Enabled check box.**

   All user synchronization configuration options appear.
3. **In the Synchronize Users field, enter the frequency at which up.time user information will be synchronized with the LDAP listing.**

   By default, synchronization occurs every hour.
4. **In the LDAP Group Distinguished Name field, enter the name of the LDAP group of up.time users (e.g., CN=uptime users, CN=Groups, DC=yourdomain, DC=com).**
5. **If required, enter an appropriate administrative LDAP Username and LDAP Password required to access the directory.**
6. **In the User Name field, provide the attribute used to retrieve the user name.**

   For LDAP synchronization, a user name is the minimum amount of directory information up.time needs to map to a user profile.

7 **For the remaining Field Mappings, provide attibutes for other user details you would like to synchronize with the** up.time **user profile:**

   i **First Name**

   ii **Last Name**

   iii **Location**

   iv **Email Address**

   v **Pager/Cellphone**

   vi **User's Windows Desktop Host Name**

   vii **User's Windows Desktop Workgroup**

Any user attributes chosen to be synchronized with the directory will not be editable in up.time.

8 **Select a User Role to which any newly detected users will be assigned.**

9 **Select a User Group to which any newly detected users will be assigned.**

10 **Click Save.**

Once saved, up.time will synchronize its list of users with the up.time group in the LDAP listing at the specified interval.

up.time DataStore Authentication

By default, up.time uses its own database for password storage and look-up.

If you are switching *back* to using the DataStore from a central AD or LDAP directory, all up.time users created while either was used as the authentication method will no longer have passwords. You will need to modify all existing user accounts to include passwords.

Enabling the DataStore for User Authentication

To use up.time DataStore to store passwords for user authentication, do the following:

1. **On the up.time tool bar, click Config.**
2. **In the Tree panel, click User Authentication.**
3. **Click Edit Configuration.**
4. **Select Database as the authentication method.**
5. **Click Save.**

up.time

CHAPTER 10

Alerts and Actions

This chapter covers up.time's alerting features, the monitoring periods when alerts can happen, as well as the configuration of post-alert actions:

Understanding Alerts

When a problem occurs at a Datacenter, Application, or SLA, the Enterprise Monitoring Station can send *alerts* to users. Alerts are notifications that inform users who are configured to receive alerts of the problem. The notification message contains the following information:

- the type of notification – either `Problem` or `Recovery`
- the date and time when the problem occurred
- the name of the host on which the problem occurred
- the status of the host (see "Understanding the Status of Services" on page 14 for more information)
- the name of the service that is experiencing the problem
- the current state of the service
- any output from the monitor

Whenever the status of an Element changes – for example from Critical to Warning – up.time sends an alert.

You can also configure *alert escalations* that occur if a warning is sent and is not acted upon. For example, if an alert is sent to a system administrator and the administrator does not attend to the problem within a specified amount of time, then the alert will be sent to the administrator's manager.

up.time can send alerts via:

- email messages to a cell phone or a pager, or to one or more email addresses
- a Windows popup

The following is a sample email alert:

```
Notification type: Problem
1/12/2008 10:52
Host: UK Lab2 (OK)
Service: Datacenter Health Check
Service State: CRIT
Output: Can't connect to remote database link.
```

The following is a sample pager alert:

```
subject:
    CRIT Alert
content:
    5/7/2005 13:22
    Type: Problem
    Service: FTP (CRIT)
    Host: filter (CRIT)
```

For more information on alerts, see “Monitor Alert Settings” on page 148.

Understanding the Alert Flow

Alerts in up.time follow a specific flow. When up.time detects a problem with a host, it issues an alert. up.time then continues to check the host at specific intervals and reports on the status of the host.

Considering the following example:

- up.time checks the host system every 15 minutes
- alerts are sent continually every check interval until up.time detects a change in the state of the host system
- whenever an error is encountered, up.time rechecks the system every minute
- if all rechecks up to the maximum number of rechecks fails, up.time issues an alert

up.time encounters a critical error on a host. up.time performs three rechecks at one minute intervals – all of which return a critical error – and then sends an alert after the third recheck.

up.time then checks the host every two hours. While up.time encounters two critical errors, it does not send an alert. Then, the status of the host changes from critical to warning. When this change is detected, up.time sends an alert informing recipients of the change in status. When the status of the host changes to OK, up.time issues an alert informing recipients that the host has recovered.

This alert flow is illustrated in the following diagram:

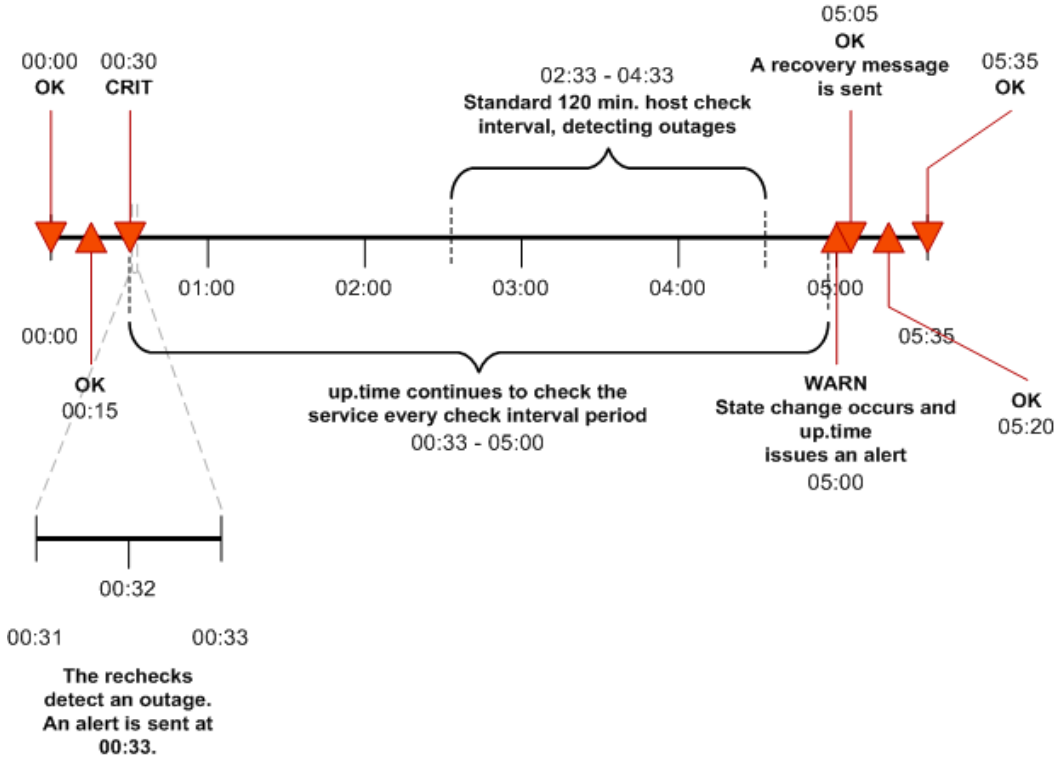

Alert Profiles

Alert Profiles are templates that tell up.time how to react to various alerts that are generated by service checks. Alert Profiles enable up.time to execute a series of actions in response to the failure of a service check or when a threshold is exceeded. The following diagram illustrates how an Alert Profile works:

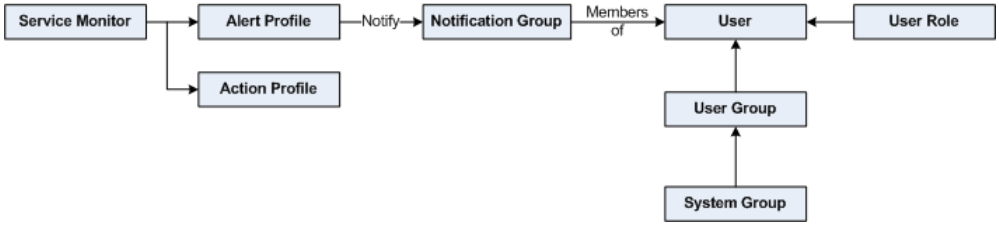

An Alert Profile can send an alert via email, or to a pager or a cell phone, or a Windows popup alert. You can configure any or all of these actions to occur simultaneously. For example, if a Web server process stops responding, the system administrator can be notified.

Enabling the Windows Messaging Service

In order to receive popup alerts from up.time, the Windows messaging service must be enabled on the recipient's computer.

To enable the Windows messaging service, do the following:

1. **In Windows, select Start > Control Panel.**
2. **In the Control Panel, double click Administrative Tools, and then double click Services.**

   The **Services** window appears.
3. **Find and then double click Messenger in the list of services.**

   The **Messenger Properties** dialog box appears.
4. **In the Messenger Properties dialog box, select Automatic from the Startup type dropdown list.**
5. **Click Apply.**

Creating Alert Profiles

To create Alert Profiles, do the following:

1. **On the up.time tool bar, click Users.**
2. **In the Tree panel, click Add Alert Profile.**

   The **Add Alert Profile** window appears.
3. **Type a descriptive name for the profile in the Name of Alert Profile field.**
4. **In the Start alerting on notification number field, enter the number of times an error must occur before up.time sends an alert notification.**
5. **Enter the number of times to re-send the notification in the End alerting on notification number field.**

   Optionally, click the **Never Stop Notifying** option to have up.time continually send notifications.
6. **Select one of the following notification options:**

   - Email Alert

     Sends the alert to the email addresses of the members of a Notification Group.

   - Pager Alert

     Sends the alert to the pagers of the members of a Notification Group.

   - Script Alert

     Uses a script to send the alert via SMS to the mobile phones of the members of a Notification Group.

     Since this alert option relies on a script or batch file, you must enter its name and path in the **Script Path** field (for example, `/usr/local/uptime/scripts/scriptAlert.sh`).

     When the alert is triggered, up.time runs the script and passes the script or batch file a set of parameters. The script is run for each up.time user who will receive the SMS message.

For details on how to create the script, see the Client Care Web site Knowledge Base article “Creating Custom Alert Scripts in up.time Alert Profiles”.

- Windows Popup Alert

  Sends the alert via the Windows messaging service to the desktops of the members of a Notification Group.

7 **Select one or more groups that will receive the notifications from the Available Notification Groups list, and then click Add.**

8 **Click Save.**

Viewing Alert Profiles

To view Alert Profiles, do the following:

1 **On the up.time tool bar, click Users.**

2 **In the Tree panel, click View Alert Profiles.**

The **Alert Profiles** subpanel appears. The subpanel displays the settings that you configured when you created the profile, as well as a list of the services that are attached to the profile.

3 **To test whether or not the profile will send alerts, click the Test Alert Profile button.**

A popup window appears, and the alert is sent using the notification method – email, pager, script, or Windows popup – that is specified in the profile. The following is an example of an email alert:

```
Notification type: Problem 27/4/2006 09:19
Host: Test Host (OK)
Service: Test Monitor
Service State: OK
Output: This is a test notification; please ignore.
```

When the alert is sent, the message `Alert Profile Tested` appears in the popup window. If an error message appears in the popup window, edit the profile and test it again.

Editing Alert Profiles

To edit Alert Profiles, do the following:

1. **On the up.time tool bar, click Users.**
2. **In the Tree panel, click View Alert Profiles.**
3. **Click the Edit Alert Profile icon ( ) beside the name of the profile that you want to edit.**

   The **Edit Alert Profile** window appears.
4. **Edit the Alert Profile fields, as described in the section "Creating Alert Profiles" on page 146.**

Associating Alert Profiles to Elements

You can associate an Alert Profile to any Application or SLA if their state changes from OK to Warning or Critical. Alert Profiles are normally associated with any of these monitored items at the time of their configuration.

See "Working with Applications" on page 62 and "Adding and Editing SLA Definitions" on page 111 for more information about configuring Applications and SLAs, respectively.

Working with Custom Alert Formats

up.time's standard alert format is well suited for most alerting needs. However, you can modify the content of the alert. up.time comes with three custom alert templates. You can change the content of the alert by adding or removing variables from the template.

To define a custom alert format, do the following:

1. **Define an Alert Profile, as described on page 146.**
2. **In the Custom Format Options section, click Custom Formats.**
3. **From the dropdown list, select one of the following options:**

   - Small Template

     Contains the date and time of the alert, as well as the names and status of the service and host for which the alert was generated. This corresponds to the template used for pager alerts.

   - Medium Template

     Contains the information in the small template, as well as an expanded subject line, the type of notification, and output from the service monitor. This corresponds to the template used for email alerts.

   - Long Template

     Contains the information in the medium template, as well as the status of the host.

4. **Click Fill.**

The variables associated with the template appear in the subject and body fields.

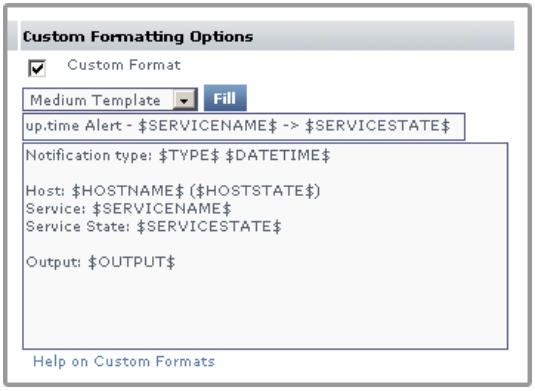

5 **Add or remove variables (see) as needed. You can also add other information to the body of the alert, such as paths to custom scripts or the names of alternative contacts.**

6 **Click Save.**

Custom Alert Format Variables

The variables are the building blocks of a custom alert format. You can add or remove variables to suit your needs.

These alert variables are also available as input parameter values when configuring an Action Profile to initiate a VMware vCenter Orchestrator workflow.

The table below explains the variables available in custom alerts, as well as Orchestrator input parameters :

Variable	Definition
`$DISPLAYNAME$`	The name of the Element as it appears in the up.time Web interface. A system can have a different display name than the hostname. For example, you can assign the display name `Toronto Mail Server` to a system with the host name `10.1.1.6`.

Variable	Definition
$DATETIME$	The date and time at which the alert was generated. This appears in the subject line of the message.
$SERVICENAME$	The name of the service, along with the name of the host for which the alert was generated. For example, if the alert was generated by the ping check for the server MailHub, then PING-MailHub appears in the alert. This appears in the subject line of the message.
$SERVICESTATE$	One of the following: •OK •WARN •CRIT •MAINT •UNKNOWN This appears in the subject line of the message.
$DATE$	The date on which the alert was generated.
$TIME$	The time at which the alert was generated.
$HOSTNAME$	The name of the host (as saved in up.time) for which this alert was generated.
$HOSTSTATE$	The status of the host, which can be one of the following: •OK •WARN •CRIT •MAINT •UNKNOWN

Variable	Definition
$TYPE$	The type of notification, which can be one of the following: •Problem •Recovery
$OUTPUT$	The output of the monitor that generated the alert. For example, Ping completed: 1 sent, 100.0% loss, 0.0ms average round trip time

Action Profiles

Action Profiles are templates that direct up.time when it encounters a problem on a monitored system. You can associate an Action Profile to any Application or SLA if their state changes from OK to Warning or Critical. Action Profiles are normally associated with any of these monitored Elements at the time of their configuration.

See "Working with Applications" on page 62 and "Adding and Editing SLA Definitions" on page 111 for more information about configuring Applications and SLAs, respectively.

Actions include one of the following tasks:

- write an entry to a log file
- run a recovery script that can reboot a non-responsive server; or restart an application, process, or service
- stop, start, or restart a Windows server
- initiate a VMware vCenter Orchestrator workflow
- send an SNMP trap to a specific traphost and trap community

As templates, Action Profiles can be reused for any number of Service Monitor configurations. This means you can create a series of them as standard actions used to respond to typical types of problems you may encounter, depending on what role a Service Monitor is playing (e.g., availability or performance).

VMware vCenter Orchestrator Workflow Actions

If an administrator has integrated up.time with VMware vCenter Orchestrator (see "VMware vCenter Orchestrator Integration" on page 297, you can configure Action Profiles to initiate Orchestrator workflows.

Orchestrator is a VMware vCenter Server add-on that allows its administrators to create workflows that automate vCenter management tasks. These Orchestrator workflows are open ended: all vCenter actions are available for automation through the processing of parameters and runtime arguments. up.time Action Profiles can be configured to provide input parameters to specific workflows, thus integrating vCenter management with up.time's monitoring and alerting capabilities.

For example, if up.time is monitoring memory, CPU, and hard disk use for a virtualized server, the passing of performance thresholds can trigger an Action Profile that, in turn, triggers an Orchestrator workflow that creates a new virtual machine to alleviate resource strain. In a converse example, if up.time is monitoring a virtualized server for long periods of inactivity, a triggered Action Profile can initiate an Orchestrator workflow that shuts down the instance to free up resources.

By tightly integrating up.time's monitoring and alerting with VMware vCenter Orchestrator's automated virtual environment administration, you can accelerate your organization's reaction time with virtual systems management, and map established policies to automated actions.

When configuring Action Profiles, up.time communicates with Orchestrator and dynamically produces a list of all available workflows. (This includes any third-party workflow packages that have been installed on the Orchestrator server, including the up.time Orchestrator package.)

When a workflow is selected, and the **Get Parameters** button is clicked, the corresponding input parameter fields are dynamically displayed, allowing you to specify parameter values required to completely configure the workflow for execution should an up.time alert initiate it.

Orchestrator Input Parameter Variables

When configuring a VMware vCenter Orchestrator workflow, you have at your disposal a set of up.time-specific variables that can be entered as parameter variables, and whose ensuing runtime values will be passed to the Orchestrator workflow during execution. The variables available to you are those that are used when creating a custom alert format. See "Custom Alert Format Variables" on page 150 for information.

SNMP Trap Actions

You can also configure an Action Profile to send an SNMP trap to a particular host. An SNMP trap is notification that is issued by a system that is running SNMP when a problem occurs. The host to which the SNMP trap is sent must be running an SNMP trap listener.

If you use SNMP traps, the trap message will be sent in the format specified by the up.time MIB. This MIB is found in the `scripts` directory. The uptime software enterprise OID is `.1.3.6.1.4.1.24216`.

Creating Action Profiles

To create Action Profiles, do the following:

1. **On the up.time tool bar, click Users.**
2. **In the Tree panel, click Add Action Profile.**

   The **Add Action Profile** window appears.
3. **Enter a name for this profile in the Name of Action Profile field.**
4. **Specify the number of times an error must occur before up.time sends a notification in the Start action on notification number field.**
5. **Specify the number of times action will be carried out in the End action on notification number field.**

   Optionally, select the **Never Stop Notifying** option to continually carry out the action in this profile until the problem is resolved.
6. **If VMware vCenter Orchestrator integration has been enabled, and you would like the Action Profile to drive an Orchestrator workflow, do the following:**

   i. **In the Select Workflow field, input a workflow to configure.**

      You can either scroll through and select the workflow from the drop-down list, or begin typing the workflow's name.

   ii. **Click Get Parameters.**

      up.time will retrieve information from the Orchestrator server and dynamically display configuration fields for the chosen workflow's input parameters.

   iii. **Configure the input parameter fields for the workflow.**

      For information on the specific configuration parameters available for the chosen workflow, consult the appropriate developer's documentation.
4. **If you would like the Action Profile to write to a log, in the Log File field, enter the name and path to a log file on the Enterprise Monitoring Station to which error information will be written.**
5. **If you would like the Action Profile to run a recovery script, in the Recovery Script field, enter the name and path to a script that**

will reboot a server, or restart an application, process, or service.

The recovery script will also have the following information appended to it:

- the date and time on which the error occurred
- the type of error notification that was sent
- the name of the host on which the error occurred
- the state of the host
- the name of the service that threw the error
- the state of the service
- the output that was generated by the error

For example:

```
"/usr/local/uptime/recover.sh" "24/12/2007 5:01:05"
"Problem" "printserver" "null" "WinSrv-Print Spooler"
"CRIT/threshold error" "servicestatus: Not Running does
not match Running (Service 'Print Spooler' found, status:
Not Running, took 12ms)"
```

You can also use the recovery script to file trouble tickets with a system like Remedy, or to interact with third party software packages.

6 **If you are setting up an Action Profile for a Windows server on which a Datacenter and up.time agent are running, you can also leave the Windows Service as Agent, and complete the following fields:**

- Windows Host

  The name of the host on which the service is running.

- Agent Port

  The port on which the up.time agent that is installed on the system is listening. The default is 9998.

- Use SSL

Select this option if up.time will securely communicate with the host using SSL (Secure Sockets Layer).

- Agent Password

  Enter the password that is required to access the agent that is running on the system that is being monitored.

- Windows Service

  The name of the specific Windows service to which the Action Profile will apply.

- Action

  Select one of the following actions:

  - None
  - Start
  - Stop
  - Restart

7 **If you are setting up an Action Profile for a Windows server that is using a WMI implementation, and on which a Datacenter is running, you can also select the Windows Service as WMI, and complete the following fields:**

- WMI Host:

  The name of the host on which the service is running.

- Windows Domain:

  The Windows domain in which WMI has been implemented.

- Username:

  The name of the account with access to WMI on the Windows domain.

- Password:

  The password for the account with access to WMI on the windows domain.

- Windows Service

The name of the specific Windows service to which the Action Profile will apply.

- Action

  Select one of the following actions:

  - None
  - Start
  - Stop
  - Restart

8 **If you want to send SNMP traps to a particular host, complete the following fields:**

- SNMP Trap Host

  The name of the host that monitors SNMP traps.

- SNMP Trap Port

  The port number on the trap host to which the SNMP trap is sent.

- SNMP Trap Community

  The name which acts as a password for sending trap notifications to the trap host.

- SNMP Trap OID (optional)

  The object identifier (OID) that identifies the SNMP trap – for example, `.1.3.6.1.2.1.34.4.1.7`.

9 **If Splunk integration has been enabled, and you would like the Action Profile to write to the Splunk log, complete the following fields:**

- Splunk Hostname

  The host name of the server on which Splunk is running.

- Logging Port

  The port on which the Splunk server is listening for logging requests. This port is configured in Splunk, and you will need to contact the Splunk administrator for this information.

Click the **Use SSL** option to securely access the Splunk server using SSL.

For more information on Splunk integration, see “Splunk Settings” on page 300.

10 **Click Save.**

Viewing Action Profiles

To view Action Profiles, do the following:

1 **On the up.time tool bar, click Users.**

2 **In the Tree panel, click View Action Profiles.**

The **Action Profiles** subpanel appears, displaying the settings that you configured when you created the profile, as well as a list of the services that are attached to the profile.

3 **To test whether or not the profile works, click the Test Action Profile button.**

A popup window appears, and the Enterprise Monitoring Station tries to carry out the action defined in the profile. When the action is completed, the message `Action Profile tested` appears in the popup window.

If an error message appears in the popup window, edit the profile and test it again.

Editing Action Profiles

To edit Action Profiles, do the following:

1 **On the up.time tool bar, click Users.**

2 **In the Tree panel, click View Action Profiles.**

3 **Click the Edit Action Profile icon ( ) beside the name of the profile that you want to edit.**

The **Edit Action Profile** window appears.

4 **Edit the Action Profile fields as described in the section “Creating Action Profiles” on page 155.**

Monitoring Periods

Monitoring Periods are the times when up.time sends alerts

up.time comes with the following Monitoring Periods:

- 24x7

  Monitoring is performed 24 hours a day, seven days a week.

- 9am to 5pm weekdays

  Monitoring is performed from 9 a.m. to 5 p.m., Monday to Friday.

- Never

  No monitoring is carried out.

You can add Monitoring Periods that suit your needs. For example, you can create a Monitoring Period called `Weekends` that only monitors a host from 12:00 a.m. on Saturday to 11:59 p.m. on Sunday.

Adding Monitoring Periods

To add Monitoring Periods, do the following:

1. **On the up.time tool bar, click Users.**
2. **In the Tree panel, click Add Monitoring Period.**

   The **Add Monitoring Periods** window appears.
3. **Type a name in the Monitoring Period Name field.**
4. **In the Definition section, enter one or more time period expressions that combine to create a full Monitoring Period definition.**

   See “Time Period Definitions” on page 321 for information on the types of time period expressions that are valid in up.time.
5. **Click Save.**

up.time

CHAPTER 11

Understanding Report Options

This chapter is an overview of the options available for generating reports in up.time, and contains the following sections:

Overview

up.time can generate reports on the status of the servers in your environment, based on criteria that you specify. A report uses data that up.time has collected from a system, over a period of time that you specify. You can configure reports to run between certain hours of the day.

Reports are useful when you need to pinpoint the source of a problem within you environment. With a report, you can visually analyze how individual critical resources – such as memory, CPU, and disk resources – are being consumed. You can dynamically generate and view reports, schedule and email reports to other up.time users.

This chapter looks at the options that you can set to generate, save, and schedule reports. For more information about the individual reports and how to configure them, see “Using Reports” on page 175.

Generating Reports

You can generate reports either dynamically or in the background. Dynamic reports are reports that up.time displays in a new Web browser window. Dynamic reports appear within several seconds or several minutes, depending on the type of report that you are generating and on the information that the report collects.

Background reports are reports that you schedule to be run at specific intervals using the up.time report queue. When it is time for a scheduled report to run, up.time puts the report into the report queue and determines that status of the report based on the following states:

- Pending

  The report is in the queue and is waiting to run.

- Running

  The report is being generated.

- Completed

  The report has been generated, and has been sent (via email) to the users configured to receive that report.

For information on how to schedule reports, see “Scheduling Reports” on page 169.

If you do not receive a scheduled report, check the Report Log (see “The Report Log” on page 172) or contact your system administrator.

Report Generation Options

up.time can generate reports in four ways:

- Print to Screen

  Displays the report in a new window. This is the default option.

- PDF to Screen

  Converts the report to a PDF document, and displays it in a new window. You can save the PDF document to a local or network drive, or print it.

- XML to Screen

  Displays the report, as an unformatted XML document, in a new window.

- Email Address

  Enables you to email the report, as a PDF document attached to an email message, to:

  - A specific up.time user, for example a system administrator.

    Click **User** and then select the name of an up.time user to whom you want to send the report from the dropdown list.

  - The members of one or more up.time user groups.

    Click **Group** and then select the name of an up.time user group to which you want to send the report from the dropdown list.

  - One or more email addresses.

    Click the **Email Address** option, and then type the email address of the person to whom you want to send the report in the field. To send the report to multiple recipients, type their email addresses in the field separated by commas or semi-colons. For example:

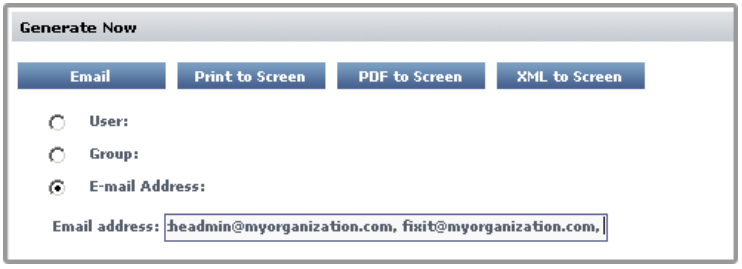

Reports that are sent by email have a file name that consists of the type of report and the date and time range it covers. For example, a CPU Utilization Ratio report might be named:

```
ReportCPUUtilizationRatio_2006-01-10_00-00-2006-01-10_14-53.pdf
```

If you choose to output the report to the screen, a message appears while the report is being generated. When the report has been generated, it is displayed in the report window. If up.time cannot connect to a host, the following error message appears in the report window:

```
An error occurred while running this report. Verify the configuration of up.time and try again.
```

Saving Reports

If you find that you need to generate reports on a regular or frequent basis, you can save the parameters for the report to the DataStore. A link to the report appears in the **My Portal** panel. Click the link to generate the report.

You can also schedule reports to be generated and sent by email at particular intervals. See "Scheduling Reports" on page 169 for more information.

To save reports, do the following:

1. **In the Save Report area of the Report subpanel, select one of the following options:**
   - HTML
   - PDF
   - XML
   - Email
2. **If you selected Email in step 1, specify one of the email options.**
3. **Type a name for the report in the Save to My Portal As field.**
4. **Optionally, type a description for the report in the Report Description field.**
5. **Click Save Report.**

Saving Reports to the File System

You can save reports to the file system of a server in your environment so others in your organization can view the reports. You can, for example, save a report to a Web server for viewing on your Intranet. The reports are saved as either PDF or HTML files. The system administrator can specify the

directory on the server in which reports will be saved by adding the following entry to the file `uptime.conf`:

```
publishedReportRoot=<directory_name>
```

Where `<directory_name>` the directory into which up.time will write reports – for example, `C:/Program Files/uptime software/uptime/`. The report files are saved to a subdirectory named `GUI/published`. You need permissions to write to the `published` directory.

up.time automatically names each report file. The file name contains the following information:

- name of the report, taken from the **My Portal** panel
- date on which the report was run
- user name of the person who ran the report

The following is an example of a report file name:

```
Service Outages_2006-01-24_rfripp.pdf
```

To save reports to a file system, do the following:

1. **In the Save Report area of the Report subpanel, enter a name for the report in the Save to My Portal As field.**
2. **Optionally, enter a description of the report in the Description field.**
3. **Select either HTML or PDF from the list of options.**
4. **Click the Publish Report option.**
5. **Click the Scheduled Report option, and then select a a date and time for the report to run.**

   For more information on scheduling reports, see “Scheduling Reports” on page 169.
6. **Click Save Report.**

Viewing Saved Reports

You can quickly view any reports that were generated on the Enterprise Monitoring Station and saved to the file system. To do so, do the following:

1. **On the tool bar, click Reports.**

2 **Click Published Reports in the Tree panel.**

The **Report Library** window appears. The **Report Library** window lists the reports that were generated on the Enterprise Monitoring Station in descending order by date.

Using the Search Function

The **Report Library** window includes a search function that enables you to find specific reports.

To use the search function, do the following:

1 **In the Published Reports window, click the Search button.**

The **Search Options** appear in the window.

2 **Select one of the following options from the Search Column dropdown list:**

- Year
- Month
- Name
- Date
- User

3 **Specify the criteria for the search, and then click the Search button to view the results on the Report Library page.**

Scheduling Reports

If you need to run a report at a particular interval – for example, daily or weekly – you can schedule when the report should be generated. up.time generates the report and emails it to a user or group of users.

For example, you generate a File System Capacity Growth Report – which charts the amount of disk usage for a system. However, the system for which you are generating the report schedules backups from midnight to 4:00 a.m. Due to the gap caused by the backup, the CPU usage and disk activity statistics are not indicative of the overall system load. You can specify that the report does not cover the periods of time over which the backups occur.

To schedule reports, do the following:

1. **In the Reports subpanel, select the Email option in the Save Report section of the subpanel, and then select one of the following options:**
   - User
   - Group
   - E-mail Address
2. **Type a name for the report in the Save to My Portal As field.**
3. **Optionally, type a description for the report in the Report Description field.**
4. **Click the Scheduled Reports checkbox, and then select the time at which to run the report from the dropdown lists.**

   For example, to run the report at 3:30 p.m., select `15` from the first dropdown list and `30` from the second dropdown list, as shown below:

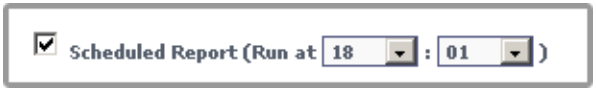

5 **Select one of the following options:**

- Daily

Do one of the following:

- Click the **Every** option, and select the number of days from the dropdown list.
- Click the **Every Weekday** option.

- Weekly

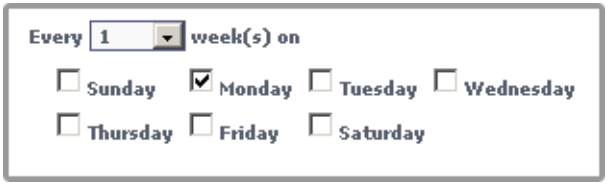

Do the following:

- Select a number of weeks from the **Every week(s) on** dropdown list. If, for example, you select 2 from the list, the report will be run every two weeks.
- Select one or more days of the week on which the report will be run.

- Monthly

Do one of the following:

- Select the **Day** option. From the first dropdown list, select the day (from 1 to 31) on which to run the report. Then, select the month (from 1 to 12) during which to run the report.

For example, if you select 3 and 7 from the dropdown lists, the report will be run on the third day of every seventh month.

- Select the second option, then do the following:
  - select first, second, third, fourth, or last from the first dropdown list
  - select a day of the week on which the report will run from the second dropdown list
  - select a number from 1 to 12 from the third dropdown list

  For example, if you select second, Tuesday, and 9 from the dropdown lists, the report will be run on the second Tuesday of every ninth month.

If you are saving an existing report after editing it or saving a new report with the name of an existing one, up.time displays a warning dialog box. Click **OK** on the dialog box to overwrite the report. Or, click **Cancel** on dialog box to give the report a different name.

The Report Log

The Report Log tracks the progress and status of scheduled reports, or reports that are running in the background. Using the Report Log, you can quickly determine whether or not reports have been successfully generated. If they have not, then you can use the log to determine why report generation failed.

The **Report Log** subpanel tracks the status of reports in the following sections:

- Pending Reports

  Reports that are in the report queue, and are waiting to run. This section contains the following information:

  - the name of the report
  - the description of the report, if available
  - whether or not the report is scheduled
  - the date and time on which the report will be run

  The following image illustrates the **Pending Reports** section:

- Running Reports

  Reports that are being run. This section contains the same information as the **Pending Reports** section, as illustrated below:

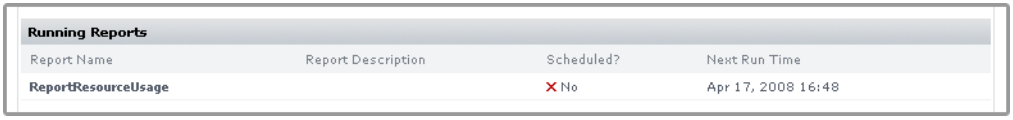

If the running report is not a scheduled report, `Emailing report in PDF format` appears in the **Report Name** column.

- Completed Reports

  Reports that have finished running, whether they were successfully generated or not. This section contains the following information:

  - the name of the report
  - the date and time on which the report run was started
  - the date and time on which the report run ended
  - the status of the report – for example, `finished`
  - a status message – for example, `Email sent` or `Address list is empty`

The following image illustrates the **Completed Reports** section:

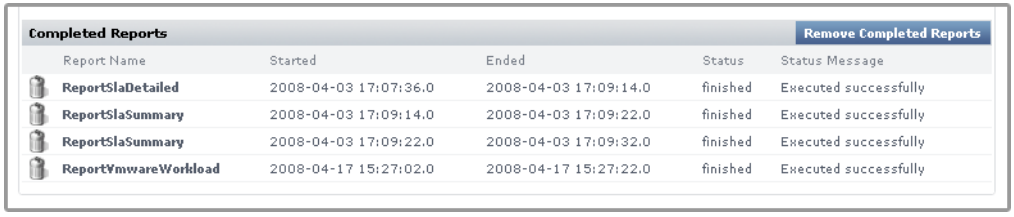

Viewing Report Logs

To view report logs, do the following:

1. **On the up.time tool bar, click Reports.**
2. **In the Tree panel, click Report Log.**

   The report log appears in the **Reports** subpanel.

   If there are no reports in the queue, up.time displays a message similar to the following ones in the **Pending Reports** and **Running Reports** sections of the **Report Logs** subpanel:

   ```
   No reports are pending
   ```

   ```
   No reports are running
   ```

Deleting Report Log Entries

Completed reports are stored in a table in the up.time DataStore. To free space in the DataStore, or to remove report log entries that you no longer need, you can delete entries in the report log from the **Report Log** subpanel.

To delete entries in the Report Log, do one of the following:

- Click the **Delete** icon () beside the entry that you want to delete.
- If you want to delete all entries in the Report Log, click the **Remove Completed Reports** button.

When prompted to confirm whether or not you want to delete the report log entry, click **OK**.

up.time

CHAPTER 12

Using Reports

This chapter describes the reporting features of up.time in the following sections:

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
- Multi-System CPU Report
- File System Capacity Growth Report
- CPU Utilization Ratio Report
- Wait I/O Report
- Service Monitor Metrics Report

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” on page 16.

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.

This graph does not appear when you generate a 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 in a hubbed environment 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.

This graph does not appear when you generate a 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.

- 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

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.

This graph does not appear when you generate a 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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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” on page 16.
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, as shown below:**

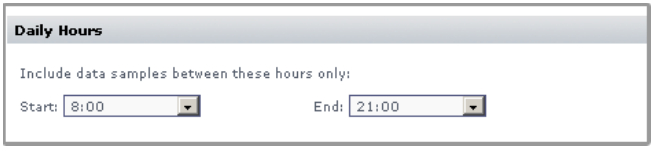

   For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select `1:00` from the **Start** dropdown list, and `13:00` from the **End** dropdown list.
4. **If you want to generate reports for systems in specific groups, select the groups from the List of Groups area.**
5. **To generate reports for one or more views, select the groups from the List of Views area.**

   See “Working with Views” on page 69 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" on page 164 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" on page 166 and "Scheduling Reports" on page 169 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" on page 16.

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.

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.

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” on page 69 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" on page 164 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" on page 166 and "Scheduling Reports" on page 169 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" on page 16.

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, as shown below:**

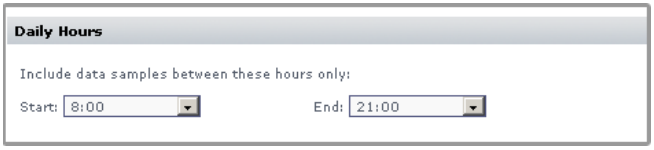

For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select `1:00` from the **Start** dropdown list, and `13:00` from the **End** dropdown list.

4 **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.

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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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:

1. **In the Reports Tree panel, click Wait I/O.**
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” on page 16.
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, as shown below:**

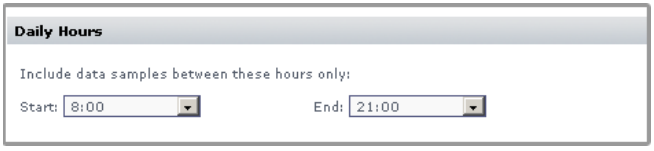

   For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select `1:00` from the **Start** dropdown list, and `13:00` from the **End** dropdown list.
4. **Optionally, enter a value in the Highlight average WIO over threshold field.**

   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.
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” on page 69 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” on page 164 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” on page 169 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” on page 16.
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” on page 69 for more information about views.
5. **If you are generating reports for specific systems in your environment, select them from the List of Entities.**
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 as shown below:**

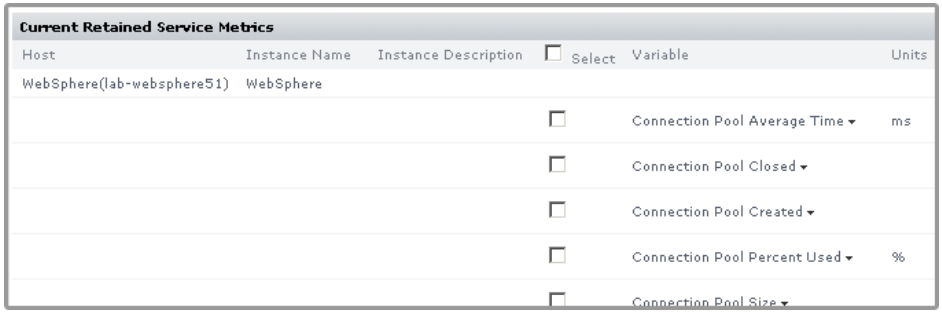

8. **Optionally, select one of the following:**

   - 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” on page 166 and “Scheduling Reports” on page 169 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
- File System Capacity Growth Report
- Server Virtualization Report
- Solaris Mutex Exception Report
- Network Bandwidth Report
- Disk I/O Bandwidth Report
- CPU Run Queue Threshold Report
- File System Service Time Summary Report

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” on page 16.
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, as shown below:**

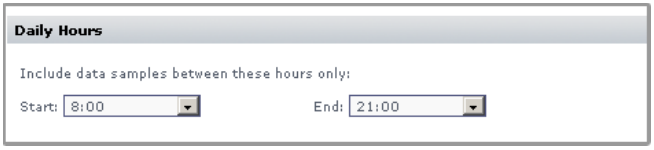

   For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select `1:00` from the **Start** dropdown list, and `13:00` from the **End** dropdown list.
4. **Select one of the following options from the Sort by dropdown list to sort the results that up.time returns:**

   - 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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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

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” on page 16.

   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 filed. 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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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” on page 16.
3. **Click the Display entity 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.
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 POWER5 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” on page 69 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” on page 164 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” on page 169 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” on page 16.

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, as shown below:**

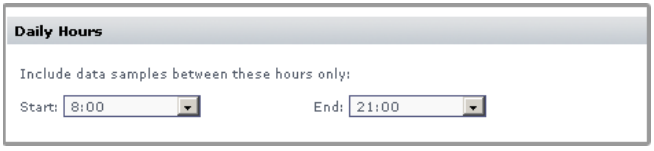

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select `8:00` from the **Start** dropdown list, and `18:00` from the **End** dropdown list.

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” on page 69 for more information about views.

7 **If you are generating reports for specific Applications in your environment, select them from the List of Entities.**

Only Solaris systems with two or more CPUs are show in the List of Entities.

8 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 for more information.

Using the Solaris Mutex Exception Report

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

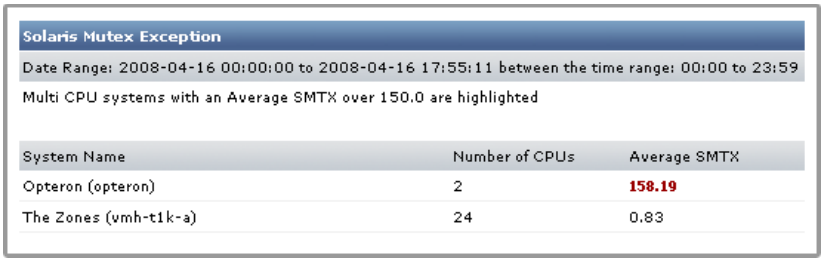

Solaris Mutex Exception

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:

- Multi-CPU Usage (see page 257 for more information)
- Run Queue Length (see page 255 for more information)
- Run Queue Occupancy (see page 255 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” on page 16.

   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, as shown below:**

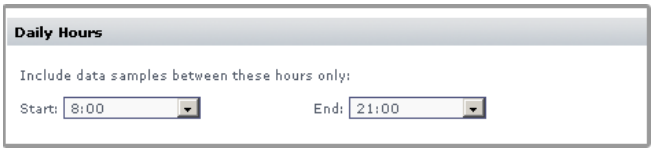

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select `8:00` from the **Start** dropdown list, and `18:00` from the **End** dropdown list.

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” on page 69 for more information about views.

7 **If you are generating reports for specific Applications in your environment, select them from the List of Entities.**

8 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 for more information.

Using the Network Bandwidth Report

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 page 273 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 the 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:

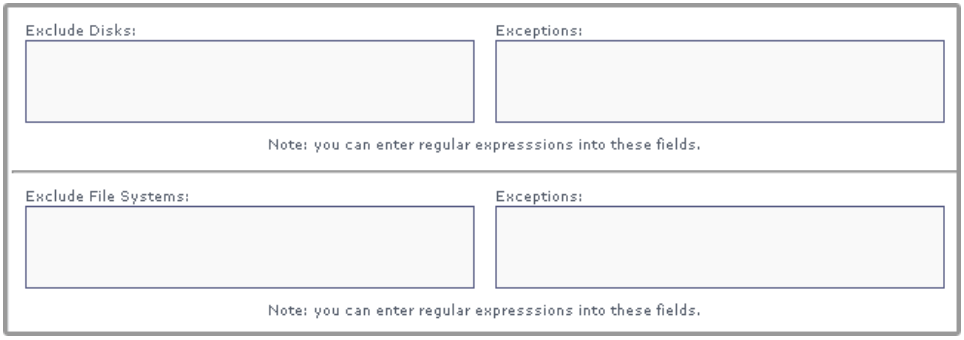

Using regular expressions, you can focus on particular disks or file systems on a server and also decrease the length of your report.

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” on page 16.

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, as shown below:**

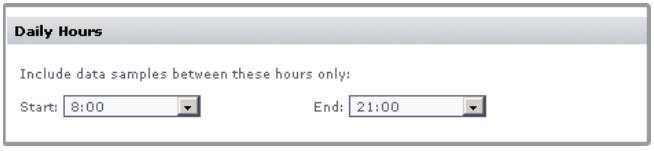

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select `8:00` from the **Start** dropdown list, and `18:00` from the **End** dropdown list.

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” on page 204 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” on page 204 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” on page 69 for more information about views.

9 **If you are generating reports for specific Applications in your environment, select them from the List of Entities.**

10 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 for more information.

Using the Disk I/O Bandwidth Report

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

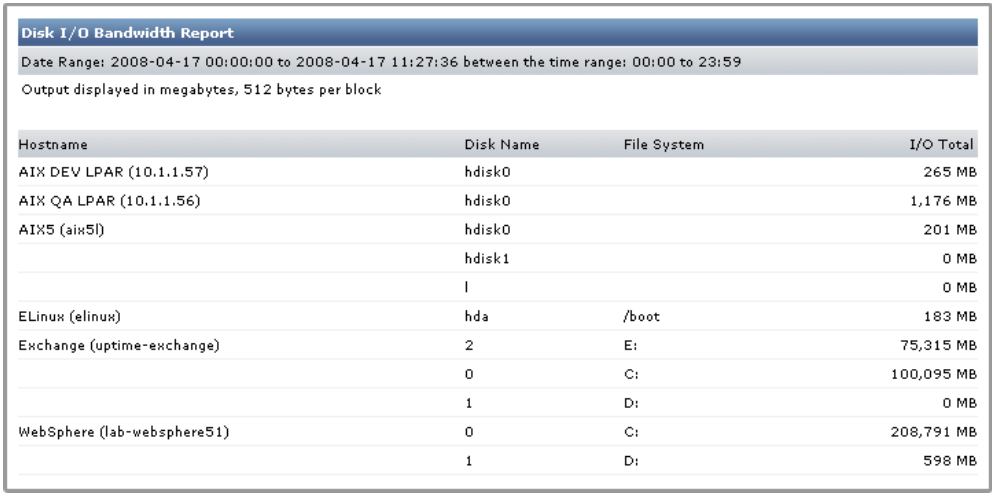

Disk I/O Bandwidth Report

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 (aix5l)	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 page 276 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" on page 16.

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, as shown below:**

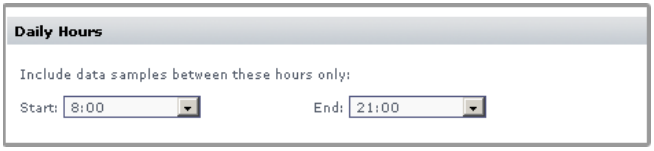

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select `8:00` from the **Start** dropdown list, and `18:00` from the **End** dropdown list.

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.

7 **If you want to include a list of processes that are in the run queue in the report, click Show Processes.**

8 **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.

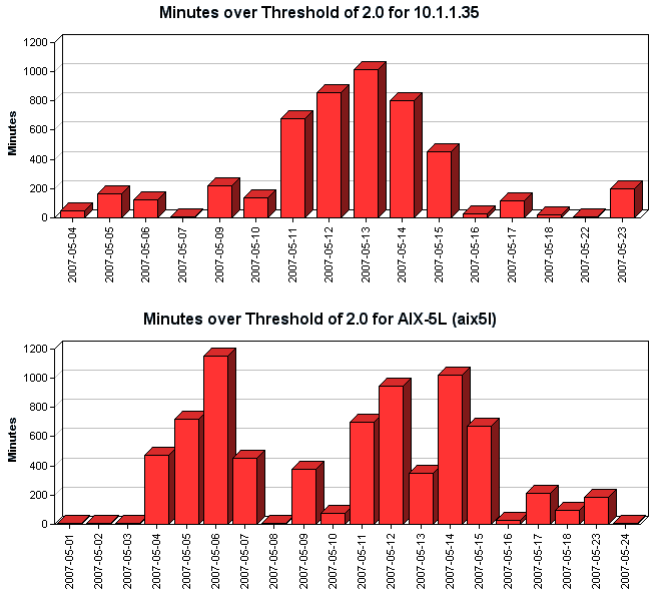

9 **To generate reports for groups of systems, select the groups from the List of Groups area.**

10 **To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 69 for more information about views.

11 **If you are generating reports for specific Applications in your environment, select them from the List of Entities.**

12 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 for more information.

Using the CPU Run Queue Threshold Report

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

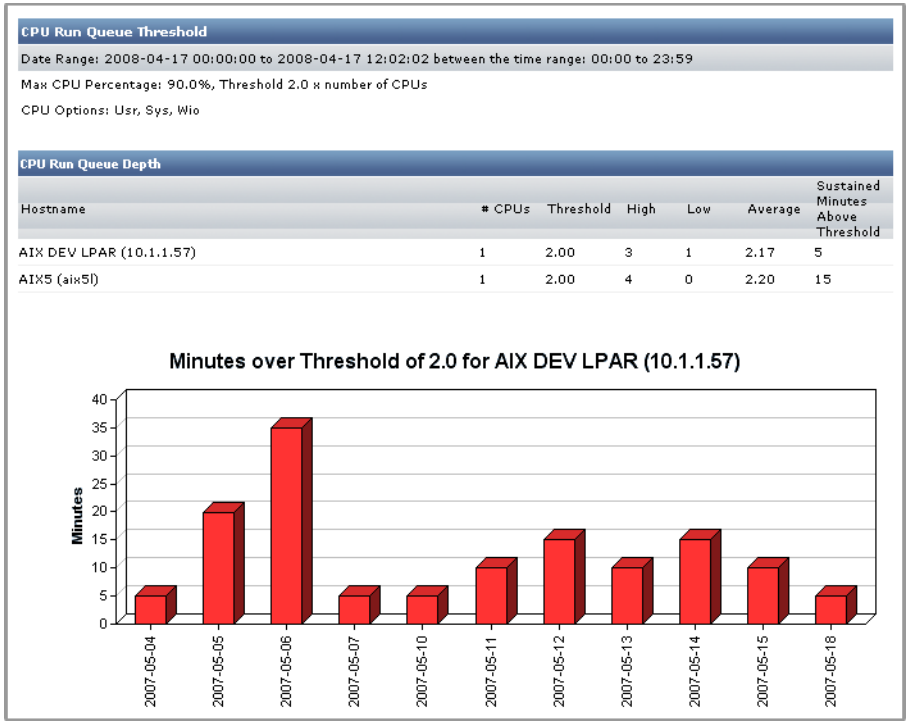

CPU Run Queue Threshold

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

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 page 253 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 n^{th} percentile for each disk or file system (e.g., although a file system may have had a high service time of 100ms, its 95^{th} 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” on page 16.

   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, as shown below:**

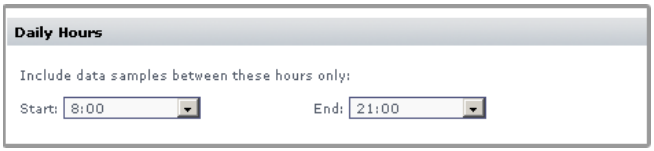

For example, if you want to report to cover the hours from 8:00 a.m. to 6:00 p.m., select `8:00` from the **Start** dropdown list, and `18:00` from the **End** dropdown list.

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” on page 204 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” on page 204 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” on page 69 for more information about views.

13 **If you are generating reports for specific Applications in your environment, select them from the List of Entities.**

14 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 for more information.

Using the File System Service Time Summary Report

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

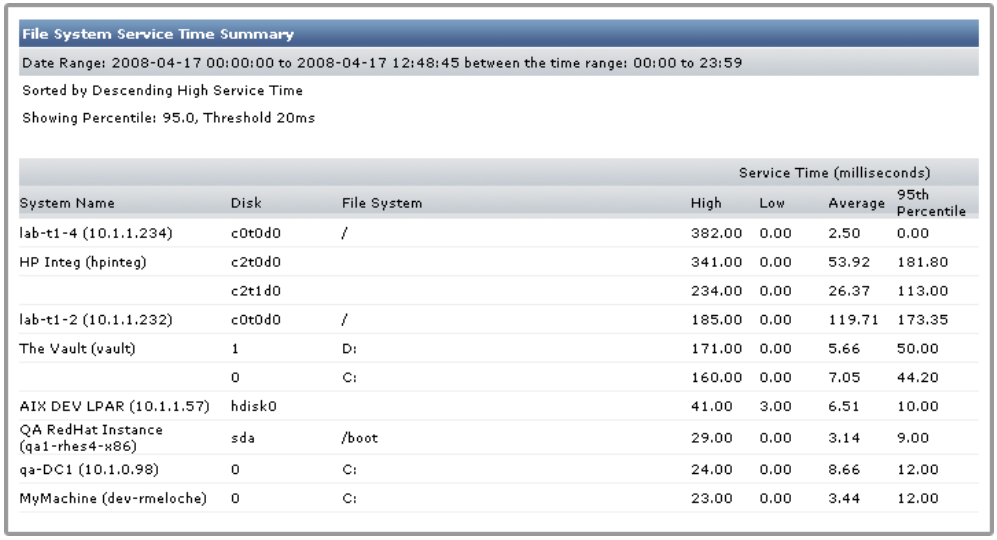

File System Service Time Summary

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

			Service Time (milliseconds)			
System Name	Disk	File System	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
- SLA Detailed Report

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 "Working with Service Level Agreements" on page 97.

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” on page 16.

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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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 “Working with Service Level Agreements” on page 97.

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” on page 16.
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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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
- Incident Priority Report
- Service Monitor Availability Report
- Service Monitor Outages Report

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” on page 62.

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” on page 16.
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.**

If you do not select this option, then a summary of the status of all Applications appears on a single line, as shown below:

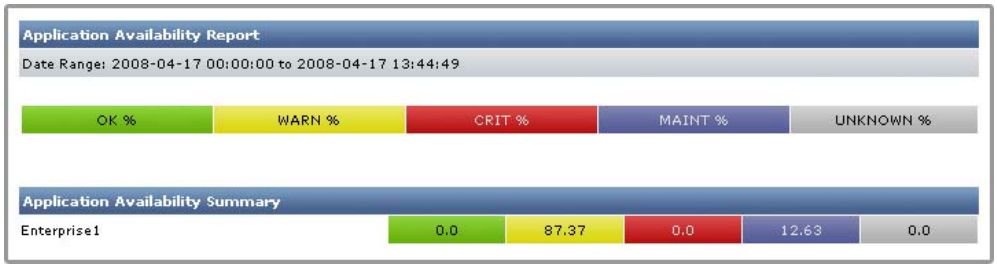

4 **If you want to generate reports for groups of systems, select the groups from the List of Groups area.**

5 **To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 69 for more information about views.

6 **If you are generating reports for specific Applications in your environment, select them from the List of Applications.**

7 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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 affectsdowntime 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” on page 143 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” on page 16.

   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.**
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” on page 69 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" on page 164 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" on page 166 and "Scheduling Reports" on page 169 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.

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" on page 16.

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" on page 69 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" on page 164 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” on page 166 and “Scheduling Reports” on page 169 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” on page 16.
3. **Select one of the following options from the Sort by dropdown list:**

   - Sample Time by Entity.
   - Service Name by Entity.
   - 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” on page 69 for more information about views.
7. **If you are generating reports for specific systems in your environment, select them from the List of Entities.**
8. **Select a report generation option. See “Report Generation Options” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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
- WebLogic Report

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” on page 16.
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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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.

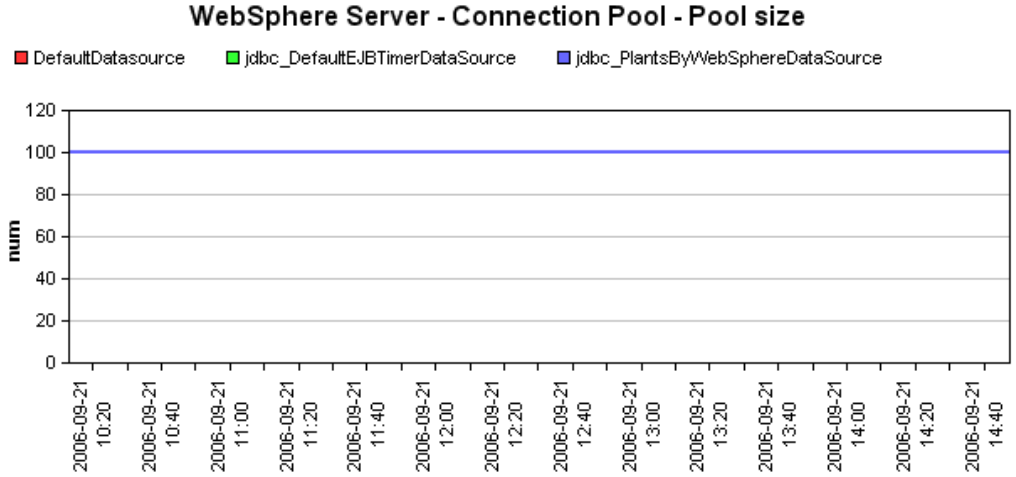

You can then adjust the size of the database connection pool to allow more connections.

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” on page 203 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” on page 16.
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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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.

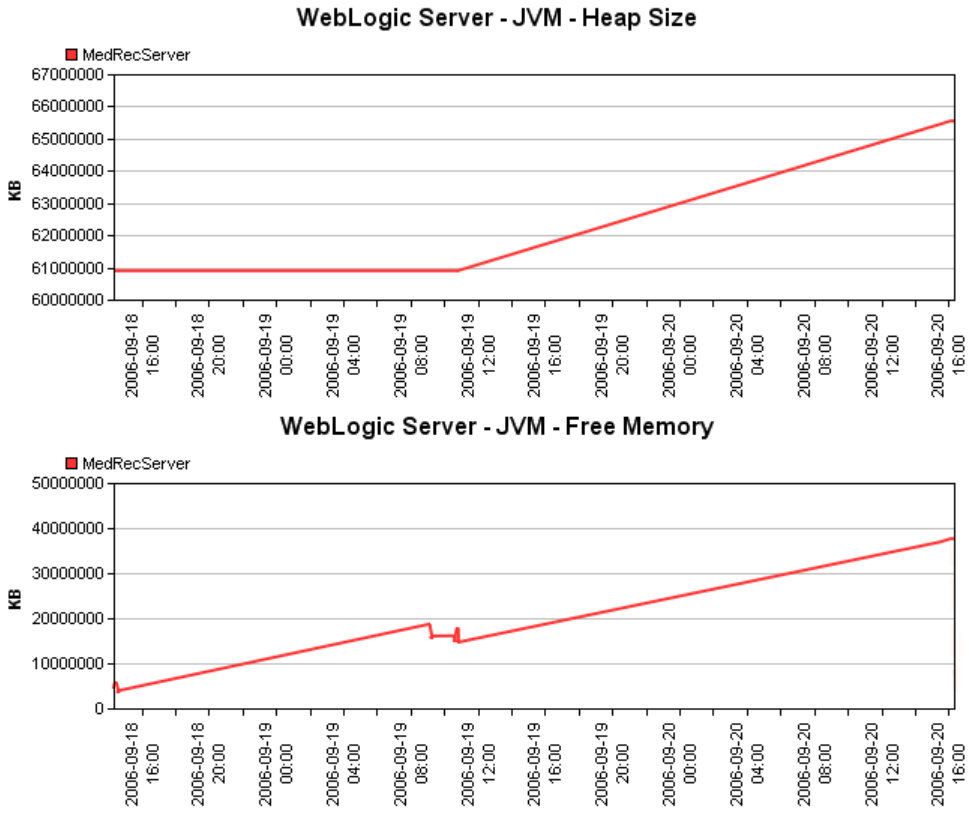

If there are increases or sudden spikes in the heap size or memory usage of the JVM, then you can tune the JVM to ensure that it is working at optimal levels.

Reports for Virtual Environments

The following reports enable you to visualize the performance of systems that are consolidated on virtual machines, whether using VMware or IBM pSeries Logical Partitions (LPARs):

- VMware Workload Report
- VMware Infrastructure Density Report
- LPAR Workload Report

VMware Workload Report

VMware ESX enables you to consolidate several servers or applications in a virtual environment. Using 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. Virtual Infrastructure 3 (VI3, or VirtualCenter) is a software suite that manages multiple, physical VMware ESX v3 servers. The latest version that supports ESX 4 is called vSphere 4 (or vCenter). VI3 or vSphere 4 enable you to manage and monitor virtual servers, as well as allocate resources among virtual machines.

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:

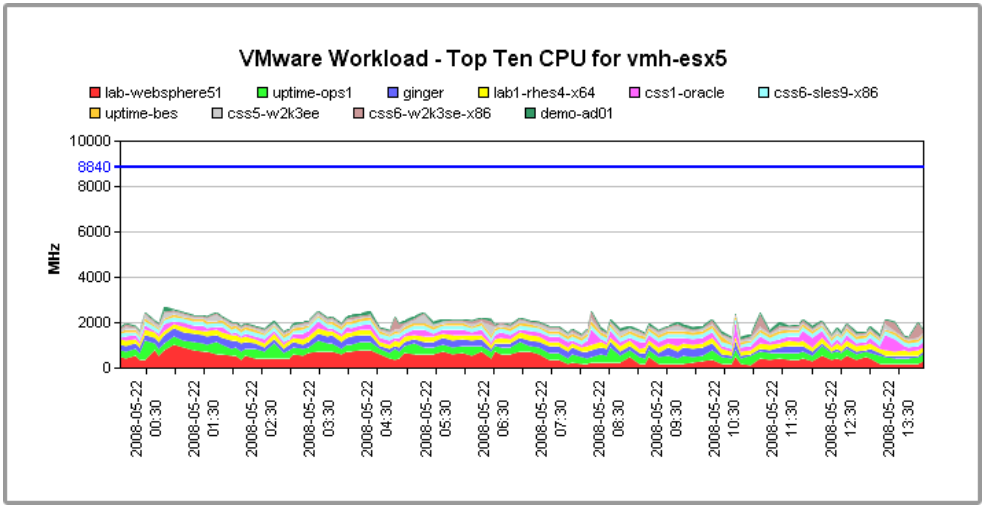

This graph indicates that, on average, the ten most CPU-intensive instances use only 20% of the server's CPU capacity. The PU on the server can handle up to three to four times its current load.

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

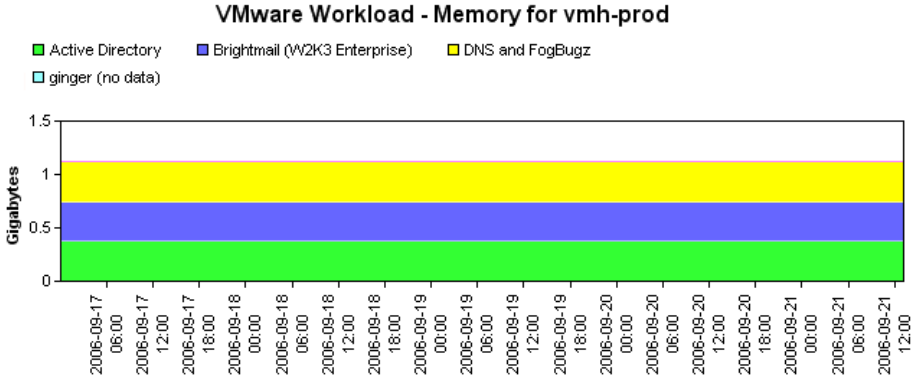

The server appears to have an ample amount of memory available.

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” on page 16.
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” on page 69 for more information about views.

6 **If you are generating reports for specific systems in your environment, select them from the List of Entities.**

7 **Select a report generation option. See “Report Generation Options” on page 164 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” on page 169 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:

- **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.
- **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.

- **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” on page 16.
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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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” on page 16.
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” on page 69 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” on page 164 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” on page 166 and “Scheduling Reports” on page 169 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.

up.time

CHAPTER 13

Understanding Graphing

This chapter introduces the graphing features of up.time in the following sections:

Graphing in up.time

You can graph performance information to learn about the behavior of a system in your environment. Graphs visualize information about CPU, memory, and process usage; as well as network, disk, and user activity. For more information about specific graphs, see “Using Graphs” on page 249.

up.time can generate performance data graphs in two ways:

- In Internet Explorer, the graph is generated using an ActiveX graphing control, as shown below:

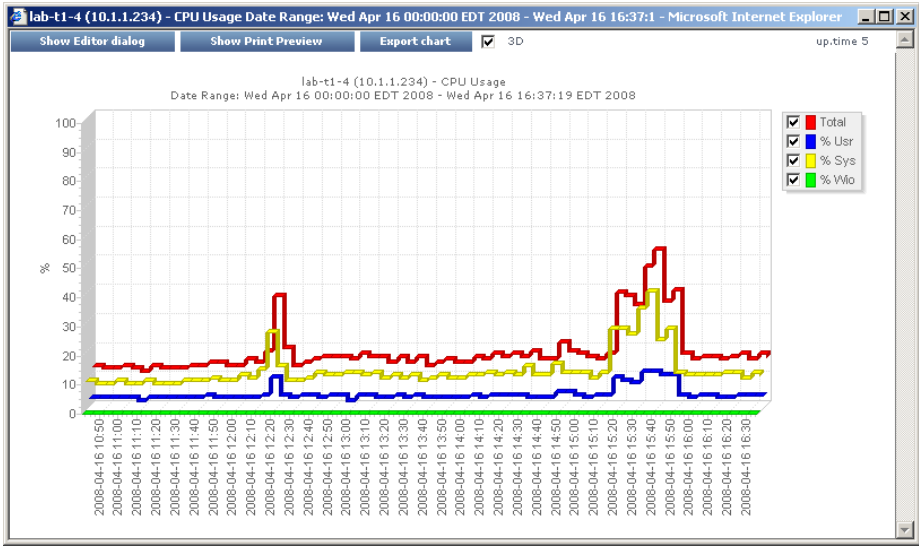

- In any Java-enabled Web browser on any operating system – for example, in Firefox, on Linux – the graph is generated using a Java graphing applet, as shown below:

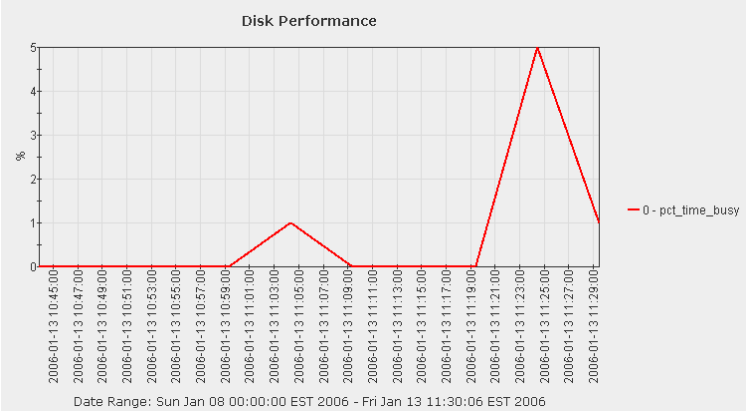

You can click any line in the graph or any item in either axis to zoom in on a particular time period or value. Click the R key on your keyboard to return to the original view.

You can modify ActiveX graphs after they have been generated. You cannot modify Java graphs.

Graphing Tool

After you generate an ActiveX graph, you can customize it using up.time's graphing tool. With the graphing tool, you can do the following:

- apply graphing line styles
- apply graphing and charting formats
- apply titles, text, and dimensioning
- manipulate a graphing axis
- apply dynamic motion to a graph

Using the Graph Editor

The Graph Editor enables you to manipulate the presentation of your graphs, as well as apply a variety of effects to a graph to change its overall look. The following image illustrates the Graph Editor:

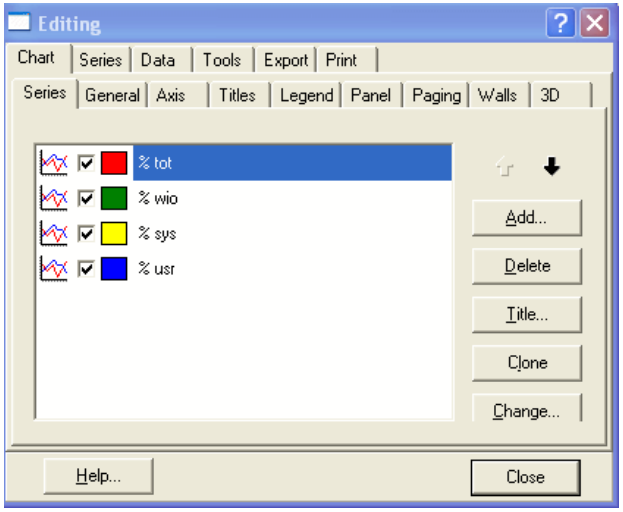

Use the Graph Editor to do the following:

- exclude graph lines
- change the style of the graph
- re-arrange the order of lines on your graph, or the actual data, to highlight specific entities in your data
- copy lines
- change the title of a line or of the graph
- change the style of graph lines, margins, titles, and the X and Y axis information

The Graph Editor contains the following subtabs:

- Series subtab

  Enables you to select the data series that the graph will display. If, for example, you have a graph that displays the following data series:

  - total memory
  - percentage of memory used by system processes
  - percentage of memory used by user processes

  You can choose to display any or all of the data series.

- General subtab

  Adjusts the graphs margins, and controls the focus and scrolling functions.

- Axis subtab

  Manipulates the graph axis, inverts the graph, scales the data points on the axis, and sets the position of the graph.

- Titles subtab

  Enables you to add, delete, or modify all labels and titles in the graph. You can, for example, change the generic title `LRX-234` to `Main Email Server`.

- Legend subtab

  Enables you to manipulate the legend – which describes the graphed information – for a graph. You can add, adjust, and delete legend information. You can also change position of the legend, and manipulate its size and format.

- Panel subtab

  Enables you to add, delete, and change the graph's background; add images or color; and apply logos to customize the look of your graph.

- Paging subtab

  Enables you to define the number of pages that your graph contains; choose to display a numeric index; and determine the number of data points that will be displayed on each page.

- Walls subtab

  Enables you to adjust the left, right, bottom, and back walls of your graph.

- 3D subtab

  Enables you to apply the following effects to graphs:

  - rotation, elevation, and zoom to adjust the depth of the graph
  - horizontal and vertical offsets
  - changes to perspective

Working with Trend Lines

A trend line is a line on a graph that indicates a statistical trend. Typically, a trend line connects multiple points on a graph. A trend line extends into the future, and you can use it to identify current and potential increases or decreases in server performance.

You can create a trend line when you need to clarify graphed information. A trend line can help you obtain a comprehensive view of the data and pinpoint any tendencies in server performance.

The following image illustrates a trend line:

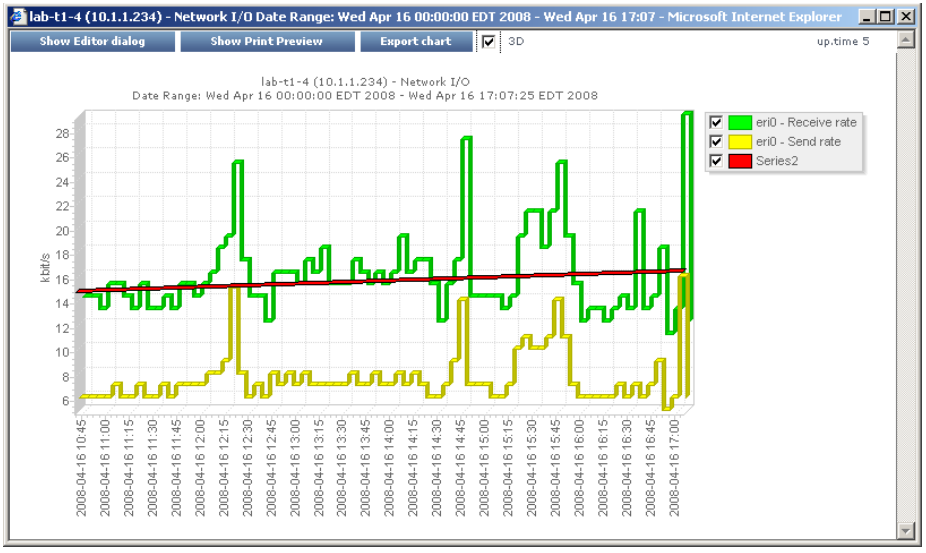

Creating a Trend Line

To create a trend line, do the following:

1. **Create a graph.**

   See “Using Graphs” on page 249 for more information.

2. **In the graph window, click Show Editor Dialog.**

3. **Click Add.**

   The **Chart Gallery** dialog box appears.

4. **Click the Functions tab, and then click the Extended subtab.**

5. **Click Trend and then click OK.**

   The **Editing** dialog box appears.

6. **In the Source Series subtab, select one or more of the available data series and then click the Add button.**

   The data series that you select are the ones for which a trend line will be generated.

7. **Click Apply.**

   up.time creates a trend line for each data series that you selected in step 6.

Formatting Individual Graph Elements

You can format individual graph Elements using the options available on the **Series** tab, and apply a different graph chart style to each Element.

Using your graphed line data, perform any of the following activities:

- Apply styles

  Changes the style of lines – for example, solid, variety of dashes, variety of dots, line thickness, visible, not visible, shape, and width.

- Apply colors and color styles

  Applies any color, image, or logo to your graphed data.

- Apply data point effects

  Makes data points visible or invisible, or displays them in two or three dimensions. You can change the following attributes of data points: style, width, height, color, border, and pattern, and image.

- Apply value formatting styles and masking

  Applies formats and masks to your data by value, percentages, horizontal axis, vertical axis, and cursor.

- Marks

  Graphs any of the following: every data point of every statistic, every data point of any statistic, and every n^{th} data point.

- Data Source

  Lists all data points by value and time. Using Data Source you can perform calculations on retrieved statistics and graph the result. You can import, perform calculations, perform contrasts and comparisons, and graph external data with collected statistics.

Exporting Graphs

Using the **Export** tab, you can send your graph by e-mail, or save it to a directory on your computer or network. You can export your graph in three ways:

- A one of the following formats: Bitmap, Metafile, SVG, Postscript, PDF, PCX, GIF, PNG, or JPEG.
- In the native up.time graph format.
- In one of the following data formats: text, HTML table, XML, or Excel.

Changing the Look and Feel of a Graph

Using the Themes tab, you can change the appearance of a graph. You can select one of eight styles for the graph, as well as specify whether the graph should be in 3D or if it should be to scale.

up.time

CHAPTER 14

Using Graphs

This chapter describes each up.time graph in the following sections:

Overview

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

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

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

UNIX vs. Windows Performance Monitoring

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

Viewing the Status of a System

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

The Quick Snapshot is typically used as a preliminary step toward root cause analysis. When you first acknowledge an issue by clicking an Element name on either **Global Scan** or the **My Alerts** section of **My Portal**, you are shown the Quick Snapshot for that Element. From here, you can scan the information provided in the charts and tables, and begin further investigation. (For example, if you notice problem while viewing the Quick Snapshot, you can generate a report to obtain more information about the problem.)

The Quick Snapshot contains the following information:

System Status Charts	Top 10 Processes	File System Statistics
•CPU Usage	•Process name	•Device
•Memory Usage	•Process ID	•Mount
•Disk I/O (transfers/sec)	•% CPU usage	•Size
•Network I/O rates	•% memory usage	•Used space
•Outages		•Available space
•Disk usage		•% used

Viewing a Quick Snapshot

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

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

Monitoring CPU Performance

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

- Usage (% busy)
- Run Queue Length
- Run Queue Occupancy

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a CPU Performance Graph” on page 256.

Usage (% busy)

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

CPU Usage in Windows

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

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

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

Optionally, you can monitor the following metrics:

- Processor: % Privileged Time

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

- Processor: %User Time

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

- Processor: % Interrupt Time

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

- System: Processor Queue Length

  The number of threads that are waiting for processor time.

CPU Usage in UNIX and Linux

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

- User Time per CPU

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

- System Time per CPU

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

- Wait I/O Time per CPU

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

Run Queue Length

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

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

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

Run Queue Occupancy

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

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

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

Generating a CPU Performance Graph

To generate a CPU performance graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click one of the following options:**
   - Usage (% busy)
   - Run Queue Length
   - Run Queue Occupancy
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click Generate Graph.**

Multi-CPU Usage

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

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

```
This system is currently listed as only having one
CPU
```

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

Generating a Multi-CPU Usage Graph

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

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click Multi-CPU Usage.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.

5 **Click one of the following options:**

- User %

  The percentage of CPU user processes that are in use. For Windows systems, this option is **% User Time**.

- System %

  The percentage of CPU kernel processes that are in use. For Windows systems, this option is **% System Time**.

- % Privileged Time

  On Windows systems, the percentage of time that the CPU spends executing kernel commands.

- Wait I/O %

  The percentage of time that a process which can be run must wait for a device to perform an I/O operation.

- SMTX

  The number of read or write locks that a thread was not able to acquire on the first attempt, as reported by the `mpstat` command.

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

- XCAL

  The number of interprocess cross-calls.

  In a multi-processor environment, one processor sends cross-calls to another processor to get that processor to do work. Cross-calls can also be used to ensure consistency in virtual memory. Heavy file system activity – such as NFS – can result in a high number of cross-calls.

- Interrupts

  The number of CPU interrupts. For Windows systems, this option is **% Interrupt Time**.

  Interrupts are a mechanism that a device uses to signal to the kernel that it needs attention, and that immediate processing is required on its behalf.

- Interrupts/sec

  On Windows systems, rate at which CPU handles interrupts from applications or hardware each second. If the value for Interrupts/sec is high, there could be problems with the hardware on the system.

- Total %

  On Solaris systems, the total amount of User %, System %, and Wait I/O %.

  On Windows systems, this option is **% Total** and is the total amount of % User Time, % Privileged Time, and % Interrupt Time.

6 **Select the CPUs to graph from the Choose CPUs to graph list.**

7 **Click Generate Graph.**

Graphing Memory Usage

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

- Used
- Cache Hit Rate
- Paging Statistics
- Free Swap

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a Memory Usage Graph” on page 262.

Used

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

Cache Hit Rate

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

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

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

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

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

Paging Statistics

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

```
scan threshold = handspreadpages ÷ residence time
```

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

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

Free Swap

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

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

On Solaris, swap space is separated into:

- Physical swap space

  The actual space on a disk available for swapping.

- Virtual swap space

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

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

Linux swaps data to a dedicated swap partition.

Generating a Memory Usage Graph

To generate a memory usage graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click one of the following options:**
   - Used
   - Cache Hit Rate
   - Paging Statistics
   - Free Swap
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click Generate Graph.**

Graphing Processes

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

- Number of Processes
- Process Running, Blocked, Waiting
- Process Creation Rate

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a Process Graph” on page 264.

up.time also has other process graphs, which collect more detailed information. For information on the other process graphs, see:

- “Displaying Detailed Process Information” on page 286.
- “Workload Graphs” on page 267.

Number of Processes

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

Process Running, Blocked, Waiting

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

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

Process Creation Rate

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

Generating a Process Graph

To generate a process graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click one of the following options:**
   - Number of Processes
   - Process Running, Blocked, Waiting
   - Process Creation Rate
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click Generate Graph.**

Graphing TCP Retransmits

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

- A header

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

- A payload

  Contains the data that is being sent.

- A trailer

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

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

Generating a TCP Retransmits Graph

To generate a TCP retransmits graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click TCP Retransmits.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click Generate Graph.**

Graphing User Activity

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

- Login History

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

- Sessions

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

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

Generating a User Activity Graph

To generate a user activity graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click either Login History or Sessions.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.

5. **Click Generate Graph.**

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

Workload Graphs

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

You can generate the following workload graphs:

- Workload - User

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

- Workload - Group

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

- Workload - Process Name

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

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a Workload Graph” on page 268.

Each workload graph captures the following metrics:

- CPU %

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

- Memory Size

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

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

- RSS

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

Workload graphs that are generated for SNMP agents only chart the Memory Size metric.

Generating a Workload Graph

To generate a workload graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click one of the following options:**
   - Workload - User
   - Workload - Group
   - Workload - Process Name
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click one of the following metrics:**
   - CPU %
   - Memory Size or Virtual Bytes (on UNIX and Windows, respectively)
   - RSS or Working Set (on UNIX and Windows, respectively)

You can only graph one metric at a time.

6. **Select one or more of the available users, groups, or processes from the list.**

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

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

7 Click Add.

8 Click Generate Graph.

Workload Top 10 Graphs

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

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

Generating a Workload Top 10 Graph

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

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click one of the following options:**
   - Workload Top 10 - User
   - Workload Top 10 - Group
   - Workload Top 10 - Process Name
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click one of the following options:**
   - CPU %
   - Memory Size
   - RSS

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

LPAR Workload Graphs

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

- Workload - CPU

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

- Workload - Memory

  The total amount of memory being used by an LPAR.

- Workload - Disk

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

- Workload - Network

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

You can also graph the CPU entitlement of individual LPARs using the CPU Utilization graph. See “LPAR CPU Utilization Graphs” for more information.

Generating an LPAR Workload Graph

To generate an LPAR Workload graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the pSeries server which is hosting the LPARs whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click one of the following options:**

   - Workload - CPU
   - Workload - Memory
   - Workload - Disk
   - Workload - Network

4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 16.

5 **Click Generate Graph.**

LPAR CPU Utilization Graphs

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

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

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

1 **In the Global Scan or My Enterprise panel, click the name of the pSeries server which is hosting the LPAR whose information you want to graph.**

2 **In the Tree panel, click the Graphing tab.**

3 **Under the LPAR Workload heading, click Workload - CPU Utilization.**

4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 16.

5 **Select the name of the LPAR whose information you want to graph.**

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

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

6 **Click Generate Graph.**

Network Graphs

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

- I/O
- Errors
- NetFlow

The I/O and Errors graphs use the same input criteria, but return different data. NetFlow graphs are available if up.time is integrated with Scrutinizer. For information on how to generate these graphs, see “Generating a Network Graph” on page 274.

I/O

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

The I/O graph captures the following statistics:

- In bytes

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

- Out bytes

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

Errors

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

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

- Excessive noise.

- Cabling problems.
- Problems with backbone connections.

The Errors graph captures the following statistics:

- In Errors

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

- Out Errors

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

- Collisions

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

NetFlow

The NetFlow graphing function transfers you to your Scrutinizer instance.

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

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

See Generating a Network Graph for information on how to generate this graph.

Generating a Network Graph

To generate network graphs, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**

3 **Click one of the following options:**

- I/O
- Errors
- NetFlow (available if up.time has been integrated with Scrutinizer)

4 **For I/O and Errors graphs, select the start and end dates and times for which the graph will chart data. For NetFlow, select one of the set time frames.**

For more information, see “Understanding Dates and Times” on page 16.

5 **For I/O and Errors graphs, select one or more network interfaces from the Available Interfaces list, and then click Add.**

6 **Click Generate Graph.**

Disk Performance Statistics Graph

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

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

Generating a Disk Performance Statistics Graph

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

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click Disk Performance Statistics.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Select one of the following options:**

   - Percent Busy

     The percentage of the disk capacity that is being used.

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

   - Average Queue

     The average number of processes that are waiting to access the disk.

     The length of the queue is affected by how busy the system is and the amount of time that each transaction requires to perform a disk operation. A complete transaction must occur before the next transaction can start. Longer disk operations per transaction increases the average length of the queue.

- Read/Writes

  The number of read/write requests, per second, from or to a disk.

- Throughput (blks/s)

  The amount of disk traffic, in blocks of 512 bytes, that is flowing to and from a disk each second.

- Average Wait Time

  The average time, in milliseconds, that a transaction is waiting in a queue. The wait time is directly proportional to the length of the queue.

- Average Serve Time

  The average time, in milliseconds, required to perform a task.

- All of the above for one disk

  up.time graphs all of the metrics listed above for a single disk.

6 **Select the disks for which you want to collect information from the list.**

If you select multiple disks and selected **All of the above for one disk** in step 5, then up.time only graphs information for the first disk that you selected.

7 **Click Generate Graph.**

Top 10 Disks Graph

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

Generating a Top 10 Disks Graph

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

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click Top 10 Disks.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Select one of the following options:**

   - Percent Busy

     The percentage of the disk capacity that is being used.

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

   - Average Queue

     The average number of processes that are waiting to access the disk.

     The length of the queue is affected by the amount of time that each transaction requires to perform a disk operation. For both sequential and random disk transactions, a complete transaction must occur before the next transaction can begin. Longer disk operations per transactions increase the average length of the queue.

- Read/Writes

  The number of read/write requests per second from or to a disk.

- Throughput (blks/s)

  The amount of traffic, in 512 byte blocks, that is flowing to and from a disk.

- Average Wait Time

  The average time, in milliseconds, that a transaction is waiting in a queue. The wait time is directly proportional to the length of the queue.

- Average Serve Time

  The average time, in milliseconds, required to perform a task.

6 **Click Generate Graph.**

File System Capacity Graph

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

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

The File System Capacity graph visualizes the following statistics:

- Total Size

  The total amount of space available on the system.

- Space Used

  The amount of space on the file system that has been used.

Generating a File System Capacity Graph

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

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click File System Capacity.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see "Understanding Dates and Times" on page 16.

5. **Select one or more file systems from the list.**

   If you are generating a graph for a Windows system, you will only be able to generate a graph for the `C:\` drive.

6. **Click Generate Graph.**

VXVM Stats Graph

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

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

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

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

In the **Info & Rescan** panel, verify that the entry **Has a Logical Volume Manager?** is set to **Yes**. If it is, then ensure that you can connect to the host from the Enterprise Monitoring Station. See "Viewing System and Service Information" on page 38 for more information.

Generating a VXVM Stats Graph

To generate a VXVM Stats graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click VXVM Stats.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see "Understanding Dates and Times" on page 16.

5 **In the Available Disk Groups and Volumes area, select one or more volumes on which to report.**

The disk groups or volumes that appear in this area will vary from system to system. You must select at least one disk group or volume.

6 **Select one of the following options:**

- I/O Operations

  The number of times, per second, that data is written to and read from the volume.

- Block Throughput

  The amount of disk traffic, in blocks of 512 bytes, that is flowing to and from the volume.

- Average Service Times

  The average amount of time, in milliseconds, that is required for a request to be carried out.

7 **If necessary, uncheck either of the Read or Write checkboxes.**

Depending on the option you chose in step 6, the Read and Write options chart the following information in the graph:

- If you selected **I/O Operations** in step 6, the number of read and write operations to and from the volume.
- If you selected **Block Throughput** in step 6, the number of blocks that were read and written to and from the volume.
- If you selected **Average Service Times** in step 6, the amount of time requires to read and write data to and from the volume.

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

8 **Click Generate Graph.**

Novell NRM Graphs

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

- Available Memory

  The amount of memory that is not allocated to any service.

- DS Thread Usage

  The number of server threads that Novell eDirectory uses. The server thread limit ensures that server threads are available for other functions as needed.

- Work To Do Response Time

  The amount of time that a Work To Do process requires to run from the time a process is scheduled.

- Allocated Server Processes

  How the service processes are allocated on the NRM system.

- Available Server Processes

  The number of available processes on the NRM system.

- Abended Thread Count

  The number of threads that have *abended* (ended abnormally) and that are suspended because of abended recovery.

- Packet Receive Buffers

  The status of Packet Receive Buffers (which transmit and receive packets) for the NRM system.

- Available ECBs

  The status of available Event Control Blocks (ECBs), which are Packet Receive Buffers that have been created but which are not currently being used.

- LAN Traffic

  Whether or not the NRM system can transmit and receive packets.

- Available Disk Space

  The status of the available disk space on a server.

- Disk Throughput

  The status of amount of the data being read from and written to the storage media on the server.

- Connection Usage
- The number of connections that are being used, and the peak number of connections used on this server.

Generating a Novell NRM Graph

To generate a Novell NRM graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the Novell NRM system whose information you want to graph.**

   Novell NRM systems are denoted by this icon:  .

2. **In the Tree panel, click the Graphing tab and then click one of the metrics in the list.**
3. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.

4. **Click Generate Graph.**

Instance Motion Graphs

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

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

Generating an Instance Motion Graph

To generate an Instance Motion graph, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the ESX instance whose motion you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click Instance Motion.**
4. **Select the start and end dates and times for which the graph will chart data.**

   For more information, see “Understanding Dates and Times” on page 16.
5. **Click Generate Graph.**

Displaying Detailed Process Information

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

- Process

  The name of the process, which is taken from its executed path name.

- PID

  The number that identifies the process.

- PPID

  The number that identifies the parent process. The PPID can help identify possible relationships between processes.

  On Windows systems, the PPID is called the *Creating Process ID*.

- UID

  The ID of the user or account that has been consuming CPU time.

  On Windows systems, the UID is called the *Owner*.

- GID

  The ID of the group that has been consuming CPU time.

  On Windows systems, the GID is called the *Group Name*.

- Memory Used

  The amount of memory, expresses as a percentage of total available memory, being consumed by a process.

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

  The **Memory Used** value can be misleading because shared memory between processes is counted multiple times. For example, if five Oracle processes are using 10% of available memory, this does not indicate that Oracle is consuming 50% of system memory.

- RSS

  Run Set Size – the amount of physical memory that is being used.

  On Windows systems, RSS is called the *Working Set*.

- CPU %

  The percentage of the CPU time used by the process, calculated by dividing total used CPU Time by the process' running time; if applicable, the result is further divided by the number of CPUs for the Element on which the process is running.

  On Windows systems, the CPU % is called *% Processor Time*.

- User Time

  The amount of time (in seconds) that a particular user, group, or account has been using the CPU.

  This value is not displayed for Windows systems.

- User System Time

  The amount of time (in seconds) that a process has been consuming system time on the CPU.

  This value is not displayed for Windows systems.

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

- Start Time

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

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

Generating Detailed Process Information

To display detailed process information, do the following:

1. **In the Global Scan or My Enterprise panel, click the name of the system whose information you want to graph.**
2. **In the Tree panel, click the Graphing tab.**
3. **Click Detailed Process Information.**

4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 16.

5 **Click Display Process Information.**

A window containing a chart that lists the process information for the time period that you specified appears. The following image illustrates process information for a Solaris system:

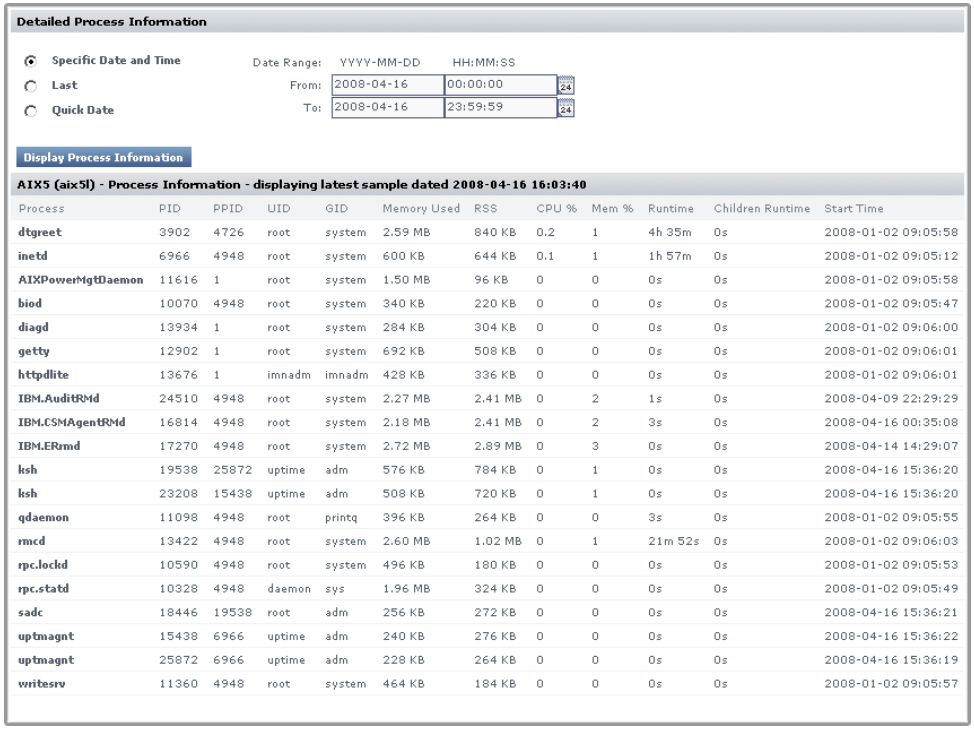

Detailed Process Information

Specific Date and Time — Date Range: YYYY-MM-DD HH:MM:SS
Last — From: 2008-04-16 00:00:00
Quick Date — To: 2008-04-16 23:59:59

Display Process Information

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

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

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

CHAPTER 15

Configuring and Managing up.time

The configuration and management of up.time, mainly through the **Config Panel** and `uptime.conf` file, is described in the following sections:

Overview

up.time includes user-definable parameters that can control some aspects of its behavior including the following:

- Database Settings
- Mail Server Settings
- Global Scan threshold settings
- Resource Scan threshold settings
- Proxy settings
- Remote reporting settings
- RSS feed settings
- Splunk integration settings
- Web monitor settings

From a configuration perspective, there are two types of parameters:

- parameters whose modification does not require a restart of the Core service (also known as the up.time Data Collector service); these parameters can be modified in up.time, on the **Config** panel
- parameters whose modification requires a restart of the Core service; these parameters are found in the `uptime.conf` file

Modifying up.time Config Panel Settings

Configuration parameters that are not directly tied to, thus do not require a restart of, the up.time Core service can be modified directly in the up.time GUI (shown below):

In general, to edit these configuration settings in the up.time interface, do the following:

1. **On the up.time tool bar, click Config.**
2. **In the Tree panel, click up.time Configuration.**
3. **Enter the configuration variable and new value.**
4. **Click Update to save your changes.**

Only the variables whose default values have been modified appear in **up.time Configuration**.

Modifying uptime.conf File Settings

Configuration parameters that are directly tied to the up.time Core service are found in the `uptime.conf` file. `uptime.conf` is a text file that you can modify in any text editor, and can be found in the root up.time installation directory.

In addition to the up.time database, `uptime.conf` parameters affect a variety of up.time behavior.

Not all of the settings listed in this section will necessarily be found in your particular `uptime.conf` file.

Stopping and Restarting up.time Services

In addition to the Web interface, the up.time Enterprise Monitoring Station consists of the following services:

- DataStore
- Web server
- Data Collector (also called the Core)

These services run in the background and start automatically after the operating system on the server hosting up.time starts. However, system administrators may need to stop the up.time services – for example, before making configuration changes to the `uptime.conf` file, performing an upgrade, or archiving the DataStore.

Stopping the up.time Services

To stop the up.time services in Windows, do the following:

1. **Select Start > Control Panel.**
2. **Double click Administrative Tools, and then double click Services.**
3. **In the Services window, find the following entries and click Stop the service:**
   - up.time Web Server
   - up.time Data Collector
   - up.time Data Store

To stop the up.time services on Solaris or Linux, do the following:

1. **Log into the Enterprise Monitoring Station as user root.**

2. **Type the following command to stop the Web server:**

   ```
   /etc/init.d/uptime_httpd stop
   ```

3. **Type the following command to stop the Data Collector:**

   ```
   /etc/init.d/uptime_core stop
   ```

4. **Type the following command to stop the database:**

   ```
   /etc/init.d/uptime_datastore stop
   ```

Starting the up.time Services

To restart the up.time services in Windows, do the following:

1. **Select Start > Control Panel.**
2. **Double click Administrative Tools, and then double click Services.**
3. **In the Services window, find the following entries and click Start the service:**

   - up.time Data Store
   - up.time Data Collector
   - up.time Web Server

To restart the up.time services on Solaris or Linux, do the following:

1. **At the command line, log into the Enterprise Monitoring Station as user root.**
2. **Type the following command to start the database:**

   ```
   /etc/init.d/uptime_datastore start
   ```

3. **Type the following command to start the Data Collector:**

   ```
   /etc/init.d/uptime_core start
   ```

4. **Type the following command to start the Web server:**

   ```
   /etc/init.d/uptime_httpd start
   ```

Interfacing with up.time

Some of the Enterprise Monitoring Station's features require integration with other elements that make up your infrastructure. In some cases configuration is mandatory (e.g., an SMTP server will need to have been set at the time of installation), while in others it is required only when particular up.time features are used. The following sections outline how to configure up.time to communicate with servers.

Monitoring Station Web Server

Monitoring Stations include a Web server component that drives the user interface. Any Monitoring Station, whether EMS or LDC, that is accessed by users or administrators requires a URL. The Web address used to access the Enterprise Monitoring Station is configured through the following `uptime.conf` parameter:

```
httpContext = http://<hostname>:<port>
```

- `<hostname>` is the host name of the server on which up.time is running (e.g., `localhost`)
- `<port>` is the port on which the up.time Web server is listening for requests (e.g., `9999`); you can optionally omit the port number

If the up.time interface is being accessed via SSL, the value for this parameter should be stated as `https` instead of `http`.

SMTP Server

up.time uses a mail server to send alerts and reports to its users. After installing up.time for the first time, the administrator was asked to enter SMTP server information. These initial values can be modified in the **Mail Servers** configuration panel.

Modifying the SMTP Server Used by up.time

To configure up.time's mail server, do the following:

1. **On the up.time tool bar, click Config.**

2. **In the Tree panel, click Mail Servers.**
3. **In the sub panel, click Edit Configuration.**
4. **Type the name of the mail server in the SMTP Server field.**

   This value was set the first time the up.time administrator logged in after installation; the default value is the name of the host on which the Monitoring Station was installed at that time.

   The name of the server could follow the "`smtp.<domain_name>`" convention, or could be its host name or IP address.
5. **Optionally, enter the port used by the mail server in the SMTP Port field.**
6. **In the SMTP Sender field, enter the email address that up.time uses to send alert notifications and reports.**

   This value was set the first time the up.time administrator logged in after installation, and should be set to your domain (e.g., `admin@mail.uptimesoftware.com`).

   A sender's name can be encapsulated with double quotes, in which case, the email address is encapsulated with angled brackets:
   `"uptime administrator" <admin@uptimesoftware.com>`
7. **In the SMTP Helo String field, enter the string that identifies the domain from which a message is being sent.**

   For example, `uptimesoftware.com`.
8. **In the SMTP User field, enter the user name that is used to authenticate connections with the SMTP server.**
9. **In the SMTP Password field, enter the password that is used to authenticate connections.**
10. **Click Save.**

    The edit window closes, and you are returned to the **Mail Server Configuration** panel.
11. **To test the mail server configuration, click the Test Configuration button.**

    The Enterprise Monitoring Station will try to send an email message containing the configuration information to the email address of the up.time administrator. If an error message appears in the subpanel, edit and then re-test the configuration.

RSS Feed Settings

up.time displays a list of recent knowledge base articles in the **My Portal** panel. This list is fed to the **My Portal** panel via RSS (Really Simple Syndication, a method for delivering summaries of and links to Web content). Clicking the title of an article opens it in your Web browser.

By default, RSS feeds are drawn directly from the uptime software Support Portal without the use of proxy server information. If your Enterprise Monitoring Station accesses the Internet through one, feeds will most likely not be available, and the following message will appear in the **My Portal** panel:

You can change the RSS feed settings to point to the proxy server rather than directly to the uptime software Web site by manually inputting settings in the **up.time Configuration** panel, as outlined in "Modifying up.time Config Panel Settings" on page 291.

Changing Proxy Server Information for RSS Feeds

You can manually configure the settings for RSS feeds through the following parameters (default values, if applicable, are shown):

- `rssFeedUrl=http://support.uptimesoftware.com/rss/kb.xml`

  The URL of the RSS feed.

- `httpProxyHost`

  The host name of the proxy server that the Enterprise Monitoring Station uses to access the Internet.

- `httpProxyPort`

  The port through which the Enterprise Monitoring Station communicates with the proxy server.

- `httpProxyUsername`

  The user name required to use the proxy server.

- `httpProxyPassword`

  The password required to use the proxy server.

VMware vCenter Orchestrator Integration

Administrators can configure Action Profiles to automatically carry out tasks in the event of an up.time alert. One such task is the initiation of contact with VMware vCenter Orchestrator, and the execution of a workflow. To have access to this functionality, up.time needs to know how to communicate with Orchestrator.

For information about Action Profiles and VMware vCenter Orchestrator, see “Action Profiles” on page 153.

Integrating up.time with VMware vCenter Orchestrator

To configure up.time integration with Orchestrator to execute workflows, do the following:

1. **On the up.time tool bar, click Config.**
2. **In the Tree panel, click VMware vCenter Orchestrator.**
3. **In the sub panel, click Edit Configuration.**
4. **Ensure the VMware Orchestrator Enabled check box is selected.**
5. **In the VMware Orchestrator Server field, enter the host name of, or IP address assigned to the Orchestrator server when it was configured.**
6. **In the VMware Orchestrator Port field, enter the port the Orchestrator server was configured to use in order to communicate with other systems.**
7. **Optionally select the Use SSL check box if Orchestrator was configured to use an SSL certificate.**
8. **Enter the Username and Password of an appropriate user account on the Orchestrator server.**

For proper integration, an Orchestrator account with View and Execute permissions is required.

9 **Click Save.**

The configuration window closes, and you are returned to the **VMware vCenter Orchestrator Configuration** panel.

10 **To ensure the settings you provided are correct, click the Test Configuration button.**

The Enterprise Monitoring Station will try to communicate with the VMware vCenter Orchestrator server. If an error message appears in the subpanel, edit and then re-test the configuration.

Remote Reporting Settings

If you are using a reporting instance (an up.time instance that only generates and serves reports), the remote reporting settings enable you to specify the location of the reporting instance, and the port on which it is listening.

Modifying the Remote Reporting Server Settings

To configure the remote reporting instance used by up.time, do the following:

1 **On the up.time tool bar, click Config.**

2 **In the Tree panel, click Remote Reporting.**

3 **In the sub panel, click Edit Configuration.**

4 **Ensure the Reporting Instance Enabled check box has been selected.**

5 **In the Remote Reporting Server field, enter the host name or IP address of the server on which the remote reporting instance is found.**

6 **Enter the port used to communicate with the server.**

7 **Click Save.**

The edit window closes, and you are returned to the **Remote Reporting Instance Configuration** panel.

8 **To test the remote reporting server configuration, click Test Configuration.**

A pop-up window appears, indicating whether up.time was able to connect to the remote reporting instance. If an error message is displayed, correct your configuration and re-test it.

Note that the modification of these values is one of a series of steps performed to correctly set up a remote reporting instance. Refer to the Knowledge Base article entitled "*Setting up a reporting instance*" for more information.

User Interface Instance Settings

A UI instance is an up.time installation that does not perform any data collection tasks, and is primarily used for real-time monitoring and report generation. UI instances can divert traffic from a standard Monitoring Station implementation, and are helpful when there are many up.time users who do not need to perform full administrative tasks.

You can manually configure UI instance settings with the following `uptime.conf` parameters:

- `uiOnlyInstance = true`

  Determines whether the Monitoring Station functions only as a user interface instance.

- `uiOnlyInstance.monitoringStationHost = HOSTNAME`

  The host name or IP address of the up.time Monitoring Station that is performing data collection, and to which this UI instance will connect.

- `uiOnlyInstance.monitoringStationCommandPort = 9996`

  The port through which the UI instance can communicate with the data-collecting Monitoring Station.

A Monitoring Station that is acting as a UI instance must have the same database settings as the data-collecting Monitoring Station.

Scrutinizer Settings

Scrutinizer is a NetFlow analyzer that can be installed to monitor network traffic managed by compatible switches and routers. Scrutinizer can be integrated with **Global Scan**, as well as up.time's graph generation for node-type Elements, and other hosts that are also monitored with Scrutinizer.

In order to access Scrutinizer, up.time needs to be pointed to your installation.

Modifying the Scrutinizer Settings

You can configure Scrutinizer's integration with up.time through the following parameters:

- `netflow.enabled`

  Determines whether Scrutinizer is integrated with the Enterprise Monitoring Station.

- `netflow.hostname`

  The host name or IP address of your Scrutinizer installation.

- `netflow.port`

  The HTTP port through which Scrutinizer sends and receives communication.

- `netflow.username`

  The user name required to log in to Scrutinizer.

- `netflow.password`

  The password required to log in to Scrutinizer.

Splunk Settings

Splunk is a third-party search engine that indexes log files and data from the devices, servers, and applications in your network. Using Splunk, you can quickly analyze your logs to pinpoint problems on a server or in a network, or ensure that you are in compliance with a regulatory mandate or

Service Level Agreements. You install Splunk on a server in your datacenter.

When values are provided for the Splunk settings listed below, the Splunk icon (splunk>) will appear in the **My Portal** panel beside the names of services that are in WARN or CRIT states. When you click the Splunk icon, you will be automatically logged in to your Splunk search page.

You can change your up.time-Splunk integration by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 291.

Changing Splunk Server Information for up.time

You can enable automatic login to the Splunk search page, or modify an existing configuration through the following parameters:

- `splunk.url`

  The URL of the server on which your Splunk search page is hosted (e.g., `http://webportal:8000`).

- `splunk.username`

  The user name required to log in to your Splunk search page.

- `splunk.password`

  The password required to log in to your Splunk search page.

- `splunk.soapurl`

  The URL that points to the SOAP management port that Splunk uses to communicate with the splunk daemon (e.g., `https://webportal:8089`).

  In the URL, you must include the port on which the Splunk server listens for requests. See the Splunk Admin Manual for more information.

- `splunk.version`

  The version of Splunk you are using.

Archiving the DataStore

Depending on the amount of disk space available for the continuously growing DataStore, administrators can set an archive policy that determines how many month's worth of data is retained. Old performance data is automatically archived and removed from the DataStore. This archiving procedure works with all databases that are compatible with up.time.

The existing archive policy can be viewed and modified on the **Archive Policy** subpanel, which is accessed from the main **Config** panel. Here, the main archive categories are shown, along with the number of months for which collected data is retained in the DataStore.

Every month, up.time checks the DataStore's entries; data that is older than the limit set in the archive policy are written to XML files. The XML archives use the following format:

```
<table_name>_<date>.xml.gz
```

The archives created reflect the database table structure used to store performance data, as well as the date that the stored data represents:

```
performance_cpu_2006-09-13.xml.gz
```

The DataStore is trimmed and the XML files are compressed and stored in the `/archives` directory.

For example, if you installed up.time in the default location, the path to the archived data will be:

- Linux: `/usr/local/uptime/archives`
- Solaris: `/opt/uptime/archives`
- Windows: `C:\Program Files\uptime software\uptime\archives`

Windows Vista users can find the DataStore archive in the Virtual Store instead of the default location (i.e., C:\Users\uptime\AppData\Local\VirtualStore\Program Files\<uptime-install-directory>

Once backed up, archives can be stored offline. If required, they can be temporarily imported into the DataStore.

Archive Categories

The following table lists the statistical categories whose archiving can be configured, along with the corresponding DataStore database table:

Archive Policy Category	Database Table
Overall CPU/Memory	performance_cpu
Multi-CPU	performance_aggregate
Detailed Process	performance_psinfo
Disk Performance	performance_disk
File System Capacity	performance_fscap
Network	performance_network
User Information	performance_who
Volume Manager	performance_vxvol
Retained Data	erdc_int_data erdc_decimal_data erdc_string_data

Configuring an Archive Policy

To set an archive policy, do the following:

1. **On the up.time tool bar, click Config.**
2. **In the Tree panel, click Archive Policy.**
3. **For the following categories, specify the number of months worth of data that will be retained in the DataStore before being removed and archived:**
   - Overall CPU/Memory Statistics
   - Multi-CPU Statistics
   - Detailed Process Statistics

- Disk Performance Statistics
- File System Capacity Statistics
- Network Statistics
- User Information Statistics
- Volume Manager Statistics
- Retained Data

4 **Ensure the Enable Archiving checkbox is selected.**

5 **Click Set Archive Policy.**

6 **Optionally, you can click the Archive Now button to immediately create archives of the data in your DataStore.**

up.time will check the DataStore entries and archiving anything that is older than the limits you have configured.

Restoring Archived Data

If you need to generate graphs or reports on older data that has already been archived, and is no longer in the DataStore, you can import specific archives using the `restorearchive` command line utility. The command's parameters allow you to import archives in the following manner:

- a single archive that represents a specific archive category and date; the collected data for each archive category and 24-hour period is exported to individual XML files
- all archives for a specific date (i.e., 24-hour period)

Importing Archived Data into the DataStore

To import archived data into the DataStore, do the following:

1 **At the command line, navigate to the following directory:**

- Linux: `/usr/local/uptime/scripts/`
- Solaris: `/opt/uptime/scripts/`

- Windows: `C:\Program Files\uptime software\uptime\archives`

2 **Run the restorearchive command with one or more of the following options:**

- `-f <filename>`

  Imports a single file (i.e., an archive category's data for a single date). You must specify the full path to the file name.

- `-d <date>`

  Imports all files with the specified date (in YYYY-MM-DD format).

- `-D <directory>`

  The directory containing the archived files. Note that you must specify this option when using the `-d` option.

- `-c <directory>`

  The full directory path to the file `uptime.conf`.

For example, enter the following command to import all of the data archived on September 18, 2006 which are located in the default directory for archived data:

```
restorearchive -d 2006-09-18 -D /usr/local/uptime/
archives/ -c /usr/local/uptime
```

up.time Diagnosis

The following options assist you with diagnostic steps that you may need to perform should you encounter problems with up.time. You have access to two types of logs: system logs and audit logs that track user actions. Additionally, you can generate a problem report for up.time Customer Support if further analysis is required.

System and audit logs are written to the `/logs` directory, and problem reports are found in the `/GUI` directory, both of which are found in the up.time installation directory:

- Linux: `/usr/local/uptime/`
- Solaris: `/opt/uptime/`
- Windows: C:\Program Files\uptime software\uptime

Windows Vista users can find the audit log in the Virtual Store instead of the default location (i.e., C:\Users\uptime\AppData\Local\VirtualStore\Program Files\<uptime-install-directory>

System Event Logging

up.time automatically logs system events to the `/logs` directory. These weekly logs follow the `uptime.log.<year>-<week>.log` naming format. You can determine the type of system information up.time writes to the log by using one of the following values:

- `DEBUG`
- `INFO`
- `WARN`
- `ERROR`
- `FATAL`
- `ALL`

- OFF

The default setting, DEBUG, essentially logs all system event types. To reduce the number of log entries, you can limit logging to events with a higher level of severity, from INFO to FATAL. Note that each severity level is a subset of higher levels (e.g., setting loggingLevel to WARN means any WARN-, ERROR- or FATAL-level events are written to the log).

Logging is configured through the following uptime.conf parameter:

```
loggingLevel = DEBUG
```

Audit Logs

up.time can record changes to the application's configuration in an audit log. The details of the configuration changes are saved in the audit.log file, which is found in the /logs directory.

There are many uses for the audit log. For example, you can use the audit log to track changes to your up.time environment for compliance with your security or local policies. You can also use the audit log to debug problems that may have been introduced into your up.time installation by a specific configuration change; the audit log enables you to determine who made the change and when it took effect.

The following is an example of an audit log entry:

```
2006-02-23 12:28:20,082 - kdawg: ADDSYSTEM [cfgcheck=true,
port=9998, number=1, use-ssl=false, systemType=1,
hostname=10.1.1.241, displayName=MailMain,
systemSystemGroup=1, serviceGroup=, description=,
systemSubtype=1]
```

Audit Logging is enabled or disabled, with "yes" or "no" values, respectively, through the following uptime.conf parameter:

```
auditEnabled = yes
```

Problem Reporting

When you encounter a problem with up.time, Client Care needs specific information to diagnose and fix the problem. up.time can automatically collect this information and compress it in an archive which you can send to Client Care.

The archive contains the following: up.time configuration files; system information; log files; database information and error files; and a listing of the DataStore directory. Optionally, the archive will also contain a copy of the configuration data from your DataStore.

The archive is saved to the `GUI/problemreports` directory on the Enterprise Monitoring Station and has a file name with the following format:

`prYYYYMMDD-HHMMSS.zip`

- `YYYYMMDD` is the date on which the report was generated (e.g., `20061212`).
- `HHMMSS` is the time at which the report was generated (e.g., `142306`).

Generating a Problem Report

To generate a problem report, do the following:

1. **On the up.time tool bar, click Config.**
2. **In the Tree panel, click Problem Reporting.**

   If you have generated problem reports in the past, they appear in the subpanel.
3. **If you do not want to include a copy the configuration data from your DataStore, click the Include config database dump option.**
4. **Click the Generate Report button.**

   A message such as the following appears in the subpanel:

   `Problem report created : pr20061017-094927.zip`

   Click the name of the problem report to download it to your local file system, then send the archive to uptime software Client Care.

up.time Measurement Tuning

In some cases, you can make measurement adjustments to up.time's default values. Changes can be made to the following:

- the number of threads allocated to service monitors
- status thresholds in the **Resource Scan** and **Global Scan** panels
- how often performance and status are checked for monitored hosts

Service Monitor Thread Counts

By default, the number of Java threads allocated to service and performance monitors is 100. This can be modified with the following uptime.conf parameter:

```
serviceThreads = 100
```

Status Thresholds

The **Global Scan** threshold settings determine when a cell in the **Global Scan** panel changes state to reflect a host's status change: green represents normal status, yellow represents Warning status, and red represents Critical.

The Resource Scan threshold settings determine the size of the gauge ranges on the **Resource Scan** view: green represents normal status, yellow represents Warning status, and red represents Critical status.

You can change the thresholds used to determine status by manually inputting settings in the **up.time Configuration** panel, as outlined in "Modifying up.time Config Panel Settings" on page 291.

Changes to Global Scan thresholds are not retroactively applied to all Elements; only Elements added after threshold changes will reflect those changes.

Changing Global Scan Threshold Settings

You can modify the **Global Scan** threshold settings through the following parameters (default values are shown):

- `globalscan.cpu.warn=70`

  A Warning-level status is reported when CPU usage is at 70% or greater.

- `globalscan.cpu.crit=90`

  A Critical-level status is reported when CPU usage is at 90% or greater.

- `globalscan.diskbusy.warn=70`

  A Warning-level status is reported when a disk on the host is busy for 70% or more of a five-minute time frame.

- `globalscan.diskbusy.crit=90`

  A Critical-level status is reported when a disk on the host is busy for 90% or more of a five-minute time frame.

- `globalscan.diskfull.warn=70`

  A Warning-level status is reported when 70% or more of the disk space on the host is used.

- `globalscan.diskfull.crit=90`

  A Critical-level status is reported when 90% or more of the disk space on the host is used.

- `globalscan.swap.warn=70`

  A Warning-level status is reported when 70% or more of the swap space on a disk is in use.

- `globalscan.swap.crit=90`

  A Critical-level status is reported when 90% or more of the swap space on a disk is in use.

Resource Scan Threshold Settings

You can modify the **Resource Scan** threshold settings through the following parameters (default values are shown):

- `resourcescan.cpu.warn=70`

  The Warning-level range in the **CPU Usage** gauge begins at this value (70%), and ends at the Critical-level range.

- `resourcescan.cpu.crit=90`

  The Critical-level range in the **CPU Usage** gauge is between this value (90%) and 100%.

- `resourcescan.memory.warn=70`

  The Warning-level range in the **Memory Usage** gauge begins at this value (70%), and ends at the Critical-level range.

- `resourcescan.memory.crit=90`

  The Critical-level range in the **Memory Usage** gauge is between this value (70%) and 100%.

- `resourcescan.diskbusy.warn=70`

  The Warning-level range in the **Disk Busy** gauge begins at this value (70%), and ends at the Critical-level range.

- `resourcescan.diskbusy.crit=90`

  The Critical-level range in the **Disk Busy** gauge is between this value (70%) and 100%.

- `resourcescan.diskcapacity.warn=70`

  The Warning-level range in the **Disk Capacity** gauge begins at this value (70%), and ends at the Critical-level range.

- `resourcescan.diskcapacity.warn=90`

  The Critical-level range in the **Disk Capacity** gauge is between this value (70%) and 100%.

Report Storage Options

When an up.time user generates a report, that report is stored in the `/GUI/reportcache` directory; when a scheduled report is automatically generated and published, it is stored in the `/GUI/published` directory. Both of these directory paths are found in the up.time installation directory:

- Linux: `/usr/local/uptime/`
- Solaris: `/opt/uptime/`
- Windows: C:\Program Files\uptime software\uptime

Windows Vista users can find the audit log in the Virtual Store instead of the default location (i.e., C:\Users\uptime\AppData\Local\VirtualStore\Program Files\<uptime-install-directory>

By default, generated reports are cached on the Enterprise Monitoring Station for 30 days; additionally, the location for published reports is also on the local Enterprise Monitoring Station file system. Both options can be modified. In the latter case, automatically publishing reports to a publicly accessed directory on the network is an ideal way for non-IT staff to view them. See “Saving Reports to the File System” on page 166 for more information.

Changing the Number of Days Reports Are Cached

You can change a report’s expiry time limit by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 291.

Change the expiry limit through the following parameter (the default value is shown):

```
reportCacheExpiryDays=30
```

Changing the Published Report Location

This can be modified with the following `uptime.conf` parameter:

```
publishedReportRoot=<location>
```

If the intended published report directory is on a system other than the Enteprise Monitoring Station, the provided location should be a full network path to the system in addition to the directory path on that system.

Resource Usage Report Generation

Due to the large number of options available for the Resource Usage report, generating an extensive report for a large group of Elements can take several minutes. If exhaustive report generation is necessary, but taking too long, you can increase the number of report images (the default being "6") that up.time concurrently generates for this type of report.

Note that the default number is optimal in most cases; increasing the amount may improve performance, but the law of diminishing returns applies, as too many concurrent threads can tax the PDF generation process overall.

Logging is configured through the following `uptime.conf` parameter:

```
reporting.prefetch.images.threads = 6
```

Monitoring Station Interface Changes

Some configuration options affect the Enterprise Monitoring Station interface. These can be modified by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 291.

Status Alert Acknowledgement

When services reach a warning or critical state, administrators can flag an alert as “acknowledged,” which prevents subsequent alerts from being broadcasted, giving them time to investigate the issue. See “Acknowledging Alerts” on page 73 for more information.

Service status alert acknowledgements can be reported in the status tables on the **Global Scan** panel. By default, status alert acknowledgement counts are not shown; if enabled a new column (labelled ACK) appears in the **Service Status** section of **Global Scan**. When the current status of a monitor is acknowledged, it appears in the ACK column instead of in the WARN or CRIT column.

You can enable or disable status acknowledgement (i.e., add or remove the ACK column from the status tables) through the following parameter (the default value is shown):

```
acknowledgedSeparate=false
```

3D Graphs

When performance and availability graphs are generated, the Graph Editor is used to manipulate the appearance of graphed data (see “Using the Graph Editor” on page 244. Transformations from a three-dimensional perspective are possible if the user account permits it (see “Adding Users” on page 121), and the user is connecting to the EMS using Internet Explorer.

This 3D presentation option can be disabled outright. You can determine whether ActiveX graphs are displayed in 3D for users with Internet Explorer through the following parameter (the default value is shown):

```
default3DGraphs=true
```

Custom Dashboard Tabs

Custom dashboards can be added to **My Portal** to display custom content that is relevant to the particular user who is currently logged in. Up to 50 dashboards can be added, each of which is accessed through, and viewed in, its own tab at the top of **My Portal**.

A custom dashboard tab is configured by pointing up.time to a custom Web page, and indicating which User Group will be able to view it. You can enable and configure the first dashboard through the following parameters:

```
myportal.custom.tab1.enabled=true
myportal.custom.tab1.name=<DashboardNameOnTab>
myportal.custom.tab1.URL=<URLtoCustomPage>
myportal.custom.tab1.usergroups=<UserGroupName>
```

Values for the first three parameters are required. If no name is specified for the User Group parameter (or, if no User Groups have been defined), the custom dashboard will be visible to all up.time users. Thus, a User Group parameter is only required if you want to restrict or refine user access to a particular custom dashboard.

To create additional tabs, add the same set of parameters, but increment the tab count:

```
myportal.custom.tab2.enabled=true
myportal.custom.tab2.name=<DashboardNameOnTab>
myportal.custom.tab2.URL=<URLtoCustomPage>
```

License Information

If your up.time package did not come with a license key, then either contact your sales representative to request a key or send an email to support@uptimesoftware.com. You will need the host ID for the system so that a permanent license key can be generated. The host ID is displayed in the **License Information** subpanel, and is similar to the following:

```
001110bf101d
```

You do not need the host ID if you are evaluating up.time. The demo licenses expire after predetermined amounts of time and can run on any system.

In addition to your up.time license, the **License Info** sub panel displays the number of individual licenses that are currently being used in your environment. This number is broken down by systems, nodes, and (if applicable) VMware ESX processors.

To install or update a license, do the following:

1. **In the Tree panel, click License Info.**

   If you currently have an up.time license, it is displayed in the **License Information** subpanel.

2. **Paste the new or updated license into the License Key text box.**

3. **Click Update.**

up.time

APPENDIX A

Reference

This appendix contains the following sections:

Frequency Definitions

To define synchronization frequencies in up.time, you input a string that represents the amount of time between actions. These units of time can be days, hours, minutes, seconds, or a combination. Frequency definitions are used when configuring user detail synchronization, when configuring up.time to use an Active Directory or LDAP listing for user authentication and management. (See “Changing How Users Are Authenticated” on page 133 for more information.) Frequency definitions are also used when adding a Replication Group to a Datacenter, and you are configuring how frequently the Local Datacenter's data will be replicated on the Enterprise Monitoring Server. (See “Adding Replication Groups to a Datacenter” on page 56for more information.)

All time units are represented by a one-letter abbreviation:

- days: d
- hours: h
- minutes: m
- seconds: s

Frequency definitions can be a combination of any of these time units and their values, in descending order, without spaces:

- 1d
- 1d12h
- 1h30m
- 30s

Time Period Definitions

When defining new, or editing existing, Maintenance Profiles and Monitoring Periods, you need to use precise definitions that up.time can correctly interpret. Time period definitions use a controlled vocabulary that allow you to precisely define, combine, and exclude time periods.

Although all examples listed in the following sections are written in mixed case (e.g., “Every Oct 28”), none of the terms used in time period definitions is case sensitive.

Building Blocks

The following tables outline the basic components of all time period definitions.

Time Units

- Units of time that act as building blocks in definitions include times of day, days of the week, months, years, and exact dates.

Times		
Required	• hour of day • 12-hour clock suffix, inputted as “AM” or “PM”	correct: `8:00 PM`
Optional	• minutes of the hour • spaces	correct: `8 PM`, `8:00PM`, `8PM`

Not Accepted	• missing 12-hour clock suffix
	• 24-hour clock convention
		incorrect:
		`8:00`
		`20:00`, `20:00 PM`

Days		

Required	three-letter abbreviation	correct:
		`Sun`, `Mon`, `Tue`
		`Wed`, `Thu`
		`Fri`, `Sat`

Not Accepted	• full spellings
	• other abbreviation styles
		incorrect:
		`S`, `M`, `T`
		`We`, `Th`
		`Friday`, `Saturday`

Dates		
Required	single- or two-digit number	correct: `8`, `09`, `10`
Not Accepted	• ordinal suffixes • full spellings	incorrect: `8th`, `9th`, `tenth`
Months		

Required	three-letter abbreviation	correct:
		`Jan`, `Feb`, `Mar`, `Apr`
		`May`, `Jun`, `Jul`, `Aug`
		`Sep`, `Oct`, `Nov`, `Dec`

Not Accepted	other abbreviation styles	incorrect:
		`J`, `F`, `M`, `A`
		`June`, `July`, `August`
		`Se`, `Oc`, `No`, `De`

Years		
Required	full year	correct: `2008`
Not Accepted	any abbreviation of the year	incorrect: `08`, `‘08`, `Y2K+8`

Lists and Ranges

Days can be inputted as a list:

- each day is separated by a comma (e.g., "`mon, tue, wed`")
- spaces are optional (e.g., "`mon,tue,wed`")

Times and days can be inputted as ranges:

- Elements in the range must be separated by hyphens
- spaces are optional; the following examples are correct:
  - `8AM-8PM`
  - `8:00 AM - 8:00 PM`
  - `Fri - Mon`
  - `Fri-Mon`
- ranges wrap around day and week boundaries:
  - "`10PM - 2AM`" is interpreted as 10:00 p.m to 11:59 p.m. on one day, and 12:00 a.m. to 2:00 a.m. the following calendar day
  - "`Fri-Mon`" is interpreted as Friday through Saturday on one week, then Sunday through Monday the following week
- up.time converts day ranges to lists (e.g., "`Fri-Mon`" becomes "`Fri, Sat, Sun, Mon`")
- day ranges and lists can be mixed; the following examples are correct:
  - `Fri - Sun, Mon`
  - `Fri-Sun,Mon`

Basic Expressions

Using the building blocks outlined in the previous section, use the following templates to create basic expressions that are used to define time periods in up.time. Note that shaded components of a template are optional.

Fixed Dates

<month>	<date>	,	<year>	<time range>

Basic example:

`Oct 28, 2008`

Spaces are optional:

Oct28,2008

Time ranges are optional:

Oct 28, 2008 7 PM - 11 PM

Oct28,20087PM-11PM

Note: Fixed dates that do not include a time range are interpreted to include the entire day (i.e., 12:00 a.m. through 11:59 p.m.), although this will not automatically appear in the defined time period.

Fixed Date Ranges

from	<month>	<date>	<year>	<time range>
to	<month>	<date>	<year>	<time range>

Basic example:
`From Oct 28, 2008 to Oct 29, 2008`

Spaces are optional:
`FromOct28,2008toOct29,2008`

Time ranges are optional:
`From Oct 28, 2008 7 PM to Oct 29, 2008 2 AM`

Note: A fixed date without a time that is at the end of a date range is interpreted to include the first minute of the next day (e.g., up.time converts "`From Oct 28, 2008 to Oct 29, 2008`" into "`From Oct 28, 2008 12:00AM to Oct 30, 2008 12:00AM`").

Note: The time range in a fixed date range merely acts as a more precise start point and end point; a fixed date range is a contiguous block of time that has no gaps.

Weekly Recurrence

every	<day> / <day range / list>	<time range>

Basic example:

```
Sun
Sun - Tue
Every Sun, Mon, Tue
```

Spaces are optional:

```
Sun-Tue
EverySun,Mon,Tue
```

Time ranges are optional:

```
Sun 9 AM - 5 PM
Sun - Tue 9AM - 5PM
EverySun,Mon,Tue9AM-5PM
```

Note: Recurring days that do not include a time range are interpreted to include the entire day (i.e., 12:00 a.m. through 11:59 p.m.), although this will not automatically appear in the defined time period.

Yearly Recurrence

every	<month>	<date>	<time range>

Basic example:

```
Every Oct 28
```

Ordinal suffixes are optional:

```
Every Oct 28th
```

Time ranges are optional:

Every Oct 28 7PM - 11PM

Note: You cannot define a date range within a yearly recurrence; instead, combine a separate yearly recurrences for each date in the date range.

Monthly Recurrence

every month on the	<date>	<time range>

Basic example:

```
Every month on the 28
```

Ordinal suffixes are optional:

```
Every month on the 28th
```

Time ranges are optional:

```
Every month on the 28 6 PM - 11 PM
Every month on the 28th 6PM-11PM
```

Monthly Ordinal Recurrence

every month on the	<ordinal_as_word>	<day>	<time range>

Basic example:

```
Every month on the last Fri
```

Time ranges are optional:

```
Every month on the last Fri 6 PM - 11 PM
Every month on the last Fri 6PM-11PM
```

Note: The ordinal must be stated as a word: `first`, `second`, `third`, `fourth`, and `last`.

Combining Expressions and Excluding Time Periods

Elaborate time period defintions are built from a combination of the basic expressions defined in the previous section:

- fixed dates
- fixed date ranges

- weekly recurrences
- monthly recurrences
- monthly ordinal recurrences
- yearly recurrences

Combinations

Combine basic expressions by writing each one on a new line in the **Definition** box when defining a Maintenance Profile or Monitoring Period. The following examples demonstrate combinations of different basic expressions used to define a maintenance window:

Combining fixed dates:

```
Dec 25, 2008 12AM - 12PM
Jan 1, 2009 12AM - 12PM
```

Combining a fixed date and a fixed date range:

```
Dec 25, 2008 12AM - 12PM
From Dec 31, 2008 11PM to Jan 1, 2009 12PM
```

Combining weekly recurrences:

```
Mon-Fri 1AM-3AM
Sat 1AM-5:30AM
Sun
```

Combining yearly recurrences:

```
Every Dec 25 12AM-12PM
Every Dec 31 11PM-11:59PM
Every Jan 1st 12AM-12PM
```

Combining monthly recurrences:

```
Every month on the 2
Every month on the 16th
```

Combining monthly ordinal recurrences:

```
Every month on the first Fri
Every month on the third Fri
Every month on the last Fri
```

Note that when a time period consists of more than one component time period expression, a condition met within *any* of those component time periods applies to the entire time period. For example, if a Monitoring Period named “Open Hours” is defined as:

```
Mon-Fri 9AM-5PM
Sat 10AM-5PM
Sun 12PM-5PM
```

An alert-worthy event that occurs on Sunday at 1:00 p.m. means the entire time period definition has been fulfilled.

Exclusions

Time periods can be excluded from greater time period definitions by using the term “`exclude`” as a prefix to the exclusionary expression. The following examples demonstrate the use of exclusions in time periods:

Excluding a monthly recurrence from a weekly recurrence:

```
Sun 3PM-5PM
Exclude every month on the last Sunday
```

Defining two yearly recurrences to exclude from a weekly recurrence:

```
Mon-Fri 2AM-3AM
Exclude every Jan 1
Exclude every Jan 2
```

up.time

APPENDIX B

End User License Agreement

Before downloading up.time, obtaining a license key, or using up.time, please read the following End User License Agreement for up.time. The up.time End User License Agreement defines the rights, permissions, and limitations that you agree to by choosing up.time.

The up.time End User License Agreement is detailed in the following sections:

NOTICE TO USER

This End User License Agreement (the “Agreement”) is a legal contract between you, as either an individual or a business entity, and Uptime Software Inc. (“Uptime”).

PLEASE READ THIS CONTRACT CAREFULLY BEFORE DOWNLOADING UPTIME’S PROPRIETARY SOFTWARE (the “SOFTWARE”) OR OBTAINING A LICENSE KEY TO THE SOFTWARE OR USING THE SOFTWARE. BY CLICKING ON THE “I ACCEPT” BUTTON AND BY DOWNLOADING THE SOFTWARE OR OBTAINING A LICENSE KEY TO THE SOFTWARE YOU REPRESENT AND WARRANT THAT YOU ARE EITHER THE REPRESENTATIVE OF THE COMPANY WITH THE AUTHORITY TO ENTER INTO THIS AGREEMENT AND TO BIND THE COMPANY OR YOU ARE AN INDIVIDUAL OVER THE AGE OF 18 AND THAT YOU HAVE READ THIS AGREEMENT, THAT YOU UNDERSTAND IT, AND THAT YOU ACCEPT AND AGREE TO BE BOUND BY ITS TERMS. IF YOU ARE UNWILLING TO BE BOUND BY THE TERMS OF THIS AGREEMENT YOU SHOULD CLICK THE “I DO NOT ACCEPT” BUTTON BELOW, TERMINATE THE DOWNLOAD PROCESS AND REFRAIN FROM ACCESSING OR USING THE SOFTWARE. THIS AGREEMENT REPRESENTS THE ENTIRE AGREEMENT BETWEEN YOU AND UPTIME CONCERNING THE SOFTWARE AND THIS AGREEMENT SUPERSEDES AND REPLACES ANY PRIOR PROPOSAL, REPRESENTATION, COMMUNICATION, ADVERTISEMENT OR UNDERSTANDING YOU MAY HAVE HAD WITH UPTIME RELATING TO THE SOFTWARE.

1. License

1.1 Grant of License.

Uptime hereby grants to you and you accept, a limited, non-exclusive license to use the Software in machine-readable, object code form only and the user manuals accompanying the Software (the “Documentation”), only as authorized in this Agreement. For purposes of this Agreement, the “Software” includes any updates, enhancements, modifications, revisions or additions to the Software made by Uptime and made available to end

users through Uptime's web site. Notwithstanding the foregoing, Uptime shall be under no obligation to provide any updates, enhancements, modifications, revisions or additions to the Software.

1.2 Scope of Use

You may use the Software activated by a license key on a single server designated by you as the monitoring station. If you have multiple license keys for the Software each key will be activated on a designated server. For purposes of this Agreement, "use" of the Software means loading the Software into the temporary or permanent memory of a computer. The Software may not be used on or distributed to a greater number of servers than you have license keys. There is no restriction on the number of users who may access the designated servers and use the Software.

1.3 Copies and Modifications

You may not reverse engineer, decompile, disassemble or otherwise translate the Software or attempt to derive the source code of the Software or any license keys you have obtained. You may not modify or adapt the Software or any license keys that you have obtained in any way. You may make one (1) copy of the Software, the Documentation and any license keys that you have obtained, solely for backup or archival purposes. Any such copies of the Software, Documentation or license keys shall include any copyright or other proprietary notices that were included on such materials when you first received them. Except as authorized in this Section 1.3, no copies of the Software, Documentation or license keys, or any part thereof, may be made by you or any person under your authority or control.

1.4 Assignment of Rights

The license granted under this Agreement is personal to you. You are not permitted to grant access to, distribute, sell, transfer, publish, disclose, display, sublicense, lease, rent or lend your rights in the Software, Documentation or license keys as granted by this Agreement for any purpose or in any manner.

1.5 Licenses Required for Third-party Software

The Software enables you to monitor multiple instances of third-party operating systems and application programs. You are responsible for obtaining and complying with any licenses necessary to operate any such third-party software, including Operating Systems and/or application programs.

2. Intellectual Property and Confidentiality

2.1 Use Reporting, License Violations and Remedies

Uptime reserves the right to gather data on key usage including license key numbers, server IP addresses, domain counts and other information deemed relevant to ensure that its products are being used in accordance with the terms of this Agreement. Uptime expressly prohibits simultaneous, multiple installations of its licensed products and domain count overrides without its prior written approval. Any unauthorized use shall be considered by Uptime to be a violation of this Agreement. Uptime reserves the right to remedy violations immediately upon discovery, by charging the then-current list price of unauthorized keys to the end user or by any other means necessary. You agree not to block, electronically or otherwise, the transmission of data required for compliance with this Agreement. Any blocking of data required for compliance under this Agreement is considered to be violation of this Agreement and will result in immediate termination of this Agreement pursuant to Section 4.

2.2 License Automatic Update and Expiration

Your license may include an expiration date that can result in the termination of the license. For permanent license keys, the license updates will be available to you upon payment of the appropriate, then-current Uptime license fees. You must contact Uptime to take the appropriate steps to obtain the permanent key. If your license key is stolen or if you suspect any improper or illegal usage of your license outside of your control you should promptly notify Uptime of such occurrence. A replacement license will be issued to you and the suspect license will be allowed to expire. For your convenience Uptime provides license expiration warnings in the product interface should there be any issues that would cause the license to

expire. It is your responsibility to contact Uptime regarding any potential expiration. Uptime is not liable for any damages or costs incurred in connection with an expiring license.

2.3 Proprietary Rights to Software and Trade Marks

You acknowledge that the Software and the Documentation are proprietary to Uptime and the Software and Documentation are protected under Canadian copyright law and international treaties. You further acknowledge and agree that, as between you and Uptime, Uptime owns and shall continue to own all right, title and interest in and to the Software and Documentation including associated intellectual property rights under copyright, trade secret, patent or trade mark laws. This Agreement does not grant you any ownership interest in or to the Software or the Documentation but only a limited right of use that is revocable in accordance with the terms of this Agreement. Any and all trade marks or service marks that Uptime uses in connection with the Software or with services rendered by Uptime are marks owned by Uptime. This Agreement does not grant you any right, license or interest in such marks and you shall not assert any right, license or interest in such marks or any words or designs that are confusingly similar to such marks.

2.4 Confidentiality

You shall permit only authorized users who possess rightfully obtained license keys to use the Software or to view the Documentation. Except as expressly authorized by this Agreement you shall not make the Software, Documentation or any license key available to any third party. You will use your best efforts to co-operate with and assist Uptime in identifying and preventing any unauthorized use, copying or disclosure of the Software, Documentation or any part thereof.

3. License Fees

The Software will be available to you for use upon your receipt of one or more license keys. Upon acceptance of this Agreement you may obtain one or more temporary license keys and permanent license keys using the procedure set forth on Uptime's web site including, but not limited to, payment of Uptime's license fees. The license fees paid by you are paid in

consideration of the license granted under this Agreement. Uptime does not refund license fees. By accepting this Agreement you fully understand that once license fee payment is made to Uptime you will have no recourse for receiving a refund of any part of the fees.

4. Term and Termination

This Agreement is effective upon your acceptance of the Agreement or upon your downloading, accessing and using the Software, even if you have not expressly accepted this Agreement. This Agreement shall continue in effect until terminated. Without prejudice to any other rights, this Agreement will terminate automatically if you fail to comply with any of the limitations or other requirements described herein. If you have a temporary key and fail to pay the applicable license fees for continuation of use the key will expire. You may terminate this License Agreement at any time by: (i) providing written notice of your decision to terminate the Agreement to Uptime; and, (ii) either returning the Software, Documentation, all copies thereof and all license keys that you have obtained to Uptime or destroying all such materials and providing written verification of such destruction to Uptime. Uptime reserves the right to physically verify that the Software has been removed. Uptime may terminate this License Agreement if you breach any term of the Agreement by giving you written notice of your breach and Uptime's decision to terminate the Agreement. Upon termination by Uptime you agree to either return the Software, Documentation and all copies thereof and all license keys that you have obtained to Uptime or to destroy all such materials and provide written verification of such destruction to Uptime.

5. Remedies and Indemnification

5.1

If you learn of any actual or threatened infringement or piracy of the Software or, if any infringement or piracy claim is made against you by a third party in connection with your use of the Software, you shall notify Uptime in writing of the infringement, piracy or claim as soon as is reasonably possible. Uptime shall, in its sole discretion, determine what action, if any, to take with respect to the foregoing and shall assume the

defense or bear the expenses of any such action (except to the extent, if any, to which such dispute or costs arise from your negligence, willful misconduct or modification of the Software). In the event that the use of the Software in accordance with the provisions of this Agreement is declared by a court of competent jurisdiction to infringe the rights of any third party, as your sole remedy, Uptime, at its option may: (i) procure for you the right to use the Software; or, (ii) modify the Software to render it non-infringing.

5.2

You will, at your expense, indemnify and hold Uptime and all its officers, directors and employees, harmless from and against any and all claims, actions, liabilities, losses, damages, judgments, grants, costs and expenses, including reasonable lawyer fees (collectively “Claims”) arising out of any use of the Software by you, any party related to you or any party acting upon your authorization in a manner that is not expressly authorized by this Agreement.

6. Disclaimer

THE SOFTWARE, DOCUMENTATION AND ANY (IF ANY) SUPPORT SERVICES ARE LICENSED “AS IS” AND UPTIME AND ITS SUPPLIERS DISCLAIM ANY AND ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. WITHOUT LIMITING THE GENERALITY OF THE FOREGOING, UPTIME EXPRESSLY DOES NOT WARRANT THAT THE SOFTWARE WILL MEET YOUR REQUIREMENTS OR THAT OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR FREE. YOU ASSUME RESPONSIBILITY FOR SELECTING THE SOFTWARE TO ACHIEVE YOUR INTENDED RESULTS AND FOR THE RESULTS OBTAINED FROM YOUR USE OF THE SOFTWARE. YOU SHALL BEAR THE ENTIRE RISK AS TO THE QUALITY AND THE PERFORMANCE OF THE SOFTWARE.

7. Limitation of Liability

UPTIME’S CUMULATIVE LIABILITY TO YOU OR ANY PARTY RELATED TO YOU FOR ANY LOSS OR DAMAGES RESULTING

FROM ANY CLAIMS, DEMANDS OR ACTIONS ARISING OUT OF OR RELATING TO THIS AGREEMENT INCLUDING, WITHOUT LIMITATION, UPTIME'S INTELLECTUAL PROPERTY INDEMNIFICATION OBLIGATIONS SHALL BE LIMITED TO THE AMOUNT OF LICENSE FEES PAID TO UPTIME BY YOU UNDER THIS AGREEMENT. BUT, IN NO EVENT SHALL SUCH LIABILITY EXCEED CDN. $2,000.00 IN THE AGGREGATE FOR ALL OCCURRENCES. THIS LIMITATION APPLIES TO ALL CAUSES OF ACTION OR CLAIMS IN THE AGGREGATE, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, INDEMNITY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS. IN NO EVENT SHALL UPTIME BE LIABLE TO YOU OR ANY PARTY RELATED TO YOU FOR ANY INDIRECT, INCIDENTAL, CONSEQUENTIAL, SPECIAL, EXEMPLARY OR PUNITIVE DAMAGES OR LOST PROFITS EVEN IF UPTIME HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE FOREGOING LIMITATIONS, EXCLUSIONS AND DISCLAIMERS SHALL APPLY TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW EVEN IF ANY REMEDY FAILS ITS ESSENTIAL PURPOSE.

8. General Terms

8.1 Governing Law and Choice of Forum

This Agreement shall be governed by and interpreted in accordance with the laws of the Province of Ontario, Canada, without regard to the conflicts of law rules thereof. Any claim or dispute arising in connection with this Agreement shall be resolved in the federal or provincial courts situated with the City of Toronto, Ontario. To the maximum extent permitted by law, you hereby consent to the jurisdiction and venue of such courts and waive any objections to the jurisdiction or venue of such courts. To the extent any terms and conditions on a purchase order or other ordering document submitted to Uptime by you conflicts with the terms of this Agreement, the terms of this Agreement shall control and notwithstanding any term of your order which states to the contrary.

8.2 Severability

If any term or provision of this Agreement is declared void or unenforceable in a particular situation by any judicial or administrative authority this declaration shall not affect the validity or the enforceability of the remaining terms and provisions hereof or the validity or enforceability of the offending term or provision in any other situation.

8.3 Survival

Sections 2, 5, 6, 7 and 8 of this Agreement and all subsections thereof shall survive the termination of this Agreement regardless of the cause for termination and shall remain valid and binding indefinitely.

8.4 Headings

The Article and Section headings contained in this Agreement are incorporated for reference purposes only and shall not affect the meaning or interpretation of this Agreement.

8.5 No Waiver

The failure of either party to enforce any rights granted hereunder or to take action against the other party in the event of any breach hereunder shall not be deemed a waiver by that party as to subsequent enforcement of rights or subsequent actions in the event of future breaches.

8.6 Amendment

Uptime reserves the right, in its sole discretion, to amend this Agreement from time to time. If there is a conflict between this Agreement and the most-current version of this Agreement posted at www.uptimesoftware.com, the most-current version will prevail. If you do not accept amendments made to this Agreement then this license will be immediately terminated pursuant to Section 4.

8.7 Taxes

You shall, in addition to the license fees required under this Agreement, pay all applicable sales, use, transfer or other taxes and all duties, whether national, provincial or local, however, designated, that are levied or imposed by reason of the transaction contemplated under this Agreement excluding income taxes on the net profits of Uptime. You shall reimburse Uptime for the amount of any such taxes or duties paid or incurred directly by Uptime as a result of this transaction.

Index

L

M

N

O

P

R

S

T

U

V

W