

Understanding up.time

Understanding the up.time Interface

up.time Tool Bar

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

Global Scan

The *Global Scan* panel provides information about the status of your resources. You can drill down by system group, system, or alert status to manage the resources in your infrastructure.

For more information about using the *Global Scan* panel, see [Overseeing Your Infrastructure](#).

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 [My Portal My Infrastructure](#)

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

- Systems
- Groups
- Applications
- Service Level Agreements
- Views

For more information about using the *My Infrastructure* panel, see [Managing Your Infrastructure](#).

Services

The *Services* panel enables you to manage and configure services, which are provided by an application to perform a specific task. *up.time* monitors both services and applications to ensure that performance and availability are maintained.

In the *Services* panel, you can manage and configure the following:

- service instances and service groups
- Alert Profiles and Action Profiles
- host checks
- topological dependencies
- scheduled maintenance

For more information about using the *Services* panel, see [Using Service Monitors](#).

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 [User Management](#).

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](#).

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 some *up.time* system configuration options from the *Config* panel. For more information about using the *Config* panel, see [Configuring and Managing up.time](#).
Search

As you enter text in the *Search up.time* field, *up.time* configuration actions, as well as the display names of Elements, are displayed.

For example, initially entering the string "serv" will display the "Add Service Monitor" and "Add Service Group" commands. The string could also display Element names such as "QA Server 1" and "Active Directory". In the latter example, the match would occur if the Element's host name were "AD-Server".

For all Elements, the string entered is compared with the Element's host name, display name, architecture, and custom fields.

System List

The system list (Syslist) is a popup window that contains the following information:

- the display names in *up.time* and the host names of systems in your environment, arranged in alphabetical order
- the name of the group to which, if any, the system belongs

You access the system list by clicking the *Syslist* icon in the top-right corner of the *up.time* Web interface. The *Syslist* is also a tool for quick navigation within the *up.time* Web interface. Each display name is a hyperlink. Click a display name to view the information about the system in the *System Information* subpanel.

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.

For more information, see [Understanding Reports Options](#) and [Understanding Graphing](#).

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](#).

If you need to regularly run certain reports, you can save them to the *My Portal* panel. See [Scheduling Reports](#) 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](#) and [Using Graphs](#)

Understanding Agents

Agents are small applications that are installed on the systems that you are monitoring. Agents do the following:

- collect information from a remote server
- send the collected service data to the Monitoring Station

Certain *up.time* monitors poll the agents for data at a frequency that you can configure. The data collector component of the Monitoring Station then stores the results in the *up.time* DataStore for use in a report or graph.

Agents enable you to collect very detailed information about a system, such as information about processes and low-level system statistics. The level of granularity of the information collected by agents is greater than that of the information collected by agentless monitors.

Each *up.time* agent is configured by default to collect and return performance information for every *up.time* agent service monitor. You do not need to configure the agent to collect information for a service.

On Windows, an agent is installed with the *up.time* Monitoring Station. However, you will need to deploy the agent on the systems you are monitoring. On other operating systems, you must download the agent from the uptime software Web site and manually install it.

Understanding Major and Minor Versions

When you install *up.time*, you install a Monitoring Station and one or more *up.time* agents. You could have different versions of Monitoring Stations and agents. For example, you could have different platforms and different *up.time* agent versions running on each system.

- Major version

Regardless of operating system platform, the major version is the number to the left of the decimal. In the diagram above the major number of the Windows agent is 3; the major number of the UNIX agent is 3; the major number of the LINUX agent is 4.0.

- Minor version

Minor version numbers follow the major version number. These numbers are used to distinguish each minor version of a major version.

On UNIX and Linux, the minor version is the first number to the right of the decimal. In the diagram above, the minor version number of the UNIX agent is 8 and the minor version number of the Linux agent is 0.

On Windows, the minor version is the last set of numbers in the complete version. In the diagram above, the minor version number of the Windows agent is 1061.

For major version 4 and later for Windows, the minor version number is the number immediately after the decimal that follows the major number. For example, for Windows agent version 4.0, the minor number is 0.

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 and system information for the hosts that you are monitoring
- the performance data gathered by monitors, which is used for generating graphs and reports
- user information, including user names and passwords (encrypted if it is sensitive information)
- the settings for service monitors, Alert and Action Profiles, scheduled maintenance, and host checks
- reports that 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. For more information about these accounts, see the uptime software Knowledge Base article "Securing MySQL Database and Adding Users".

up.time can also use either an Oracle or MS SQL Server database as its DataStore.

Connecting to the DataStore Using ODBC

You can extract data from the DataStore for use in custom reporting or data warehousing by connecting to the DataStore using an ODBC connection. Once the connection is established, you can import the contents of the DataStore into such tools as MySQL Query Browser, Microsoft Excel and Crystal Reports.

Before you can connect to the DataStore using ODBC, the client system that is accessing the database must have the MySQL ODBC driver installed. The ODBC driver enables the client system to communicate with the DataStore.

For detailed information on installing and configuring the MySQL ODBC driver, see the uptime software Knowledge Base article "Connecting to the *up.time* DataStore via ODBC".

Understanding Service Monitors

up.time service monitors ensure the performance and availability of services in your environment. Using service monitors, you can ensure that the systems in your environment - including databases, mail servers, networking protocols, and file systems - are operating as required. *up.time* also captures performance metrics collected from hardware profiles of physical systems in your environment and can present this data in a graph.

up.time can track the performance of services using over 30 monitors. As well, *up.time* enables you to configure custom monitors that you can use to extend your service monitoring capability.

For detailed information on service monitors, see [Using Service Monitors](#).

Understanding Database Monitors

There are two types of monitors for MySQL, Oracle, and SQL Server databases:

- Basic Checks

These monitors determine whether or not the database is running and listening on the expected port. You can also run queries against the databases using scripts.

- Advanced Metrics

These monitors collect detailed information about database processes, which you can later use for reporting and graphing.

Understanding Agentless Monitors Using Net-SNMP

Net-SNMP suite of command line and graphical applications that interact with SNMP agents that are installed on hosts. Net-SNMP presents a set of SNMP MIBs (Management Information Base, which is a listing that defines variables needed by the SNMP protocol to monitor and control network equipment). The MIBs are used to collect system performance information for use by the *up.time* Monitoring Station.

The Net-SNMP monitor uses the *HOST-RESOURCES* MIB to collect the following data:

- Configuration
- System name.
- Number of CPUs.

- The size of the system memory.
- The network interfaces on the system, as well as their MTU, speed, and physical address.

Note - The HOST-RESOURCES MIB can collect other configuration data, but the Monitoring Station does not use this information.

- *Performance Data*
- CPU
- CPU user time
- CPU system time
- CPU wait I/O time
- Memory
- the amount of free memory
- the amount of free swap space
- Processes
- the name of a process
- the ID of a process (PID)
- the amount of memory used by a process
- process run time (in centi-seconds on the CPU)
- the number of running processes
- Network
- the name of the network interface
- the number of kilobytes flowing into the interface per second
- the number of kilobytes flowing out of the interface per second
- the number of inbound errors
- the number of outbound errors
- File System
- the name of the file system
- the size of the file system
- the amount of the file system that is being used
- User
- the number of users who are logged into the system

For more information on SNMP and Net-SNMP, see .

Understanding Services

Services are specific tasks, or sets of tasks, performed by an application in your environment. For example, network services such as FTP or TCP transmit data in a network. Database services, such as Oracle, SQL Server, MySQL or Sybase store and retrieve data in a database. *up.time* service monitors continually check the condition of services to ensure that they are providing the functions required to support your business.

up.time service monitors use a common template to ensure that the configuration of service monitors is the same across all monitors. For more information on services, see [Using Service Monitors](#).

Understanding Service Groups

Service groups are service monitor templates that enable you to simultaneously apply a common service check to one or more hosts. Defining and using service groups will greatly simplify the task of initially setting up and maintaining common service checks that you wish to perform across many hosts in an identical manner.

For example, you can create a service group called CPU Performance Check that is associated with 50 different servers. You can apply a common performance monitor check to 50 servers.

With service groups, you save time by not having to manually re-create an individual service monitor with the exact same service check and Alert Profile for each server you want to monitor. There is no practical limit to the number or complexity of your service groups and the underlying service monitors associated with them.

Service groups can be created for both physical infrastructure assets that are monitored by *up.time* , as well as virtual assets managed by VMware vSphere. Although they are functionally identical, vSphere service groups are automatically applied to newfound ESX hosts, VMs, and other VMware vCenter objects that are discovered through the vSync process.

See [VMware vSphere Monitoring Concepts](#) for information on vSync, and [Service Groups](#) for information on creating service groups.

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. If a service is in a warning or critical state, you can acknowledge an alert so that *up.time* does not generate subsequent notifications.

The status of the services associated with a system are displayed in the *Global Scan* panel.

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](#) 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* subpanel

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

Note - 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

Note - 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](#)) or a Service Metrics graph.

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.