



up.time User Guide version 4

Copyright © 2007 uptime software inc.

uptime software inc. considers information included in this documentation to be proprietary. Your use of this information is subject to the terms and conditions of the applicable license agreement.

Restricted Rights Legend

This product or document is protected by copyright and distributed under licenses (see “End User License Agreement” on page 489) restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of uptime software and its licensors, if any.

Third party software is copyright and licenced from uptime software suppliers.

Documentation is provided “as is” and all express or implied conditions, representations, and warranties including any implied warranty or merchantability are disclaimed, except to the extent that such disclaimers are held to be legally invalid.

Trademarks

up.time® is a registered trademark of uptime software inc.

IBM is a registered trademark of International Business Machines Corporation.

Oracle is a registered trademark, and the Oracle product names are registered trademarks or trademarks of Oracle Corporation.

Microsoft, Windows, Microsoft SQL Server, and other such trademarks are registered trademarks of Microsoft Corporation.

Sybase, PowerBuilder, and other such trademarks are the registered trademarks of Sybase Incorporated.

Solaris, ultrasparc, and other such trademarks are the registered trademarks of Sun Microsystems Incorporated.

All other trademarks belong to their respective companies, property owners, and organizations.

Contacting uptime software

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

Main telephone line: 416 868 0152

Main fax line: 416 868 4867

Contacting Sales

To contact sales, use the main telephone line: 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 Version 4 Manual/Guide, up.time Version 4 Release Notes, up.time Version 4 README, or the help system from the Help button in the application.

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

TABLE OF CONTENTS

Welcome to up.time

<i>Introducing up.time</i>	<i>2</i>
<i>Who Should Read This Guide.....</i>	<i>2</i>
<i>up.time Architecture</i>	<i>3</i>
<i>up.time Service Monitoring Concepts</i>	<i>4</i>

Understanding up.time

<i>Understanding the up.time Interface</i>	<i>6</i>
<i>up.time Tool Bar</i>	<i>6</i>
<i>System List.....</i>	<i>9</i>
<i>Icons.....</i>	<i>10</i>
<i>System Icons.....</i>	<i>11</i>
<i>Understanding Reports and Graphs.....</i>	<i>12</i>
<i>Understanding Reports</i>	<i>12</i>
<i>Understanding Graphs.....</i>	<i>12</i>
<i>Understanding Alerts</i>	<i>13</i>
<i>Understanding the Alert Flow.....</i>	<i>14</i>
<i>Understanding Alert Profiles.....</i>	<i>15</i>
<i>Understanding Actions</i>	<i>16</i>
<i>Understanding Action Profiles.....</i>	<i>16</i>

Understanding Agents.....	17
Understanding Major and Minor Versions	17
Understanding the up.time DataStore	19
Connecting to the DataStore Using ODBC	19
Understanding Service Monitors	21
Understanding Database Monitors.....	21
Understanding Agentless Monitors Using Net-SNMP	21
Understanding Services	24
Understanding Service Groups	24
Understanding the Status of Services	25
Understanding Security in up.time	26
Understanding Dates and Times.....	27
Understanding Retained Data	29

Installing up.time

Installation Plan	32
Installation Requirements.....	33
up.time Monitoring Station.....	33
up.time Agents	34
Installing the up.time Monitoring Station.....	35
Before You Begin	35
Installing the Monitoring Station on Windows.....	36
Installing the Monitoring Station on Solaris, Linux or AIX.....	38
Post-Installation Tasks	42
Upgrading to up.time 4.....	42

Installing Agents.....	43
Installing Agents on Windows	43
Installing Agents on Solaris	44
Installing Agents on UNIX	44
Installing Agents on Linux	45
Installing Agents on IBM pSeries Servers	46

Getting Started

Starting and Stopping up.time	52
Starting the up.time Services	52
Stopping the up.time Services.....	53

Accessing and Exiting up.time.....	55
Setting Up the Administrator Account	55
Accessing up.time	56
Exiting up.time.....	56

Viewing System and Service Information.....	57
Viewing System Information.....	57
Viewing Service Information.....	59

Searching and Filtering	63
Using the Search Box	63
Filtering Service Instances	64

Audit Logging.....	66
Enabling the Audit Log.....	66

Using My Portal

Overview.....	68
Assistance	68
Current Issues.....	69
My Preferences	69

<i>Latest up.time Articles</i>	<i>69</i>
<i>up.time Information</i>	<i>69</i>
<i>My Reports</i>	<i>70</i>

Using My Enterprise

Overview	72
Working with Groups	73
<i>Adding Groups</i>	<i>73</i>
<i>Adding Nested Groups</i>	<i>74</i>
Working with Systems	76
<i>Adding Systems or Network Devices</i>	<i>78</i>
<i>Auto Discovery</i>	<i>82</i>
<i>Adding VMware Instances to up.time</i>	<i>86</i>
<i>Adding Individual LPARs to up.time</i>	<i>87</i>
<i>Editing a System Profile</i>	<i>88</i>
Working with Applications	90
<i>Adding Applications</i>	<i>90</i>
<i>Viewing Details About Applications</i>	<i>92</i>
<i>Editing Applications</i>	<i>92</i>
Working with Views	93
<i>Adding Views</i>	<i>93</i>
<i>Displaying and Editing Views</i>	<i>94</i>
Deleting Systems, Applications, and Views	95
Acknowledging Alerts	96

Using Global Scan

Overview.....	100
<i>Groups and Views in the Global Scan Panel</i>	<i>102</i>
Viewing All Applications.....	103
<i>Condensed View</i>	<i>104</i>
<i>Detailed View</i>	<i>104</i>
Viewing All Elements.....	106
Viewing All Services.....	108
Viewing a Resource Scan	109
<i>Gauges.....</i>	<i>109</i>
<i>Chart</i>	<i>109</i>

Using Service Monitors

Overview.....	112
Using Service Monitors.....	113
<i>Using Agent Monitors.....</i>	<i>113</i>
<i>Using Agentless Monitors</i>	<i>114</i>
<i>Using Advanced Monitors</i>	<i>114</i>
<i>Selecting a Monitor</i>	<i>116</i>
The Monitor Template	117
<i>Monitor Identification</i>	<i>117</i>
<i>Adding Monitor Identification Information</i>	<i>118</i>
<i>Monitor Settings Configuration.....</i>	<i>119</i>
<i>Configuring Warning and Critical Thresholds.....</i>	<i>121</i>
<i>Monitor Timing Settings</i>	<i>122</i>
<i>Monitor Alert Settings.....</i>	<i>125</i>
<i>Monitoring Period Settings</i>	<i>126</i>
<i>Getting Additional Help</i>	<i>126</i>

Cloning Service Monitors.....	127
Testing Service Monitors	128
Alert Profiles	129
Creating Alert Profiles	129
Working with Custom Alert Formats.....	130
Viewing Alert Profiles	132
Editing Alert Profiles.....	132
Applying Alert Profiles to Applications.....	133
Action Profiles.....	134
Creating Action Profiles.....	134
Viewing Action Profiles.....	137
Editing Action Profiles	138
Applying Action Profiles to Applications	138
Service Groups	139
Creating Service Groups	139
Editing Service Groups.....	140
Changing Host Checks.....	142
Change a Host Check	142
The Platform Performance Gatherer.....	143
Editing the Platform Performance Gatherer	143
Topological Depend encies.....	145
Adding Topological Dependencies	146
Viewing Topological Dependencies.....	146
Scheduling Mainte nance	147
Creating Scheduled Maintenance Profiles	147
Viewing Scheduled Maintenance Profiles	149
Scheduling Maintenance for a Host	149
Scheduling Maintenance for a Service.....	150

Agent Monitors

Overview	152
File System Capacity	153
Configuring File System Capacity Monitors	153
Performance Check	155
Configuring Performance Check Monitors	155
Process Count Check	159
Configuring Process Count Check Monitors	159

Microsoft Windows Monitors

Windows Event Log Scanner	162
Configuring Windows Event Log Scanner Monitors	162
Windows Service Check	165
Configuring Windows Service Check Monitors	165
Windows File Shares (SMB)	168
Configuring Windows File Shares (SMB) Monitors	168
Active Directory	170
Configuring Active Directory Monitors	170

Application Monitors

Uptime Agent	174
Configuring Uptime Agent Monitors	174
Exchange	176
Configuring Exchange Monitors	176

IIS	179
<i>Configuring IIS Monitors.....</i>	<i>179</i>
WebLogic	182
<i>Deploying the weblogic.jar File</i>	<i>184</i>
<i>Configuring WebLogic Monitors</i>	<i>185</i>
WebSphere	187
<i>Deploying the WebSphere Performance Servlet.....</i>	<i>190</i>
<i>Configuring WebSphere Monitors</i>	<i>191</i>
ESX v3 Workload.....	193
<i>Configuring ESX v3 Workload Monitors</i>	<i>193</i>
Splunk Query	196
<i>Configuring Splunk Query Monitors</i>	<i>196</i>
Live Splunk Listener.....	198
<i>Before You Begin</i>	<i>199</i>
<i>Configuring Live Splunk Listener Monitors.....</i>	<i>200</i>

Database Monitors

MySQL (Advanced Metrics)	202
<i>Configuring MySQL (Advanced Metrics) Monitors</i>	<i>202</i>
MySQL (Basic Checks).....	209
<i>Configuring MySQL (Basic Checks) Monitors.....</i>	<i>209</i>
Oracle (Advanced Metrics)	211
<i>Configuring Oracle (Advanced Metrics) Monitors</i>	<i>211</i>
Oracle (Basic Checks)	214
<i>Configuring Oracle (Basic Checks) Monitors</i>	<i>214</i>

Oracle Tablespace Check.....	217
Configuring Oracle Tablespace Check Monitors	217
SQL Server (Basic Checks)	220
Configuring SQL Server (Basic Checks) Monitors	220
SQL Server (Advanced Metrics).....	224
Using Multiple SQL Server (Advanced Metrics) Monitors	224
Configuring SQL Server (Advanced Metrics) Monitors	225
SQL Server Tablespace Check	228
Structure of a SQL Server Database.....	228
Configuring SQL Server Tablespace Check Monitors.....	229
Sybase	233
Configuring Sybase Monitors	233

Network Service Monitors

DNS.....	238
Before You Begin	238
Configuring DNS Monitors	239
FTP.....	241
Configuring FTP Monitors	241
HTTP (Web Services)	243
Configuring HTTP (Web Services) Monitors	243
IMAP (Email Retrieval)	247
Configuring IMAP (Email Retrieval) Monitors.....	247
LDAP	249
Before You Begin	249
Configuring LDAP Monitors.....	250

NFS	253
<i>Configuring NFS Monitors.....</i>	<i>253</i>
NIS/YP	255
<i>Configuring NIS/YP Monitors.....</i>	<i>255</i>
NNTP (Network News)	257
<i>Command Implementation</i>	<i>257</i>
<i>Response Category</i>	<i>258</i>
<i>Response Codes.....</i>	<i>258</i>
<i>Configuring NNTP (Network News) Monitors.....</i>	<i>259</i>
Ping	261
<i>Configuring Ping Monitors.....</i>	<i>261</i>
POP (Email Retrieval)	263
<i>Configuring POP (Email Retrieval) Monitors.....</i>	<i>263</i>
SSH (Secure Shell)	265
<i>Configuring SSH (Secure Shell) Monitors.....</i>	<i>265</i>
SMTP (Email Delivery)	267
<i>Configuring SMTP (Email Delivery) Monitors.....</i>	<i>267</i>
SNMP	269
<i>Net-SNMP</i>	<i>269</i>
<i>SNMP MIB Browser</i>	<i>270</i>
<i>Supported Versions of SNMP</i>	<i>270</i>
<i>Using the SNMP MIB Browser.....</i>	<i>270</i>
<i>Configuring SNMP Monitors.....</i>	<i>273</i>
TCP	277
<i>Configuring TCP Monitors.....</i>	<i>277</i>

Advanced Monitors

Overview.....	280
<i>Before You Begin</i>	<i>281</i>
Custom Monitors	282
<i>Configuring Custom Monitors.....</i>	<i>282</i>
Custom with Retained Data.....	284
<i>Configuring Custom Monitors with Retained Data</i>	<i>284</i>
External Check	286
<i>Configuring External Check Monitors.....</i>	<i>287</i>

Configuring Users

Working with User Roles.....	290
<i>Adding User Roles</i>	<i>290</i>
<i>Viewing User Roles</i>	<i>291</i>
<i>Editing User Roles</i>	<i>292</i>
Working with Users.....	293
<i>Adding Users.....</i>	<i>293</i>
<i>Viewing Users</i>	<i>296</i>
<i>Editing User Information.....</i>	<i>296</i>
Working with User Groups	297
<i>Adding User Groups.....</i>	<i>298</i>
<i>Viewing User Groups</i>	<i>298</i>
<i>Editing User Groups.....</i>	<i>298</i>
<i>Deleting User Groups.....</i>	<i>299</i>
Working with Notification Groups	300
<i>Adding Notification Groups.....</i>	<i>300</i>
<i>Viewing Notification Groups</i>	<i>301</i>
<i>Editing Notification Groups.....</i>	<i>301</i>

Understanding Report Options

Overview.....	304
Generating Reports.....	305
<i>Report Generation Options</i>	<i>305</i>
Saving Reports.....	307
<i>Saving Reports to the File System.....</i>	<i>307</i>
<i>Viewing Saved Reports</i>	<i>308</i>
Scheduling Reports	310
The Report Log.....	313
<i>Viewing Report Logs</i>	<i>314</i>
<i>Deleting Report Log Entries</i>	<i>315</i>

Using Reports

Reports for Performance and Analysis	318
<i>Resource Usage Report.....</i>	<i>318</i>
<i>Multi-System CPU Report.....</i>	<i>322</i>
<i>CPU Utilization Summary Report.....</i>	<i>323</i>
<i>CPU Utilization Ratio Report.....</i>	<i>326</i>
<i>Wait I/O Report.....</i>	<i>327</i>
<i>Service Monitor Metrics Report.....</i>	<i>329</i>
Reports for Capacity Planning	332
<i>Enterprise CPU Utilization Report.....</i>	<i>332</i>
<i>File System Capacity Growth Report</i>	<i>335</i>
<i>Server Virtualization Report</i>	<i>336</i>
<i>Solaris Mutex Exception Report.....</i>	<i>340</i>
<i>Network Bandwidth Report.....</i>	<i>342</i>
<i>Disk I/O Bandwidth Report.....</i>	<i>345</i>
<i>CPU Run Queue Threshold Report</i>	<i>348</i>

<i>File System Service Time Summary Report</i>	<i>353</i>
Reports for Availability	357
<i>Application Availability Report.....</i>	<i>357</i>
<i>Service Monitor Availability Report.....</i>	<i>358</i>
<i>Service Monitor Outages Report.....</i>	<i>359</i>
Reports for J2EE Applications	361
<i>WebSphere Report.....</i>	<i>361</i>
<i>WebLogic Report.....</i>	<i>364</i>
Reports for Virtual Environments	368
<i>VMware Workload Report.....</i>	<i>368</i>
<i>VI3 Workload Profile Report.....</i>	<i>371</i>
<i>LPAR Workload Report</i>	<i>373</i>

Understanding Graphing

Graphing in up.time	378
<i>Graphing Tool.....</i>	<i>379</i>
Using the Graph Editor	380
<i>Working with Trend Lines.....</i>	<i>382</i>
<i>Formatting Individual Graph Elements.....</i>	<i>383</i>
<i>Exporting Graphs</i>	<i>384</i>
<i>Changing the Look and Feel of a Graph</i>	<i>384</i>

Using Graphs

Overview.....	386
<i>UNIX vs. Windows Performance Monitoring</i>	<i>386</i>
Viewing the Status of a System.....	387
<i>Viewing a Quick Snapshot</i>	<i>387</i>

Monitoring CPU Performance	388
Usage (% busy).....	388
Run Queue Length.....	390
Run Queue Occupancy.....	390
Generating a CPU Performance Graph	391
Multi-CPU Usage	392
Generating a Multi-CPU Usage Graph.....	392
Graphing Memory Usage	395
Free	395
Cache Hit Rate.....	395
Paging Statistics.....	396
Free Swap.....	396
Generating a Memory Usage Graph	397
Graphing Processes.....	398
Number of Processes.....	398
Process Running, Blocked, Waiting	398
Process Creation Rate	399
Generating a Process Graph	399
Graphing TCP Retransmits	400
Generating a TCP Retransmits Graph	400
Graphing User Activity	401
Generating a User Activity Graph	401
Workload Graphs	402
Generating a Workload Graph	403
Workload Top 10 Graphs.....	404
Generating a Workload Top 10 Graph	404
LPAR Workload Graphs	405
Generating an LPAR Workload Graph	405

LPAR CPU Utilization Graphs	406
Network Graphs	407
I/O	407
Errors	407
Generating a Network Graph	408
Disk Performance Statistics Graph	409
Generating a Disk Performance Statistics Graph.....	409
Top 10 Disks Graph	411
Generating a Top 10 Disks Graph.....	411
File System Capacity Graph	413
Generating a File System Capacity Graph.....	413
VXVM Stats Graph	414
Generating a VXVM Stats Graph	414
Novell NRM Graphs	416
Generating a Novell NRM Graph	417
Instance Motion Graphs	418
Generating an Instance Motion Graph	418
Displaying Detailed Process Information	419
Generating Detailed Process Information	420
 Using the Config Panel	
Archive Policy	424
Configuring an Archive Policy	424
Importing Archived Data into the DataStore	425
Mail Servers	427

License Information	428
Monitoring Periods.....	429
Adding Monitoring Periods	429
Problem Reporting	431
Generating a Problem Report	431
up.time Configuration	432

Additional Information

Adding Multiple Systems to up.time	434
Creating a Hosts File.....	434
Fields in the Hosts File.....	434
Adding Multiple Systems to up.time	437
Examples of Hosts File Entries	438
Novell NRM Systems	440
Adding a Novell NRM System to up.time	440
NRM Statistics Captured by up.time	441
Graphing Statistical Information	446
Enabling SSH on a Hardware Management Console.....	447
Enabling the Windows Messaging Service	448
Changing the Agent Password on Windows.....	449
Custom Alert Format Variables.....	450
Archiving and Restoring the DataStore	452
Archiving the DataStore	452
Restoring the DataStore.....	452

The uptime.conf File

Overview	456
Database Settings	457
Mail Server Settings	459
SNMP Trap Settings	460
Basic System Settings	462
Global Scan Threshold Settings	465
Remote Reporting Settings	466
RSS Feed Settings	467
Splunk Settings	468

SQL Server Ports and Locks

SQL Server Ports	470
Verifying the Configuration of a Port	470
Configuring a Static Port	472
Configuring a Dynamic Port	473
SQL Server Scripts	474
SQL Server Locks	475

Securing up.time

Implementing SSL for the Monitoring Station	480
Stopping the up.time Services	480
Modifying the uptime.conf file	480

Configuring SSL	481
Starting the up.time Services	484
Securing the DataStore	485
Change the Passwords for the Default Users	485
Securing Agents using SSL.....	487

End User License Agreement

END USER LICENSE AGREEMENT	490
NOTICE TO USER.....	490
1. License.....	491
2. Intellectual Property and Confidentiality.....	492
3. License Fees.....	494
4. Term and Termination.....	494
5. Remedies and Indemnification.....	495
6. Disclaimer	496
7. Limitation of Liability.....	496
8. General Terms.....	497

Index

CHAPTER 1

Welcome to up.time

.....

This chapter introduces up.time in the following sections:

<i>Introducing up.time</i>	<i>2</i>
<i>up.time Architecture</i>	<i>3</i>
<i>up.time Service Monitoring Concepts</i>	<i>4</i>

Introducing up.time

up.time is a Web browser-based application that 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:

- 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.

Who Should Read This Guide

The up.time 4 user guide is intended for:


- System administrators who want to use up.time to monitor a single system or multiple systems in a distributed environment.

Once up.time is installed and running, administrators are informed when a system is failing, and can automatically or manually fix the problem.

- 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 a monitoring station that retrieves information from client systems, either through software (called *agents*) that is installed on a system or by monitoring services running on a system. 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.

- Monitors

The service monitor templates that are bundled with [up.time](#). You use these templates to configure a service check.

- 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.

- Host Checks

Service checks that you select and assign to each host that is being monitored to test if it is functioning properly. Service checks are temporarily disabled if [up.time](#) determines that a host that is undergoing scheduled maintenance.

- 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.

- Service Groups

Service monitor templates that enable you to apply a common service check to one or multiple hosts (servers, network devices) that you are monitoring.

CHAPTER 2

Understanding up.time


.....

This chapter explains the concepts underlying the [up.time](#) software in the following sections:

<i>Understanding the up.time Interface.....</i>	<i>6</i>
<i>Understanding Reports and Graphs</i>	<i>12</i>
<i>Understanding Graphs.....</i>	<i>12</i>
<i>Understanding Alerts</i>	<i>13</i>
<i>Understanding Actions.....</i>	<i>16</i>
<i>Understanding Agents.....</i>	<i>17</i>
<i>Understanding the up.time DataStore</i>	<i>19</i>
<i>Understanding Service Monitors</i>	<i>21</i>
<i>Understanding Services.....</i>	<i>24</i>
<i>Understanding Services.....</i>	<i>24</i>
<i>Understanding the Status of Services</i>	<i>25</i>
<i>Understanding Security in up.time</i>	<i>26</i>
<i>Understanding Dates and Times</i>	<i>27</i>
<i>Understanding Retained Data</i>	<i>29</i>

Understanding the up.time Interface

up.time is a Web browser-based application that consists of six main sections. The following image displays the up.time application screen. 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
- Services
- Users
- Reports
- Config

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 obtain the lowest level of resource information in your environment.

For more information about using the **Global Scan** panel, see “Using Global Scan” on page 99.

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 67.

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:

- Systems
- Groups
- Applications
- Views

For more information about using the **My Enterprise** panel, see “Using My Enterprise” on page 71.

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:

- 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” on page 111.

Users

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

- Users.
- User groups.
- Notification groups.
- User roles.

For more information about using the **Users** panel, see “Configuring Users” on page 289.

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 317.

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.

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 “Using the Config Panel” on page 423.

System List

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

- 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 the system belongs, if any.

You access the system list by clicking the **Syslist** icon in the top-right corner of the **up.time** Web interface. A window like the following one appears:

System List		
	Display Name	Entity Group Name
	Active Directory	10.1.1.19
	AIX-5L	My Enterprise
	Brightmail	UNIX Systems
	Brightmail	brightmail2
	Cypher	Windows Systems
	Cypher	cypher
	WebLogic Server	Linux Systems
	WebLogic Server	dev-rmeloch
	ELinux	My Enterprise
	Filter	Linux Systems
	Filter	filter
	FogBugz	Solaris Systems
	FogBugz	fogbugz
	Ginger	Windows Systems
	iRobot	Linux Systems
	WebSphere Server	ginger
		irobot
		Windows Systems
		lab-websphere
		My Enterprise

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.

Icons

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

-  Clone
Makes a copy of an entry in a panel. You can then modify the entry.
-  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.

System Icons

The following icons appear in the **Global Scan** and **My Enterprise** panels, and identify the type of system that up.time is monitoring:



Linux



AIX



Solaris



Novell NRM



Windows



HP-UX



VMware ESX



Net-SNMP



WebSphere

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. You can dynamically generate and view reports, schedule and email reports to other up.time users.

For detailed information about reports, see “Using Reports” on page 317.

If you need to regularly run certain reports, you can save them to the **My Portal** panel. See “Scheduling Reports” on page 310 for more information.

Understanding Graphs

You can graph performance information when you need to collect and 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. Graph display data points in the collected performance information and when those data points were collected.

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 – for example, in Mozilla on Linux.

For more information on graphing, see “Understanding Graphing” on page 377 and “Using Graphs” on page 385

Understanding Alerts

When a problem occurs on a system that [up.time](#) is monitoring, the 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 25 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 a host or an application 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
1/12/2006 10:52
Host: filter
Host State: N/A
Service: FS Capacity - Filter
Service State: WARN/
Output: /var is 92% full
```

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 125.

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.


In 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:



Understanding Alert Profiles

Alert profiles are templates that tell **up.time** how to react to alerts that are generated by your service checks. Alert profiles trigger reactions in response to service check failures and exceeded thresholds. 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 the above actions to occur simultaneously. For example, if a Web server process stops responding, the system administrator can be notified.

For information about configuring alert profiles, see “Alert Profiles” on page 129.

Understanding Actions

You can configure **up.time** to initiate an action when it detects a problem. Actions automatically perform one of the following tasks:

- Write an entry to a log file.
- Run a script – for example, to restart a server that is not responding.
- Stop, start, or restart a Windows server.
- Send network notifications to an SNMP trap.

Understanding Action Profiles

Action profiles are templates that tell **up.time** what action to take when a monitor encounters a problem. You can apply an action profile to one or more service. Using action profiles, you have a common set of responses to a problem that you expect to encounter.

With an action profile, you can specify the following:

- When to start and stop performing the action.
- The name of the log file to which **up.time** writes error information.
- The name and location of a recovery script that can restart an application, process, or service; or reboot a server.
- The name of a Windows host and service, as well as whether to start, stop, or restart the host or service.
- An SNMP trap to be sent to a specific trap host and trap community.

For information about configuring action profiles, see “Action Profiles” on page 134.

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 and minor versions of **up.time** agents are shown in the following diagram:



- 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.

- 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 such information as:

- 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. Any sensitive information that is in the DataStore is encrypted.
- The settings for service monitors, alert and action profiles, scheduled maintenance, and host checks.
- Reports that have been saved in the up.time Web interface and which 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. For example, if you add a host to up.time, basic information about that host is written to the entity table and configuration information for the host is written to the entity_configuration table in the DataStore.

Access to the DataStore is determined by one of the three installed user accounts, which are described in the section “Database Users” on page 485. Each account gives users varying levels of access to the contents of the DataStore.

up.time 4 can also use either an Oracle or MS SQL Server database as its DataStore. If you plan to use either of these databases, refer to the Knowledge Base articles “Running up.time 4 on Oracle” and “Running up.time on Microsoft SQL Server 2000” for the additional steps required to enable up.time to work with these databases.

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” on page 111.

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.



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 “SNMP” on page 269.

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” on page 111.

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.

See “Service Groups” on page 139 for more information.

Understanding Security in up.time

up.time uses several methods to maintain the security of your data and of the up.time monitoring station.

All users of up.time must log into the monitoring station with a user name and password. As well, administrators can restrict the access a user has to the functions of up.time by configuring and assigning user roles and user groups. The functions to which users do not have access will not appear in the up.time Web interface. For more information, see “Working with User Roles” on page 290 and “Working with User Groups” on page 297.

Administrators can also restrict access to the up.time DataStore by only allowing users to connect to the DataStore from specific IP addresses or hosts. For more information, see “Starting and Stopping up.time” on page 52.

As well, you can use SSL (Secure Sockets Layer, a widely-used protocol for securing information travelling over networks) to enable the up.time monitoring station to communicate with various agents and remote hosts. See “Implementing SSL for the Monitoring Station” on page 480 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, and 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 329) or a Service Metrics graph (see “Viewing System and Service Information” on page 57).

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:

Exchange Settings		Retained Data Tracking
Port ▾	9998	
Use SSL ▾	☐ Use SSL	
☒ Web Mail Sends Per Second		
Web Mail Sends Per Second ▾		☒
Warning	is greater than or equal to ▾ 70	
Critical	is greater than or equal to ▾ 100	
☒ Web Mail Auths Per Second		
Web Mail Auths Per Second ▾		☒
Warning	is greater than or equal to ▾ 40	
Critical	is greater than or equal to ▾ 75	

CHAPTER 3

Installing up.time

.....

This chapter explains how to install [up.time](#) in the following sections:

<i>Installation Plan</i>	<i>32</i>
<i>Installation Requirements</i>	<i>33</i>
<i>Installing the up.time Monitoring Station.....</i>	<i>35</i>
<i>Installing Agents</i>	<i>43</i>

Installation Plan

Before installing **up.time** you must:

- Identify the computer that will act as a central monitoring station.
- Ensure that all client systems that you want to monitor are accessible over the network.

All communication with client systems is over TCP using port 9998. However, you can specify a different port during the installation process. All communication originates from the monitoring station. When a host that is being monitored is outside a firewall, you only need to configure outbound port access.

If you purchased the boxed version of **up.time**, the central monitoring station 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 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 monitoring station or any **up.time** agents.

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.



On HP/UX, you cannot start processes, such as agents, using the **nobody** user ID. Also, on Windows 2000 the agent must be running with Administrator privileges. If it is not, the agent will not be able to access the system performance counters.

Installation Requirements

This section describes the system requirements for the following:

- up.time Monitoring Station
- up.time Agents



Check the uptime software Web site (<http://www.uptime-software.com>) for the most up-to-date list of hardware and software requirements.

up.time Monitoring Station

The **up.time** monitoring station is a computer running the core **up.time** software that retrieves information from client systems, either through agents installed on the system or by monitoring services running on the system. The monitoring station has a self contained Web server and database that enables easy access to the application and data.

The **up.time** 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 XP	XP Professional
Microsoft Windows 2003	• Standard • Enterprise
Solaris (32-bit SPARC)	• 8 • 9 • 10
Red Hat Linux AS (x86)	4.0
Red Hat Linux ES (x86)	4.0
SuSe Linux Enterprise Server	9
AIX (on pSeries POWER architectures)	5.3

Supported Web Browsers

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

- Internet Explorer 6.0 or higher
- Firefox 1.0 or higher
- Mozilla 1.7

Minimum Hardware Configuration

The hardware configurations for a monitoring station can change depending on the number of agents that you want to monitor, the reports that you want to generate, and the amount of data that in the **up.time** DataStore.



Contact uptime software Client Care if you are monitoring more than 50 nodes.

The recommended minimum hardware for an **up.time** monitoring station is:

- 1 GHz processor.
- 1 GB of memory.
- 18 GB of disk storage.
- 100Mbps network interface.

up.time Agents

You can install and use **up.time** agents to collect data from a number of operating systems. Check the uptime software Client Care Web site for the most up-to-date list of supported platforms and architectures.



up.time can monitor Novell NetWare NRM version 6.5. Earlier versions of NRM are not supported.

up.time also supports agentless monitors on any operating system, which do not require you to install software on a system or device. See “Using Agentless Monitors” on page 114.

Installing the up.time Monitoring Station

The monitoring station is installed a single directory. On Windows, the monitoring station is installed in the directory C:\Program Files\uptime software\uptime4. On Solaris, Linux or AIX, it is installed in the directory /usr/local/uptime4.

On Windows, the up.time monitoring station is installed using a graphical installer that guides you through the steps of the installation process. On Solaris, Linux and AIX, the installer is a console application.



Before installing up.time, you must be logged in as Administrator (in Windows) or as root (in Solaris, Linux or AIX).

up.time 4 can also use either an Oracle or MS SQL Server database as its DataStore. If you plan to use either of these databases, refer to the Knowledge Base articles “Running up.time 4 on Oracle” and “Running up.time on Microsoft SQL Server 2000” for the additional steps required to enable up.time to work with these databases.

Before You Begin

There are two ways in which to install the up.time monitoring station:

1 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, Linux or AIX, mount the CD-ROM drive if you are not using automount.
- Change to the following directory on the CD:

```
up.time_MonitoringStation
```

2 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 monitoring station. For the Windows installer, extract the contents of the archive using a utility like WinZip.

Then, refer to the procedures in the sections “Installing the Monitoring Station on Windows” on page 36 and “Installing the Monitoring Station on Solaris, Linux or AIX” on page 38 for details on completing the installation for your platform.

Installing the Monitoring Station on Windows

To install the [up.time](#) monitoring station on Windows, do the following:

1 Double click on the following file:

`up.time-4.0.<build#>-win32-x86.exe`

Where <build#> is the number of the [up.time](#) build that you are installing. For example:

`up.time-4.5.100-win32-x86.exe`

2 On the Introduction screen, click Next.

3 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.

4 Click Next.

5 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\uptime4).
- 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**.

6 Do one of the following to set the location where the up.time DataStore will be installed:

- Click **Next** to accept the default location (C:\Program Files\uptime software\uptime4\DataStore).
- In the **Please Choose a Folder** field, type the name of the directory where you want to install the DataStore and then click **Next**. This should be the full path to the DataStore.



Because the DataStore can grow very large (in excess of 100 GB), you can install the DataStore in another folder on the file system if you are monitoring a large number of systems and retaining data for extended periods.

- Click **Choose** and select a directory from the **Browse for Folder** window.

7 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:

- SMTP Server

The name of the default mail host. This value is written to the file `uptime.conf` which contains configuration information for up.time.

- Email address

The email address from which the 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`.

- 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 monitoring station will listen for requests. The port number is written to the file `httpd.conf`.

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

9 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.

10 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.

11 On the Install Complete screen, click Next.

12 Click Finish.

Installing the Monitoring Station on Solaris, Linux or AIX

Installation on Solaris, Linux or AIX 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.

When you install the **up.time** monitoring station on AIX 5.3, the installer creates a user called uptime and adds the following lines to the end of the file `/etc/security/limits`:

```
uptime:
    fsize = 2097151
    core = -1
    cpu = -1
    data = -1
    rss = -1
    stack = -1
    nofiles = 2000
```

This allows **up.time** and the installer to exceed the default system resource limits, which can otherwise interfere with the ability of the application to perform monitoring and analysis.

Installing the Monitoring Station

To install the **up.time** monitoring station on Solaris, Linux or AIX, do the following:

1 Type the following command:

```
sh up.time-4.0.<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-4.5.100-rhes4-x86.bin` or `up.time-4.5.100-sles9-x86-upgrade.bin`
- Solaris: `up.time-4.5.100-solaris-sparc.bin`
- AIX: `up.time-4.5.100-aix.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.

2 On the Introduction page, press Enter to continue.

- 3 On the **License Agreement** page, carefully read the **up.time** end user license agreement. Press Enter to scroll through the agreement.
- 4 At the **DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N)** prompt, type **y** and press Enter.
- 5 Do one of the following to set the directory in which **up.time** will be installed:

- Press Enter to accept the default location (`/usr/local/uptime4/`).
- Type a new location at the command prompt (for example, `/opt/uptime4/`) and then press Enter.



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

- 6 Do one of the following to set the location where the **up.time DataStore** will be installed:
- Press Enter to accept the default location (`/usr/local/uptime4/DataStore`).
 - Type a new location at the command prompt (for example, `/opt/uptime4/DataStore`) and then press Enter. This should be the full path to the DataStore.



Because the DataStore can grow very large (in excess of 100 GB), you can install the DataStore in another folder on the file system if you are monitoring a large number of systems and retaining data for extended periods.

- 7 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 monitoring station will listen for requests. The port number is written to the file `httpd.conf`.
 - SMTP Server

The name of the default mail host. This value is written to the file `uptime.conf` which contains configuration information for [up.time](#).
 - up.time email address

The email address from which the 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`.

8 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.

9 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.

10 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 monitoring station.

Post-Installation Tasks

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

- Set up the administrator account when you first log in. See “Setting Up the Administrator Account” on page 55.
- Install the license for up.time. See “License Information” on page 428.
- Add users and systems. See “Configuring Users” on page 289 and “Using My Enterprise” on page 71.

Upgrading to up.time 4

If you are using a previous version of up.time and intend to upgrade to version 4, you can find detailed information about the upgrade process at the Client Care Web site (<http://support.uptimesoftware.com>).

Installing Agents

up.time agents are used to retrieve detailed performance statistics – such as CPU, memory, process, disk, and network usage – from the hosts that you are monitoring. The agents can also securely and remotely execute programs. The Windows agent can start and stop services, and reboot the machine.

The installation process for agents varies by operating system. On UNIX, Linux, and IBM pSeries systems installation is done at the command line using a script. On Windows, installation is done using a graphical utility.



All client systems must be accessible via a name. This name should exist in either the /etc/hosts table on the monitoring station, or be accessible via a nameserver – for example files, NIS, or DNS. If the host IP is changed then the monitoring station may send requests to the incorrect machine.

Installing Agents on Windows

The installer for Windows **up.time** agents uses a wizard that guides you through the installation process.



If the Windows installer requires unavailable service packs – for example, SiteServer or Terminal Server – send an email to support@uptimesoftware.com and request the extracted agent which can be installed without using the Windows installer.

To install an agent on Windows, do the following:

- 1 **Copy the installer (setup.exe) for the Windows agent to the system on which you want to install the agent.**
- 2 **In Windows Explorer, double click the file setup.exe.**
- 3 **On the installer Welcome screen, click Next.**

- 4 On the **Select Installation Folder** screen, type the path to the folder in which you want to install the agent in the **Folder** field.

Alternatively, click the **Browse** button and use the dialog box that appears to search for the folder.

- 5 Select the checkbox **Make available for Everyone** option.
- 6 Click **Next**.
- 7 On the **Confirm Installation** dialog screen, click **Next**.

Installing Agents on Solaris

You install up.time agents for Solaris at the command line.

To install an agent on Solaris, do the following:

- 1 Log into the system as user root.
- 2 Using telnet or FTP, transfer the archive containing the agent to the system on which you want to install the agent.

You should copy the archive to a temporary directory on the system.

- 3 Extract the archive using the following command:

```
tar -xvf uptmagnt-<version>.tar
```

Where <version> is the version of the agent, for example solaris-4.0.

- 4 Run the following command:

```
pkgadd -d
```

- 5 Follow the prompts from the pkgadd utility to select the agent package and install it.

Installing Agents on UNIX

You install up.time agents for various UNIX platforms at the command line using a shell script.

To install an agent on a UNIX system, do the following:

- 1 Log into the system as user root.

- 2 **Using telnet or FTP, transfer the archive containing the agent to the system on which you want to install the agent.**

You should copy the archive to a temporary directory on the system.

- 3 **Extract the archive.**

Depending on the version of UNIX, you will need to extract the archive using either the tar command or a combination of the gzip and tar commands. For example, to extract the agent for AIX use the following command:

```
tar -xvf uptmagnt-AIX-4.0.tar
```

To extract the agent for Tru-64, use the following commands:

```
gzip -d uptmagnt-Tru64-4.0.tar.gz  
tar -xvf uptmagnt-Tru64-4.0.tar
```

- 4 **Type the following command at the command line:**

```
./INSTALL.sh
```

- 5 **Follow the prompts to complete the installation.**

Installing Agents on Linux

You can install [up.time](#) agents for Linux using the RPM utility or the Debian package management utility (dpkg). This enables you to easily update and perform mass installations of agents.



Before trying to install an agent, ensure that the RPM or dpkg utilities are installed and are in the path by typing one of the following commands at the command line:

```
which rpm  
which dpkg
```

To install an agent on a Linux system, do the following:

- 1 **Log into the system as user root.**
- 2 **Using telnet or FTP, transfer the .rpm or .deb file containing the agent to the system.**

- 3 If you are installing the agent using the RPM utility, type the following at the command line:**

```
rpm -i <agent name>
```

Where <agent name> is the name of the .rpm file for the agent that you are installing. For example, uptimeagent-4.0.rpm.

- 4 If you are installing the agent using the dpkg utility, type the following at the command line:**

```
dpkg -i <agent name>
```

Where <agent name> is the name of the .deb file for the agent that you are installing. For example, uptimeagent-4.0.deb.

Installing Agents on IBM pSeries Servers

up.time 4 can collect workload information from IBM pSeries servers that have logical partitions (LPARs). Before up.time can do this, you must install the latest AIX or Linux agents on the LPARs whose workloads you want to profile.

There are two options for installing agents on IBM pSeries servers with logical partitions (LPARs):

- Installing the agent on a pSeries server with an HMC
- Installing the agent on a pSeries server without an HMC

In both cases, you will need to install the agent on each LPAR as well as the Virtual I/O (VIO) partition.

Installing the agent on a pSeries server with an HMC

Before you can monitor the logical partitions on an IBM pSeries server, you must install an agent on each LPAR and on the VIO. Use the following instructions to install the agent on an IBM pSeries server that is managed by an HMC.

To install an agent on an IBM pSeries server with an HMC, do the following:

- 1 Enable root login on the HMC.**

2 If Linux is running on the LPAR, do the following:

- Log into the LPAR as root.
- Copy the RPM file containing the Linux agent to the LPAR.
- Run the following command:

```
rpm -i <agent name>.rpm
```

Where <agent name> is the name of the .rpm file for the agent that you are installing. For example, UptimeAgent-Linux-4.4.1-0.src.rpm.



If you are using SuSe Linux Enterprise Server 9, you must update the kernel to the latest version using the YAST package manager. If you do not upgrade the kernel, the agent will not be able to gather workload data.

3 If AIX is running on the LPAR, do the following:

- Log into the LPAR as root.
- Copy the archive containing the agent to the LPAR.
- Extract the contents of the archive using the following command:

```
tar -xvf <agent name>
```

Where <agent name> is the name of the archive that contains the agent that you are installing. For example, uptmagnt-AIX-4.4.1.tar.

- Run the following command to install the agent:

```
./INSTALL.sh
```

4 Do the following to install the agent on the VIO:

- Log into the VIO as root.
 - Run the following command.
- ```
oem_setup_env
```
- Copy the archive containing the agent to the LPAR.

- Extract the contents of the archive using the following command:  

```
tar -xvf <agent name>
```

Where <agent name> is the name of the archive that contains the agent that you are installing. For example, `uptmagnt-AIX-4.4.1.tar`.
- Run the following command to install the agent:  

```
./INSTALL.sh
```

### Installing the agent on a pSeries server without an HMC

Before you can monitor the logical partitions on an IBM pSeries server, you must install an agent on each partition. Use the following instructions to install the agent on an IBM pSeries server that is not managed by an HMC.

To install the agent, do the following:

#### 1 If Linux is running on the LPAR, do the following:

- Log into the LPAR as root.
- Copy the RPM file containing the agent to the LPAR.
- Run the following command:  

```
rpm -i <agent name>.rpm
```

Where <agent name> is the name of the `.rpm` file for the agent that you are installing. For example, `UptimeAgent-Linux-4.4.1-0.src.rpm`.



If you are using SuSe Linux Enterprise Server 9, you must update the kernel to the latest version using the YAST package manager. If you do not upgrade the kernel, the agent will not be able to gather workload data.

#### 2 If AIX is running on the LPAR, do the following:

- Log into the LPAR as root.
- Copy the archive containing the agent to the LPAR.

- Extract the contents of the archive using the following command:

```
tar -xvf <agent name>
```

Where <agent name> is the name of the archive that contains the agent that you are installing. For example, `uptmagn-AIX-4.4.1.tar`.

- Run the following command to install the agent:

```
./INSTALL.sh
```

### 3 Do the following to install the agent on the VIO:

- Log into the VIO as root.
- Copy the archive containing the agent to the LPAR.
- Extract the contents of the archive using the following command:

```
tar -xvf <agent name>
```

Where <agent name> is the name of the archive that contains the agent that you are installing. For example, `uptmagn-AIX-4.4.1.tar`.

- Run the following command to install the agent:

```
./INSTALL.sh -vio
```





# CHAPTER 4

## Getting Started

.....

This chapter introduces you to the basic features of [up.time](#) in the following sections:

|                                                     |           |
|-----------------------------------------------------|-----------|
| <i>Starting and Stopping up.time .....</i>          | <i>52</i> |
| <i>Accessing and Exiting up.time .....</i>          | <i>55</i> |
| <i>Viewing System and Service Information .....</i> | <i>57</i> |
| <i>Searching and Filtering .....</i>                | <i>63</i> |
| <i>Audit Logging .....</i>                          | <i>66</i> |

## Starting and Stopping up.time

In addition to the Web interface, the [up.time](#) 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 and start the [up.time](#) services – for example, after making configuration changes or performing an upgrade.

### Starting the up.time Services

To start the [up.time](#) services, do the following:

**1 In Windows, do the following:**

- Select **Start > Control Panel**.
- Double click **Administrative Tools**, and then double click **Services**.
- In the **Services** window, find the following entries and click **Start the service**:
  - up.time 4 Data Store
  - up.time 4 Web Server
  - up.time 4 Data Collector

**2 On Solaris or Linux, do the following:**

- Log into the monitoring station as user root.
- Type the following command to start the database:  
`/etc/init.d/uptime4_datastore start`

- Type the following command to start the Web server:  
`/etc/init.d/uptime4_httpd start`
- Type the following command to start the Data Collector:  
`/etc/init.d/uptime4_core start`

**3 On AIX do the following:**

- Log into the monitoring station as user root.
- Type the following command to start the database:  
`/etc/rc.d/init.d/uptime4_datastore start`
- Type the following command to start the Web server:  
`/etc/rc.d/init.d/uptime4_httpd start`
- Type the following command to start the Data Collector:  
`/etc/rc.d/init.d/uptime4_core start`

## Stopping the up.time Services

To stop the up.time services, do the following:

**1 In Windows, do the following:**

- Select **Start > Control Panel**.
- Double click **Administrative Tools**, and then double click **Services**.
- In the **Services** window, find the following entries and click **Stop the service**:
  - up.time 4 Data Store
  - up.time 4 Web Server
  - up.time 4 Data Collector

**2 On Solaris or Linux, do the following:**

- Log into the monitoring station as user root.

- Type the following command to stop the database:  
`/etc/init.d/uptime4_datastore stop`
- Type the following command to stop the Web server:  
`/etc/init.d/uptime4_httpd stop`
- Type the following command to stop the Data Collector:  
`/etc/init.d/uptime4_core stop`

### **3 On AIX do the following:**

- Log into the monitoring station as user root.
- Type the following command to stop the database:  
`/etc/rc.d/init.d/uptime4_datastore stop`
- Type the following command to stop the Web server:  
`/etc/rc.d/init.d/uptime4_httpd stop`
- Type the following command to stop the Data Collector:  
`/etc/rc.d/init.d/uptime4_core stop`

## 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>`

Where <uptime\_hostname> is the name or IP address of the server that is hosting the 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 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>`

Where `<uptime_hostname>` is the name or IP address of the server that is hosting the 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:

- 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 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, for example 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 96.

- **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 25.

- 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                                  |
| Fri Mar 09 14:57:21 EST 2007 | uptime-10.1.1.35                                                                                   | CRIT        | retry          | OK        |              | up.time agent running on 10.1.1.35, up   |
| Fri Mar 09 14:56:21 EST 2007 |  uptime-10.1.1.35 | OK          |                | CRIT      | retry        | Monitor failed: Read timed out           |
| Mon Mar 05 10:36:44 EST 2007 | PING-10.1.1.35                                                                                     | CRIT        | retry          | OK        |              | Ping completed: 1 sent, 0.0% loss, 0.0   |
| Mon Mar 05 10:34:46 EST 2007 |  PING-10.1.1.35   | OK          |                | CRIT      | retry        | Monitor failed: process has not exited   |
| Mon Mar 05 09:14:59 EST 2007 | PING-10.1.1.35                                                                                     | CRIT        | retry          | OK        |              | Ping completed: 1 sent, 0.0% loss, 0.0   |
| Mon Mar 05 09:14:01 EST 2007 |  PING-10.1.1.35   | OK          |                | CRIT      | retry        | Monitor failed: process has not exited   |
| Mon Feb 26 06:44:13 EST 2007 | Configuration Update Gatherer                                                                      | UNKNOWN     | host down      | OK        |              | Agent: up.time agent 4.1 solaris 31, Re  |
| Sun Feb 25 18:42:14 EST 2007 | PING-10.1.1.35                                                                                     | OK          | recovery       | OK        |              | Ping completed: 1 sent, 0.0% loss, 0.0   |
| Sun Feb 25 18:40:16 EST 2007 | uptime-10.1.1.35                                                                                   | UNKNOWN     | host down      | OK        |              | up.time agent running on 10.1.1.35, up   |
| Sun Feb 25 18:40:16 EST 2007 | Platform Performance Gatherer                                                                      | UNKNOWN     | host down      | OK        |              | Agent: up.time agent 4.1 solaris 31, Re  |
| Sun Feb 25 18:37:15 EST 2007 | PING-10.1.1.35                                                                                     | CRIT        |                | OK        | recovery     | Ping completed: 1 sent, 0.0% loss, 0.0   |
| Sun Feb 25 18:30:36 EST 2007 | Platform Performance Gatherer                                                                      | OK          |                | UNKNOWN   | host down    | Agent: up.time agent 4.1 solaris 31, Re  |
| Sun Feb 25 18:30:36 EST 2007 | uptime-10.1.1.35                                                                                   | OK          |                | UNKNOWN   | host down    | up.time agent running on 10.1.1.35, up   |
| Sun Feb 25 18:30:36 EST 2007 |  PING-10.1.1.35   | CRIT        | retry          | CRIT      |              | Ping completed: 1 sent, 100.0% loss, 0.0 |
| Sun Feb 25 18:27:21 EST 2007 |  PING-10.1.1.35   | OK          |                | CRIT      | retry        | Ping completed: 1 sent, 100.0% loss, 0.0 |
| Sun Feb 25 18:13:17 EST 2007 | uptime-10.1.1.35                                                                                   | CRIT        | retry          | OK        |              | up.time agent running on 10.1.1.35, up   |

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 (  ) 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 DOWN | Time UNKNOWN | Total Time |
| UPTIME- <i>elinux</i> | UNKNOWN | 14m 41s | 0s        | 0s        | 0s        | 1m 18s       | 15m 59s    |
| PING- <i>elinux</i>   | CRIT    | 10m 7s  | 0s        | 4m 36s    | 0s        | 1s           | 14m 44s    |
| FTP - <i>elinux</i>   | UNKNOWN | 9m 34s  | 0s        | 0s        | 0s        | 1m 17s       | 10m 51s    |
| HTTP - <i>elinux</i>  | UNKNOWN | 0s      | 0s        | 7m 7s     | 0s        | 1m 19s       | 8m 26s     |


  

| Availability As Percent |         |           |             |             |             |                |
|-------------------------|---------|-----------|-------------|-------------|-------------|----------------|
| Monitor                 | Status  | % Time OK | % Time WARN | % Time CRIT | % Time DOWN | % Time UNKNOWN |
| UPTIME- <i>elinux</i>   | UNKNOWN | 91.85%    | 0.00%       | 0.00%       | 0.00%       | 8.15%          |
| PING- <i>elinux</i>     | CRIT    | 68.65%    | 0.00%       | 31.24%      | 0.00%       | 0.11%          |
| FTP - <i>elinux</i>     | UNKNOWN | 88.15%    | 0.00%       | 0.00%       | 0.00%       | 11.85%         |
| HTTP - <i>elinux</i>    | UNKNOWN | 0.00%     | 0.00%       | 84.36%      | 0.00%       | 15.64%         |

[Generate Graph](#)

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

**PING-*elinux***



- Manage Services

Lists the following information about the services associated with a particular host:

- The name of the service.
- The service group, if any, to which the service belongs.
- The monitors, if any, associated with the service.

If the host is part of a service group, the services for all of the hosts that are members of the group appear in the **Manage Services** subpanel.

Click the name of the service to view information about that service. You can edit the service information, as well as the alert profiles and

action profiles associated with the service by clicking the appropriate button in the subpanel.

You can add services instances by clicking the **Add Service** tab in the **Manage Services** subpanel. The services that you add do not appear in the **Manage Services** but in the **Service Instances** subpanel. For more information about adding service instances, see “Using Service Monitors” on page 111.

- **Host Check**

List the basic checks (for example, a ping) for a system.

- **Maintenance**

Lists whether or not there are any maintenance periods scheduled for the system. For more information on maintenance periods, see “Scheduling Maintenance” on page 147.

**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 29.**

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 27.
- In the **Current Retained Service Metrics** area, select the retained data variables that you want to graph, as shown below:

| Current Retained Service Monitor Metrics |                      |                                     |                           |       |           |
|------------------------------------------|----------------------|-------------------------------------|---------------------------|-------|-----------|
| Instance Name                            | Instance Description | Select                              | Variable                  | Units | Data Type |
| Ping - Matrix                            |                      |                                     |                           |       |           |
|                                          |                      | <input checked="" type="checkbox"/> | Average Round Trip Time ▾ | ms    | decimal   |
|                                          |                      | <input checked="" type="checkbox"/> | Percent Loss ▾            | %     | decimal   |
|                                          |                      | <input checked="" type="checkbox"/> | Response time ▾           | ms    | integer   |
| <div>Generate Graph</div>                |                      |                                     |                           |       |           |

- 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 which the host is running.

- Any information in the four custom fields in the system profile. This information can include, for example, the job being done by the system (router, database server, Web server), its physical location, and the like.

For more information, see “Editing a System Profile” on page 88.

## Filtering Service Instances

If you have a large number of hosts and want to view information about a particular service instance associated with those hosts, you can filter out the services that you do not want to see in the **Service Instances** subpanel.

To filter service instances, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **View Service Instances**.
- 3 Enter text in one of the following fields in subpanel:
  - **Name**  
The name of a particular service instance, for example PING-Server1.  
You can enter partial names of service instances in this field. For example, if you want to filter on instances that contain the text Mailbox, enter Mailbox in the field.
  - **Host**  
The name of a host with which the service is associated. This can be the actual name of the host or the display name in the [up.time](#) Web interface.
  - **Monitor**  
The name of a particular monitor on which you want to filter. For example, Ping or LDAP.  
You can enter partial names of monitors in this field. For example, if you want to filter on File System Capacity, enter Capacity in the field.

**4 Click Filter By.**

All service instances that you have permissions to view and that match the filtering criteria appear in the subpanel. If, for example, only 12 of the service instances match your criteria, a message like the following one appears in the subpanel:

Search found 12 out of 21 services

**5 To view all matches, click the Show All button.**

**6 To remove the filter criteria and restore the complete list of services, click Clear.**

## 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.

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, do the following:

- 1 In the **up.time** Web interface, click **Config**, and then click **up.time Configuration**.**

The contents of the file `uptime.conf` appear in a text area.

- 2 Add the following line to the file:**

```
auditEnabled=yes
```

- 3 Click the **Update** button.**



# CHAPTER 5

## Using My Portal

.....

This chapter explains the **My Portal** panel in the following section:

*Overview* ..... 68

## 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 two sections:

- Assistance
- My Preferences
- Latest up.time Articles
- My Reports


## 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](#).
- The [up.time](#) 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:



## Current Issues

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

## 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 296 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 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 Reports

The **My Reports** section lists the reports that you have scheduled and saved. For more information on scheduling reports, see “Scheduling Reports” on page 310.

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 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.

# CHAPTER 6

## Using My Enterprise

.....

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


|                                                        |           |
|--------------------------------------------------------|-----------|
| <i>Overview .....</i>                                  | <i>72</i> |
| <i>Working with Groups .....</i>                       | <i>73</i> |
| <i>Working with Systems .....</i>                      | <i>76</i> |
| <i>Working with Applications .....</i>                 | <i>90</i> |
| <i>Working with Views .....</i>                        | <i>93</i> |
| <i>Deleting Systems, Applications, and Views .....</i> | <i>95</i> |
| <i>Acknowledging Alerts .....</i>                      | <i>96</i> |

## Overview

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

- Systems or network devices.
- Groups, which are sets of systems or devices that have been combined in a meaningful way.
- Applications, which provide the overall status for one or more services.
- Views, which enable non-administrative users to view only the systems in which they are interested.

The following image illustrates the **My Enterprise** panel:



The screenshot shows the 'My Enterprise' panel with a table of monitored systems. The table has five columns: Name, Description, Architecture, OS version, and Monitored?. The systems are grouped into categories like Linux Systems, Solaris Systems, and UNIX Systems.

| Name                             | Description      | Architecture            | OS version                           | Monitored? |
|----------------------------------|------------------|-------------------------|--------------------------------------|------------|
| Active Directory (10.1.1.19)     | -                | Windows_NT x86 Famil... | Windows Server 2003                  | ✓ Yes      |
| vmob-nrond                       | -                | . VMWare.E5% Serv...    | . VMWare.E5% Serv...                 | ✓ Yes      |
| WebLogic Server (dev-rmelochs)   | -                | Windows_NT x86 Famil... | Windows XP                           | ✓ Yes      |
| WebSphere Server (lab-websphere) | -                | Linux lab-websphere ... | RedHat 3(Taroon 2.4.21-4.ELsmp i686) | ✓ Yes      |
| <b>Linux Systems</b>             |                  |                         |                                      |            |
| Cypher (cypher)                  | A Linux box      | Linux cypher 2.4.18-... | RedHat 8.0(Psyche 2.4.18-14 i686)    | ✓ Yes      |
| ELinux (elinux)                  | A Linux box      | Linux elinux 2.6.9-5... | RedHat 4(Nahant 2.6.9-5.EL i686)     | ✓ Yes      |
| Ginger (ginger)                  | Net-SNMP system  | Linux Ginger 2.4.21-... | -                                    | ✓ Yes      |
| McKay (mckay)                    | Oracle server    | Linux McKay 2.4.21-4... | RedHat 3(Taroon 2.4.21-4.ELsmp i686) | ✓ Yes      |
| Neo (neo)                        | A Linux box      | Linux css3-rhes4 2.6... | RedHat 4(Nahant 2.6.9-5.ELsmp i686)  | ✓ Yes      |
| <b>Solaris Systems</b>           |                  |                         |                                      |            |
| Filter (filter)                  | Main Solaris box | SunOS filter 5.9 Gen... | 5.9                                  | ✓ Yes      |
| Subway (subway)                  | A Solaris box    | SunOS subway 5.8 Gen... | 5.8                                  | ✓ Yes      |
| <b>UNIX Systems</b>              |                  |                         |                                      |            |
| AIX-5L (aix5l)                   | -                | AIX aix5l 1 5 003250... | 5.1                                  | ✓ Yes      |
| Tru64 (tru64)                    | -                | OSF1 tru64 V5.1 2650... | V5.1                                 | ✓ Yes      |

## 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 the 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:

- Elements, which can be systems, nodes, or applications.
- The user groups that are allowed to view the systems or elements in a group. See “Working with User Groups” on page 297 for more information on user groups.



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

## Adding Groups

To add a group, do the following:

- 1 On the My Enterprise panel, click the Add Group tab.
- 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 group to which it will be subordinate in the Available Groups list, and then click Add.



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

- 5 Select the systems that you want to add to this group from the Available Elements list, and then click Add.
- 6 Select one or more sets of users who can view this group from the Available User Groups list, and then click Add.
- 7 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 Data Centers, and then add two nested groups called Production and Disaster Recovery.

The following image illustrates nested groups in the **My Enterprise** panel:

| Name                                                                                                                | Description                  |
|---------------------------------------------------------------------------------------------------------------------|------------------------------|
|  <a href="#">Main Group</a>        | Primary servers to monitor   |
|  <a href="#">Backup</a>            |                              |
|  <a href="#">Redundant Cluster</a> | A group of redundant servers |
|  <a href="#">Secondary Group</a>   | A secondary group to monitor |
|  <a href="#">Second Test Group</a> |                              |
|  <a href="#">Filter</a>            | A server to monitor          |
|  <a href="#">Test Group</a>       |                              |
|  <a href="#">ELinux</a>          | Another server to monitor    |

The group Redundant Cluster is the parent group, and the groups Secondary Group, Test Group, and Second Test Group are the nested groups.

You can assign the following to nested groups:

- Groups of systems or elements.
- Individual systems or elements.
- The [up.time](#) user groups that are allowed to view the systems or elements in a group.



You cannot assign a parent group to a subgroup or to any other ancestor. For example, in the image on the previous page you cannot add the group Main Group to the nested group named Test Group.



Before you begin, ensure that you have at least one parent group defined. For more information, see “Adding Groups” on page 73.

## Adding a Nested Group

To add a nested group, do the following:

- 1 In the **My Enterprise** panel, click the **Add Group** tab.
- 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 Select one or more groups from the **Available Groups** list, and then click **Add**.
- 6 Select the systems or 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**.

## Working with Systems

Systems are the network devices that you will monitor using [up.time](#). You can add the following types of systems:

- **Agent**  
A system that has an [up.time](#) agent installed on it. In the **Global Scan** and **My Enterprise** panels, agent systems are denoted by this icon: 
- **Node**  
A device without an agent, but with which [up.time](#) can communicate using an IP address. In the **Global Scan** and **My Enterprise** panels, nodes are denoted by this icon: 
- **Novell NRM**  
A system that is running version 6.5 of Novell Remote Manager (NRM), a Web-based interface to newer Novell NetWare servers. Novell NRM saves server statistics in an XML file. [up.time](#) can retrieve the XML file, parse it, and then store the information in the DataStore.
- **Net-SNMP v1/2**  
Systems that use either versions 1 or 2 of the Net-SNMP protocol. While both versions of Net-SNMP offer the same core features, version 2 contains additional protocol operations.
- **Net-SNMP v3**  
Systems that use version 3 of the Net-SNMP protocol to monitor and manage systems in a network that uses TCP/IP. Net-SNMP version 3 adds security features that are lacking in Net-SNMP versions 1 and 2.
- **Virtual Node**  
In a clustered environment, a device with which [up.time](#) can communicate using a floating IP address. In the **Global Scan** and **My Enterprise** panels, virtual nodes are denoted by this icon: 
- **VMware ESX v2**  
A system that is running version 2 of the VMware ESX server software, which enables a single host to run multiple virtual servers and their applications.

You will need to install an agent on the ESX v2 server in order to collect performance data.

- VMware ESX v3


A system that is running version 3 of the VMware ESX server software, which enables a single host to run multiple virtual servers and their applications. ESX version 3 includes features like the ability to balance the computing loads of a group of virtual servers as well as backup data and better manage clusters.

You do not need to install an agent on an ESX v3 server.

- pSeries LPAR Server (VIO)

A pSeries server that is hosting multiple logical partitions (LPARs). The VIO (virtual input/output) handles the physical I/O requests from the LPARs that are on the server.

In this configuration, [up.time](#) directly polls the agents installed on the VIO and the LPARs on a pSeries server for workload and other data, as illustrated below:



You will need to install an agent on each LPAR that you want to monitor. See “Installing Agents on IBM pSeries Servers” on page 46 for more information.



You can also add pSeries servers that are managed by a Hardware Management Console (HMC) to [up.time](#). You must do this using the Auto Discovery feature. See “Using Auto Discovery to Add pSeries Servers Managed by an HMC” for more information.

You can add multiple systems to [up.time](#) in a batch operation using a text file and a command line utility. See “Adding Multiple Systems to up.time” on page 434 for more information.

## Adding Systems or Network Devices

To add systems or network devices, do the following:

- 1 In the My Enterprise panel, click the Add System/Network Device tab.**

The **Add System/Network Device** window appears.

- 2 Enter a descriptive name for the server in the Display name in up.time field.**

This name will appear 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. This way, IP addresses are stored in [up.time](#) but a more descriptive or meaningful name is displayed in the [up.time](#) Web interface.

- 3 Optionally, enter a description of the system in the Description field.**
- 4 Select one of the following options from the Type of System/Device dropdown list:**

- Agent
- Node
- Novell NRM
- Net-SNMP v1/2
- Net-SNMP v3
- Virtual Node
- VMware ESX

- VMware ESX v3



An error message appears if you do not have a license that enables you to add VMware ESX servers, or if you have exceeded the number of servers permitted by your license.

- pSeries LPAR Server (VIO)

**5 Enter the host name of the system in the hostname field.**

The host name can be the actual name of the machine that **up.time** will be monitoring. You can also enter an IP address in this field.

**6 Optionally, enter the port number at which you will be connecting to the system in the port field.**

In most cases, you can use the default port.

**7 If you selected Agent in step 4 and want to securely access the system, click the Use SSL option.**

**8 If you selected Node in step 4, click the Is Node Pingable? option to specify whether or not **up.time** can contact the node using the ping utility.**

If, for example, you have a firewall in place you might not be able to ping the node. Before selecting this option, you should try to contact the node using the ping utility. If you cannot ping the node, uncheck this option. Then, change the default host check for the node. See “Changing Host Checks” on page 142 for more information.

**9 If you selected Novell NRM in step 4, enter information in the following fields:**

- Username

The user name that is required to access the Novell NRM Web interface.

- Password

The password that is required to access the Novell Web interface.

**10 If you selected Net-SNMP v1/2 in step 4, enter information in the following fields:**

- SNMP Port

The port on which the Net-SNMP instance is listening.

- Read Community

A string that acts like a user ID or password, giving you access to the Net-SNMP instance.

Common read communities are `public` (enables you to retrieve read-only information from the device) and `private` (enables you to access all information on the device).

**11 If you selected Net-SNMP v3 in step 4, enter information in the following fields:**

- SNMP Port

The port on which the Net-SNMP instance is listening.

- Username

The name that is required to connect to the Net-SNMP instance.

- Authentication Password

The password that is required to connect to the Net-SNMP instance.

- Authentication Type (optional)

From the list, select one of the following options, which will determine how encrypted information travelling between the Net-SNMP instance and [up.time](#) will be authenticated:

- MD5

A widely-used method for creating digital signatures used to authenticate and verify the integrity of data.

- SHA

A secure method of creating digital signatures. SHA is considered the successor of MD5 and is widely used with network and Internet data transfer protocols.

- Privacy Password

The password that will be used to encrypt information travelling between the Net-SNMP instance and [up.time](#).

- Privacy Type (optional)

From the list, select one of the following options, that determine how information travelling between the Net-SNMP instance and [up.time](#) will be encrypted:

- DES

An older method used to encrypt information.

- AES

The successor to DES, which is used with a variety of software that require encryption including SSL servers.



You can set both the authentication and password types, only one of them, or neither.

**12 If you selected VMware ESX v3 in step 4, enter information in the following fields:**

- User Name

The user name required to log into the VMware ESX v3 server.

- Password

The password required to log into the VMware ESX v3 server.

**13 If you want to associate this system with a group, select the name of the group from the Group dropdown list.**

See “Overview” on page 72 for more information on defining groups.

**14 If you want to associate this system with a service group, select the name of the group from the Service Group dropdown list.**

See “Service Groups” on page 139 for more information.

**15 Click Save.**

A window listing general information about the system you have added appears.

- 16 If you want to add another system or network element, click **Add Another**. Then, repeat steps 2 to 14.**

Otherwise, click **Close**.



It can take up to 15 minutes for the monitoring station to retrieve enough samples to provide historical graphing data to the monitoring station.

- 17 Click **Save**.**

## Auto Discovery

It can be time consuming to add large numbers of systems to [up.time](#) using the **Add System/Network Device** window, especially if you do not know the exact names or IP addresses of those systems. With Auto Discovery, [up.time](#) can detect the systems on your network that have an IP address within a range that you specify.

[up.time](#) does the following to search for the systems in your environment:

- Uses the ping utility to determine whether or not systems are available on the network.
- Checks whether or not the system has already been added to [up.time](#). If the system has been added, then the button to add the system is disabled.
- Scans the systems to determine whether or not the systems have agents installed on them.
- Performs an SNMP probe to find any systems that use Net-SNMP.

Once a list of systems in the range of IP addresses that you specified is generated, you can selectively add them to [up.time](#).

You can also use Auto Discovery feature to add VMware ESX v3 systems that are being monitored by Virtual Infrastructure 3, or pSeries systems that are managed by a Hardware Management Console (HMC). For more information, see the following sections:

- “Using Auto Discovery to Add ESX v3 Systems” on page 83.
- “Using Auto Discovery to Add pSeries Servers Managed by an HMC” on page 84.



## Using Auto Discovery

To use Auto Discovery, do the following:

- 1 In the **My Enterprise** panel, click the **Auto-Discovery** tab.

The **Auto Discovery** window appears.

- 2 In the **Please enter network** field, type the range of IP addresses that you want [up.time](#) to scan for agents.

For example, typing 10.1.1 will scan all system on your network that have an IP address starting with 10.1.1.

- 3 For the **Default SNMP read community** field (which contains a string that acts like a user ID or password, giving you access to the Net-SNMP instance), do one of the following:

- Accept the default value (public).
- Enter a new value, for example private.

- 4 Click **Continue**.

[up.time](#) returns a list of the systems that have an IP address within a range that you specified.

- 5 Click the **Add** button beside the system that you want to add.

The **Add System/Network Device** window appears.

- 6 If necessary, edit the details of the system as described in the section “Adding Systems or Network Devices” on page 78.

Otherwise, click **Save** in the **Add System/Network Device** window.

- 7 Repeat steps 4 and 5 for any other systems that you want to add.

## Using Auto Discovery to Add ESX v3 Systems

Virtual Infrastructure 3 (VI3; also called VirtualCenter) is a software suite that manages multiple, physical VMware ESX v3 servers. You cannot directly add VI3 systems to [up.time](#). You can, however, use the Auto Discovery feature to point [up.time](#) to a VI3 system and then add any or all of the ESX v3 servers that the VI3 system is managing.

To use Auto Discovery to add ESX v3 systems, do the following:

**1 In the My Enterprise panel, click the Auto-Discovery tab.**

The **Auto Discovery** window appears.

**2 Click the ESX v3 Discovery option.**

**3 Complete the following fields:**

- VI3 hostname

The name of the VI3 system.

- User Name

The user name required to log into the VI3 system.

- Password

The password required to log into the VI3 system.

**4 Click Continue.**

[up.time](#) returns a list of the ESX v3 servers that are being managed by the VI3 system.

**5 Click the Add button beside the system that you want to add.**

The **Add System/Network Device** window appears.

**6 If necessary, edit the details of the system as described in the section “Adding Systems or Network Devices” on page 78.**

Otherwise, click **Save** in the **Add System/Network Device** window.


**7 Repeat steps 5 and 6 for any other systems that you want to add.**

### Using Auto Discovery to Add pSeries Servers Managed by an HMC

The Hardware Management Console (HMC) is an interface for managing and configuring pSeries servers that are hosting multiple logical partitions (LPARs). When an HMC is attached to one or more pSeries servers with LPARs, the servers are considered *managed servers*.

In this configuration, the HMC manages all I/O requests from the LPARs. you use the Auto Discovery feature to detect the managed servers and add them to [up.time](#). Through the HMC, [up.time](#) polls the agents installed on the VIO

and the LPARs on a pSeries server for workload and other data, as illustrated below:



In order to monitor the managed servers and their LPARs, **up.time** must communicate with the HMC.



Before **up.time** can communicate with an HMC, you must enable SSH on the HMC. See “Enabling SSH on a Hardware Management Console” on page 447 for more information.

To use Auto Discovery to add pSeries servers that are managed by an HMC, do the following:

- 1 **In the My Enterprise panel, click the Auto-Discovery tab.**

The **Auto Discovery** window appears.

- 2 **Click the pSeries HMC Discovery option.**

- 3 **Complete the following fields:**

- HMC hostname

The name of the system on which the HMC is running.

- Username  
The user name required to log into the HMC.
- Password  
The password required to log into the HMC.

**4 Click Continue.**

[up.time](#) returns a list of the pSeries servers that are being managed by the HMC.

**5 Click the Add button beside the server that you want to add.**

The **Add System/Network Device** window appears.

**6 If necessary, edit the details of the system as described in the section “Adding Systems or Network Devices” on page 78.**

Otherwise, click **Save** in the **Add System/Network Device** window.

**7 Repeat steps 5 and 6 for any other systems that you want to add.**

## Adding VMware Instances to [up.time](#)

VMware ESX server software enables a single host to run multiple virtual servers and their applications. [up.time](#) can monitor both the server that is running VMware ESX, and VMware instances, which are the virtual servers that are running on the VMware server.

To add VMware instances to [up.time](#), do the following:

**1 In the My Enterprise panel, click the name of the VMware server that contains instances that you want to monitor.**

A new window containing information about the system appears.

**2 Click the Info tab, and then click VMware Instances.**

A list of VMware instances appears in the sub panel, as illustrated below:

| VMware Instances                |            |                  |        |
|---------------------------------|------------|------------------|--------|
| Display Name                    | IP         | Guest OS         | Is On? |
| <a href="#">DNS and FogBugz</a> | 10.1.1.140 | winNetEnterprise | Y      |
| Brightmail (W2K3 Enterprise)    | 10.1.1.6   | winNetEnterprise | Y      |
| Active Directory                | 10.1.1.19  | winNetEnterprise | Y      |
| ginger                          |            | linux            | N      |

**3 Click the Add to up.time button.**

The **Add System** window appears.



The **Add to up.time** button is not visible if a VMware instance is not on.

**4 If necessary, you can change any of the following options:**

- Display name in up.time
- Description
- Group
- Service Group

**5 Click Save to add the instance to up.time.**

## Adding Individual LPARs to up.time

After you have added pSeries servers – whether managed by an HMC or not – to [up.time](#), you can add individual LPARs from those systems to [up.time](#). While [up.time](#) collects workload data from all LPARs on a pSeries server (whether they have been added to [up.time](#) or not), adding LPARs can help you keep track of any specific LPAR.

To add an LPAR to [up.time](#), do the following:

**1 In the My Enterprise panel, click the name of the pSeries server that contains the LPAR that you want to monitor.**

A new window containing information about the system appears.

- 2 Click the **Info** tab, and then click **Logical Partitions**.

A list of LPARs appears in the sub panel.

- 3 Click the **Add to up.time** button beside the LPAR that you want to add to [up.time](#).

The **Add System** window appears.

- 4 If necessary, you can change any of the following options:

- Display name in up.time
- Description
- Group
- Service Group

- 5 Click **Save** to add the LPAR to [up.time](#).

## Editing a System Profile

After you have added a system to [up.time](#), you might need to change some of the basic information about that system. You can do this by editing the system profile.

To edit a system profile, do the following:

- 1 In the **My Enterprise** panel, click the name of the server whose profile you want to edit.

- 2 Click the **Info** tab in the **Tree** panel, and then click **Info & ReScan**.

- 3 In the subpanel, click **Edit System Profile**.

The **Edit System** window appears.

- 4 In the **Edit System** window, change any or all of the following options:

- Display name in up.time

The descriptive name for the system that appears in the [up.time](#) Web interface.

- **Description**  
A brief functional description of the system.
- **Parent Group**  
Select the group of systems in [up.time](#) with which this system will be associated.
- **Custom Field 1 to Custom Field 4**  
These fields enable you to include additional information about the system. For example, you can record the types of reports that should be run on this system, or when maintenance is scheduled.  
  
The information in the Custom Fields is displayed when you view system information by clicking the **Info & ReScan** link in the Tree panel.
- **Number of processes to retrieve**  
The default number of processes running on the system that [up.time](#) will retrieve. If you select 10 processes, and there are 20 running on the system, [up.time](#) retrieves the 10 busiest processes.
- **Is monitored?**  
Click this checkbox to turn monitoring off for this system. If monitoring is turned off, the system will not appear in the **Global Scan** panel.

5 **Click Save.**

## 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**  
This monitor acts as an anchor with which all other monitors will be associated.
- **Regular service monitors**  
The specific service monitors that are associated with a master service monitor.

When you enable a master service monitor, all of the regular service monitors associated with it are also enabled.

For more information on services, see “Using Service Monitors” on page 111. For information on viewing information about applications, see “Viewing Details About Applications” on page 92.

## Adding Applications

To add an application, do the following:

- 1 **In the My Enterprise panel, click the Add Application tab**
- 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 Groups dropdown list.**

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

For more information on groups, see “Working with Groups” on page 73.

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

- The name of a system to which the application will be assigned.
- **All**, which associates the application that you will specify with all of the systems 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 service monitor to which the application will be assigned.
- **All**, which associates the application with all of the service monitors 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.**

The name of the application appears in the **My Enterprise** panel. The name is a hyperlink. Click the name to view details about the application.

## 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 the master service monitor.

For information about viewing more details about applications, see “Viewing System and Service Information” on page 57.

## Editing Applications

To edit an application, do the following:

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

The **Application General Information** subpanel appears.

- 2 Edit the application setting as described in the section “Adding Applications” on page 90.

## Working with Views

Not every user in your enterprise needs to view all of the systems in your environment. Some users may, for example, only need be interested in five to 10 of the available servers. You can limit the servers that one or more users will see by creating specific *views* which are subsets of the servers in your environment. By creating views, it becomes easier for users to not only monitor systems, but to also browse and compare historical data.

Views appear in their own section of the **My Enterprise** panel. The following image illustrates a view:

| Views                                                                             |                                 |                 |                         |                                   |            |
|-----------------------------------------------------------------------------------|---------------------------------|-----------------|-------------------------|-----------------------------------|------------|
|                                                                                   | Name                            | Description     | Architecture            | OS version                        | Monitored? |
|                                                                                   | Linux                           |                 |                         |                                   |            |
|  | <a href="#">Cypher (cypher)</a> | A Linux box     | Linux cypher 2.4.18-... | RedHat 8.0(Psyche 2.4.18-14 i686) | ✓ Yes      |
|  | <a href="#">ELinux (elinux)</a> | A Linux box     | Linux elinux 2.6.9-5... | RedHat 4(Nahant 2.6.9-5.EL i686)  | ✓ Yes      |
|  | <a href="#">Ginger (ginger)</a> | Net-SNMP system | Linux Ginger 2.4.21-... | -                                 | ✓ Yes      |
|  | <a href="#">McKay (mckay)</a>   | Oracle server   | Linux McKay 2.4.21-4... | RedHat 3(Taroon 2.4.21-4.EL i686) | ✓ Yes      |
|  | <a href="#">Neo (neo)</a>       | A Linux box     | Linux csa3-rhes4 2.6... | -                                 | ✓ Yes      |

## Adding Views

To add a view, do the following:

- 1 In the **My Enterprise** panel, click the **Add View** tab.
- 2 In the **Add View** window, enter a descriptive name for the application in the **View Name** field.  
This name will appear in both the **My Enterprise** panel.
- 3 Optionally, enter a description for the application in **View Description** field.
- 4 Select one or more systems from the **Available Hosts for View** list, and then click **Add**.

If you have combined your systems into groups, select a groups from the dropdown at the top of the list. Or, select **All** from the dropdown to display all of the systems in your environment

- 5 Select one or more users from the **Available Users for View** list, and then click **Add**.
- 6 Click **Save**.

## Displaying and Editing Views

You can view the details of any view and edit those details.

To view and edit views, do the following:

- 1 In the My Enterprise panel, click the name of a view.

A window containing the information about the view appears, as illustrated below:

| View                              |                           | Edit View |
|-----------------------------------|---------------------------|-----------|
| Name:                             | Solaris Gateways          |           |
| Description:                      | Mail and Database Servers |           |
| Entities that belong to this view |                           |           |
| Entity Name                       | Description               |           |
| Filter (filter)                   | Main Solaris box          |           |
| Snoopy (snoopy)                   |                           |           |
| Subway (subway)                   | A Solaris box             |           |
| Users that can see these entities |                           |           |
| Username                          |                           |           |
| nsykes                            |                           |           |
| tslothrop                         |                           |           |

- 2 To edit the view, click the Edit View button.

The Edit View window appears. You can make changes to the view as described in the section “Adding Views” on page 93.

## Deleting Systems, Applications, and Views

If you have administrator privileges, you can delete a system or network device, an application, or a 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 **Click the Delete icon .**
- 3 **On the dialog box that appears, click OK.**

# Acknowledging Alerts

When a problem occurs on a system that **up.time** is monitoring, the monitoring station sends alerts – notifications that inform a user who can receive alerts – of the problem. If the user group to which you belong is configured to do so, you can also acknowledge an alert.

When you acknowledge an alert, **up.time**:

- Sends an acknowledgement message to any **up.time** user who received the last alert.
- Turns off alert escalation. **up.time** 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 system or application 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 system or application appears in the subpanel, as shown below:

| Status             |        |                                                                                     |                     |           |                                                                |
|--------------------|--------|-------------------------------------------------------------------------------------|---------------------|-----------|----------------------------------------------------------------|
| Monitor            | Status | Ack                                                                                 | Last Check          | Duration  | Monitor Information                                            |
| Perf Check - Happy | WARN   |  | 2006-01-13 09:32:12 | + 33s     | Swap: 92.0% >= 90%                                             |
| PING-happy         | OK     |  | 2006-01-13 09:32:27 | + 23h 54m | Ping completed: 1 sent, 0.0% loss, 0.0ms average round trip →  |
| UPTIME-happy       | OK     |  | 2006-01-13 09:32:42 | + 23h 52m | up.time agent running on happy, up.time agent 3.9 solaris 1. → |

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

The acknowledgement message window appears.

Acknowledge

Please enter the reason for Acknowledging the status:

Close

Submit

- 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 comment relating to the alert or why it was acknowledged.

The following is a sample alert acknowledgement message:

up.time Administrator (jsmith@myDomain.com)

acknowledged the WARN status of File System Capacity - Filter with comment:

Initial check of problem. More information to come.

In the [up.time](#) Web interface, the acknowledge icon changes to .





# CHAPTER 7

## Using Global Scan


.....

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

|                                       |            |
|---------------------------------------|------------|
| <i>Overview .....</i>                 | <i>100</i> |
| <i>Viewing All Applications .....</i> | <i>103</i> |
| <i>Viewing All Elements .....</i>     | <i>106</i> |
| <i>Viewing All Services .....</i>     | <i>108</i> |
| <i>Viewing a Resource Scan .....</i>  | <i>109</i> |

## Overview

The **Global Scan** panel enables you to view the current status of all of the applications, elements, services, and resources in your environment. The **Global Scan** panel contains a list of all the systems or devices that have been added to **up.time**, as shown below:



This list summarizes the following information:

- The status and number of services that are associated with a system or device.
- The number of recent service outages.
- CPU usage.
- Hard disk usage.
- Memory usage.

The bar chart in the bottom left of the panel displays the number of alerts that have been in Warning (yellow), Critical (red), and Unknown (grey)

states every 15 minutes over the past 24 hours. **up.time** takes a data sample from the database every 15 minutes and charts it on the bar chart. The number of services in each state appears in the graph.

The pie chart in the bottom right of the panel visualizes the availability of systems or devices over the last 24 hours. The services for unmonitored systems in groups are not shown in the pie chart.

You can view detailed information about a system by clicking on its name. To view the details of each metric – for example, CPU usage – click the number in the column for that variable. This displays the information in the **up.time** service or performance monitoring screen related to the variable.

When you click the file folder icon (  ) to the left of a system name, an expanded view of the server information appears. The following image illustrates the expanded view:



OK

WARN

CRIT

MAINT

UNKN

1hr

12hr

24hr

USR

SYS

WIO

TOT

% Used

% Busy

Swap Used

| <b>Cypher</b>               | 2       | 0      | 0    | 0    | 0   | 0    | 0      | 0                   | 0%      | 0% | 0% | 0% | 100% | 0% | 2% |
|-----------------------------|---------|--------|------|------|-----|------|--------|---------------------|---------|----|----|----|------|----|----|
| PING: cypher UPTIME: cypher |         |        |      |      |     |      |        |                     |         |    |    |    |      |    |    |
| Name                        | mem     | rss    | cpu% | mem% | pid | ppid | uid    | stime               | usr_sys |    |    |    |      |    |    |
| apmd                        | 1.30 MB | 416 KB | 0    | 0.1  | 738 | 1    | root   | 2006-05-31 08:05:00 | 0s      |    |    |    |      |    |    |
| atd                         | 1.34 MB | 496 KB | 0    | 0.1  | 864 | 1    | daemon | 2006-05-31 08:05:00 | 1s      |    |    |    |      |    |    |
| bdfush                      | 0 KB    | 0 KB   | 0    | 0    | 6   | 1    | root   | 2006-05-31 08:04:00 | 0s      |    |    |    |      |    |    |
| crond                       | 1.47 MB | 592 KB | 0    | 0.2  | 842 | 1    | root   | 2006-05-31 08:05:00 | 6s      |    |    |    |      |    |    |
| gpm                         | 1.34 MB | 356 KB | 0    | 0.1  | 833 | 1    | root   | 2006-05-31 08:05:00 | 0s      |    |    |    |      |    |    |

Messages

No messages in log or log not found.

**up.time** displays the following information for the system in the expanded view:

- The first row displays the names of the services, and their corresponding states, associated with the system.
- The second row lists the top five CPU consuming processes for the system.
- The third row displays the last five error messages (if any) for the system.

## Groups and Views in the Global Scan Panel

If you have created one or more groups (see “Working with Groups” on page 73) or views (see “Working with Views” on page 93), the names of the groups or views, and not the names of individual systems that make up them, appear in the **Global Scan** panel. The following information is displayed:

- The names and descriptions of the groups.
- The number of systems or elements in each group.
- The status of the hosts that make up the group.
- The number of alerts per group.

When you click on the name of 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 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, which act as anchors for all other monitors within the application.
- Regular service monitors, which are associated with a master service monitor.

When you enable a master service monitor, all of the regular service monitors associated with it are also enabled.

The status of each application is color coded, as follows:

- The applications highlighted in green are functioning normally.
- The applications highlighted in yellow are in a warning state.
- The applications highlighted in red are in a critical state. When one or more applications under the master application group is in a critical state, the critical icon (  ) appears beside the name of the master.

The color coding also indicates whether an application is offline or is in scheduled maintenance, as follows:

- An application that is offline is highlighted in red, and marked by the offline icon. As well, a message indicating that the application is offline appears in the **View All Applications** subpanel.
- An application that is in scheduled maintenance is grayed out, and the message **System is in scheduled maintenance** is displayed in the **View All Applications** subpanel. As well, the application is marked with the scheduled maintenance icon (  ).

The **View All Applications** subpanel displays the status of each application that you have added to [up.time](#). This subpanel has two views: Condensed View and Detailed View.

## Condensed View

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

| Application Status |                  |                      |                           | Show Detailed View         |
|--------------------|------------------|----------------------|---------------------------|----------------------------|
|                    | Application Name | Description          | Status of Master Services | Status of Regular Services |
|                    | Databases        | --                   |                           |                            |
|                    | J2EE             | Monitor J2EE servers |                           |                            |
|                    | Mail Servers     | --                   |                           |                            |


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 status of each service in the application

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:



| Application Status |                                     | Show Condensed View                                 |  |
|--------------------|-------------------------------------|-----------------------------------------------------|--|
|                    | Application Name                    | Description                                         |  |
|                    | Databases                           | --                                                  |  |
|                    | * UPTIME-sql2000.uptimesoftware.com | Default Uptime check for sql2000.uptimesoftware.com |  |
|                    | * UPTIME-mckay                      | Default Uptime check for mckay                      |  |
|                    | PING-mckay                          | Default Ping check for mckay                        |  |
|                    | PING-sql2000.uptimesoftware.com     | Default Ping check for sql2000.uptimesoftware.com   |  |
|                    | J2EE                                | Monitor J2EE servers                                |  |
|                    | * WebSphere monitor                 | --                                                  |  |
|                    | * WebLogic Monitor                  | --                                                  |  |

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 57.

## Viewing All Elements

Elements are the systems that **up.time** is currently monitoring. In the **Global Scan** panel, you can view the status of all elements being monitored in the **View All Elements** subpanel. The **View All Elements** subpanel is the default view in the **Global Scan** panel.

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

| Elements                                                                          |                                                                                   |                                                                                                    |                |      |      |       |      |         |      |      |                                                      |     |     |      |        |        |           |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|----------------|------|------|-------|------|---------|------|------|------------------------------------------------------|-----|-----|------|--------|--------|-----------|
|                                                                                   |                                                                                   | Name                                                                                               | Service Status |      |      |       |      | Outages |      |      | CPU                                                  |     |     | Disk |        | Memory |           |
|                                                                                   |                                                                                   |                                                                                                    | OK             | WARN | CRIT | MAINT | UNKN | 1hr     | 12hr | 24hr | USR                                                  | SYS | WIO | TOT  | % Used | % Busy | Swap Used |
|  |  |  Active Directory | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 1    | 0%                                                   | 1%  | 0%  | 1%   | 73%    | 0%     | 7%        |
|  |  |  AIX-SL           | 2              | 0    | 0    | 0     | 0    | 0       | 1    | 2    | 17%                                                  | 35% | 0%  | 72%  | 96%    | 3%     | 33%       |
|  |  |  Brightmail       | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 7%                                                   | 19% | 0%  | 26%  | 76%    | 5%     | 51%       |
|  |  |  Cypher           | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 0%  | 0%  | 0%   | 100%   | 2%     | 2%        |
|  |  |  ELinux           | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 0%  | 0%  | 0%   | 22%    | 0%     | 4%        |
|  |  |  Filter           | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | No performance data available. Is the System online? |     |     |      |        |        |           |
|  |  |  irobot           | 2              | 0    | 0    | 0     | 0    | 0       | 1    | 0    | 0%                                                   | 0%  | 0%  | 0%   | 95%    | 0%     | 23%       |
|  |  |  Matrix           | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 3%  | 0%  | 3%   | 96%    | 20%    | 36%       |
|  |  |  McKay            | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 3%  | 0%  | 3%   | 96%    | 3%     | 32%       |
|  |  |  Neo              | 2              | 0    | 0    | 0     | 0    | 0       | 1    | 0    | 0%                                                   | 11% | 0%  | 11%  | 92%    | 0%     | 1%        |
|  |  |  starqate         | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 2%  | 0%  | 2%   | 77%    | 11%    | 19%       |
|  |  |  startrek         | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 3%  | 0%  | 3%   | 92%    | 7%     | 82%       |
|  |  |  Subway           | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 1    | 1%                                                   | 1%  | 0%  | 2%   | 4%     | 0%     | 48%       |
|  |  |  swordfish        | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 0%                                                   | 64% | 0%  | 64%  | 86%    | 47%    | 80%       |
|  |  |  Tru64            | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 3%                                                   | 5%  | 0%  | 8%   | 68%    | 0%     | 15%       |
|  |  |  vmh-prod         | 2              | 0    | 0    | 0     | 0    | 0       | 0    | 0    | 33%                                                  | 0%  | 0%  | 33%  | 67%    | 0%     | 34%       |
|  |  |  WebLogic Server  | 2              | 0    | 1    | 0     | 0    | 0       | 0    | 0    | 15%                                                  | 2%  | 0%  | 17%  | 21%    | 7%     | 15%       |
|  |  |  WebSphere Server | 3              | 0    | 0    | 0     | 0    | 0       | 0    | 1    | 0%                                                   | 52% | 0%  | 52%  | 37%    | 2%     | 25%       |

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

- The names of the systems in your environment.
- The status of the services that are assigned to each system.
- 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 a system, application, or entity then the following message is displayed:

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 59 for more information.
- Click any value in the **Outages** column to open the **Outages** subpanel. See “Outages” on page 60 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 388 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 413 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 409 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. For more information on services, see “Services” on page 8.

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, for example `up.time agent running on MailServer, up.time agent 3.7.2 linux`.

## Viewing a Resource Scan

A 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 a resource scan by clicking on the **View Resource Scan** tab.

The Resource Scan is divided into two sections – a set of graphical gauges and a chart.

### Gauges

There are two sets of gauges, which look like car dashboard fuel or speed gauges. The gauges are updated every 30 seconds and 24 hours, respectively and illustrate 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.



### Chart

The Resource Scan chart displays the following information for all of the systems 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 |
| 10.1.1.8 | 2         | 95           | 1         | 6             | 1624         | 1279          |
| acidburn | 3         | 84           | 0         | 16            | 698          | 196           |
| elinux   | 13        | 95           | 0         | 19            | 28507        | 36434         |
| filter   | 0         | 49           | 0         | 56            | 497          | 484           |
| irobot   | 28        | 91           | 5         | 15            | 4823         | 4078          |
| subway   | 76        | 99           | 0         | 93            | 9            | 6             |

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:



# CHAPTER 8

## Using Service Monitors

.....

This chapter introduces the common features and concepts of [up.time](#) service monitors in the following sections:

|                                               |            |
|-----------------------------------------------|------------|
| <i>Overview .....</i>                         | <i>112</i> |
| <i>Using Service Monitors .....</i>           | <i>113</i> |
| <i>The Monitor Template.....</i>              | <i>117</i> |
| <i>Cloning Service Monitors.....</i>          | <i>127</i> |
| <i>Testing Service Monitors.....</i>          | <i>128</i> |
| <i>Alert Profiles .....</i>                   | <i>129</i> |
| <i>Action Profiles.....</i>                   | <i>134</i> |
| <i>Service Groups .....</i>                   | <i>139</i> |
| <i>Changing Host Checks .....</i>             | <i>142</i> |
| <i>The Platform Performance Gatherer.....</i> | <i>143</i> |
| <i>Topological Dependencies .....</i>         | <i>145</i> |
| <i>Scheduling Maintenance.....</i>            | <i>147</i> |


## Overview

A service monitor is an [up.time](#) process that checks the performance and availability of services in your environment at regular intervals. If the monitor detects a problem, [up.time](#) issues an alert.

Before you configure a service monitor, you should determine the following:

- The host name of the system that you want to monitor.
- When you want alerts to be sent.
- The action that will be taken to fix the problem.
- When the monitor should be run.

If you have tool tips enabled (see page 295 for more information), the graphic that appears in the Service Instances panel is a clickable image map.



Click any of the icons in the image to perform a task. For example, click the **Add Service Monitors to a system** icon to configure a new service monitor.

## Using Service Monitors

There are three main types of service monitors:

- Agent Monitors  
For more information about Agent Monitors, see “Using Agent Monitors.”
- Agentless Monitors  
For more information about Agentless Monitors, see “Using Agentless Monitors.”
- Custom Monitors  
For more information about custom monitors, see “Using Advanced Monitors.”

## Using Agent Monitors

To use agent monitors, [up.time](#) requires:

- An [up.time](#) agent to be installed and running on the system on which the service that you want to monitor is running.
- The service about which you want to collect information to be installed and running on the system that you intend to monitor.

Agents enable you to collect very detailed data 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.

The monitors that require an agent are:

- |                        |                                 |
|------------------------|---------------------------------|
| • Exchange             | • SQL Server (Advanced Metrics) |
| • File System Capacity | • Uptime Agent                  |
| • IIS                  | • Windows Event Log Scanner     |
| • Performance Check    | • Windows Service Check         |
| • Process Count Check  |                                 |

### Using Agentless Monitors

Agentless monitors do not require an [up.time](#) agent to be installed and running on the system that you want to monitor. Your monitoring station communicates with the remote system to:

- Determine the status of the service that is being monitored.
- Collect information from the service that is being monitored.

The monitors that do not require an agent are:

- |                             |                                 |
|-----------------------------|---------------------------------|
| • Active Directory          | • Ping                          |
| • DNS                       | • POP (Email Retrieval)         |
| • FTP                       | • SSH (Secure Shell)            |
| • HTTP (Web Services)       | • SMTP (Email Delivery)         |
| • IMAP (Email Retrieval)    | • SNMP                          |
| • LDAP                      | • SQL Server (Advanced Metrics) |
| • MySQL (Advanced Metrics)  | • SQL Server (Basic Checks)     |
| • MySQL (Basic Checks)      | • SQL Server Tablespace Check   |
| • NFS                       | • Sybase                        |
| • NIS/YP                    | • TCP                           |
| • NNTP (Network News)       | • WebLogic                      |
| • Oracle (Advanced Metrics) | • WebSphere                     |
| • Oracle (Basic Checks)     | • ESX v3 Workload               |
| • Oracle Tablespace Check   | • Windows File Shares (SMB)     |

### Using Advanced Monitors

You can configure monitors to carry out service or performance checks that may be specific to your environment. Using advanced monitors, you can:

- Monitor any service that does not have an [up.time](#) service monitor.
- Monitor the performance of elements in your environment.



- Perform common database administration tasks.

For more information, see “Advanced Monitors” on page 279. Contact uptime software Client Care for assistance with configuring advanced monitors.

## Types of Advanced Monitors

There are three advanced monitors:

- Custom

Monitors that return the status of a monitor and an automated message to clarify the returned status.

- Custom with Retained Data

Monitors that return the following:

- Up to 10 values that you can capture and can evaluate.
- A return status.
- A message.

You can also configure these monitors to save data to the database, which you can use to generate a Service Metrics report (see “Service Monitor Metrics Report” on page 329) or a Service Metrics graph (see “Viewing System and Service Information” on page 57).

- External Check

Monitors that rely on an external event to trigger the capture of service information. External check monitors enable you to determine when to collect service data based on an external application event that you specify.

For more information on configuring and using advanced monitors, see “Advanced Monitors” on page 279.

## Selecting a Monitor

To select a monitor, do the following:

- 1 Click **Services** on the [up.time](#) tool bar.
- 2 Click **Add New Service Instance** in the **Tree** panel.

The **Add Service Monitor** window appears.

- 3 Select one of the monitors in the monitors that is listed in the window, and then click **Continue**.

See “The Monitor Template” on page 117 for information on completing the configuration of a custom monitor.

## The Monitor Template

You use a general template to configure monitors. While the specific configuration information varies from monitor to monitor, every template contains areas for:

- Monitor Identification
- Monitor Settings Configuration
- Monitor Timing Settings
- Monitor Timing Settings
- Monitor Alert Settings
- Alert Profiles
- Action Profiles

### Monitor Identification

Each service monitor template has a monitor identification information area that you use to:

- Specify the name of the monitor.
- Include an optional description of the monitor.
- Select the system, node, or virtual node that you want [up.time](#) to monitor.

The monitor identification information area is illustrated below:

|              |                                                                                                                                                                            |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service Name | <input type="text"/>                                                                                                                                                       |
| Description  | <input type="text"/>                                                                                                                                                       |
| Host         | <input checked="" type="radio"/> Single System <input type="text" value="- Select a System -"/><br><input type="radio"/> Service Group<br><input type="radio"/> Unassigned |

You must ensure that the system can be resolved by a naming service running on an operating system – for example, DNS or NIS/YP.

## Adding Monitor Identification Information

To add monitor identification information, do the following:

**1 Enter a name for the monitor in the Service Name field.**

The name can, for example, describe the purpose of the monitor – for example, Ping - Web Server.

**2 Optionally, enter a description of the monitor in the Description field.**

**3 Assign the monitor to a system by doing one of the following:**

- Click the **Single System** option, and then select the name of the system that you want to monitor from the dropdown list.
- Click **Service Group** to attach the monitor to multiple systems. Then, select the service group from the dropdown list. For more information about service groups, see “Service Groups” on page 139.
- Click the **Unassigned** option.


This step is mandatory.

**4 Complete the following fields:**

- Port  
The number of the port on which [up.time](#) is listening.
- Use SSL  
Select this option if the [up.time](#) agent is configured to use SSL (Secure Sockets Layer) for security.  
  
If you have configured your agent to use SSL but do not select **Use SSL**, [up.time](#) will not receive performance information.

## Monitor Settings Configuration

Each up.time service monitor has settings particular to the service that is it monitoring. The following image illustrates a setting from a MySQL (Basic Checks) monitor:




| IP Address ▾ |                                     | <input type="checkbox"/> |
|--------------|-------------------------------------|--------------------------|
| Warning      | Comparison Method ▾<br>[Text Input] |                          |
| Critical     | Comparison Method ▾<br>[Text Input] |                          |

Most monitors have the following views:

- **Standard View**  
Displays the basic settings that are used to configure the monitor. Not all settings for a monitor are mandatory. You only need to set the host name in the monitor identification information section.
- **Advanced View**  
Displays standard view settings plus additional settings that are used to capture more specific or more advanced service metrics from a monitor. A common Advanced View setting is **Response Time**, which determines the total amount of time that a service check takes to gather information from a service.

With certain monitors, such as SNMP, the Standard and Advanced views are same.

You can change the view by selecting an option from the **View** dropdown list, as shown below:



**Add Service - Windows Service Check - Page 2 of 2**

Service Name: [Text Input]

Description: [Text Input]

Host: ☒ Single System ☐ Service Group ☐ Unassigned


- Select a System - ▾

View: Standard View ▾  
Standard View  
Advanced View

If you create a monitor with Advanced view settings, the default view for that monitor will be Advanced when you later edit the monitor.

## Comparison Methods

You can configure settings that compare the Warning and Critical threshold values that you have set to the values that **up.time** captures. **up.time** issues an alert when these thresholds are exceeded. You choose a comparison methods from the **Select a comparison method** dropdown list, as shown below:



After selecting a comparison method, you enter a value in field beside or below the dropdown list.

The following are the available comparison methods:

- exactly matches  
The string returned by the monitor exactly matches the string that you defined.
- does not match  
The string returned but the monitor does not match the string that you defined.
- regular expression  
The string returned by the monitor exactly matches the pattern result of a regular expression that you define.
- inverse regular expression  
**up.time** accepts all of the patterns that do not correspond to the regular expression string that you define. For example, specifying `[^OK]` will return all results that are not in an OK state.
- contains  
The string returned by the monitor contains the string that you defined.
- does not contain  
The string returned by the monitor does not contain the string that you defined.

If you select a method from the dropdown list and either enter an incorrect value in the field or do not enter a value, then an error message appears and you cannot save the monitor. If you do not want to specify a comparison value, do not select an option from the **Select a comparison method** dropdown list.

## Configuring Warning and Critical Thresholds

In many instances, you must configure Warning and Critical thresholds to determine the conditions under which **up.time** issues an alert. For example, if hard disk usage on a server reaches 85% **up.time** issues a Warning alert. If disk usage reaches 95%, **up.time** issues a Critical alert.

To configure Warning and Critical thresholds, do the following:

- 1 **Enter the threshold value in the text box next to the Select a comparison method dropdown list.**
- 2 **Select an option from the Select a comparison method dropdown list.**

### Response Time

Response time is an Advanced view setting that denotes the amount of time that a monitor requires to:


- Initiate a service check.
- Transmit a request to a local or remote system, or to a service.
- Collect service information.
- Return the collected information to the monitoring station.
- Display the information on the monitoring station.

Many factors can influence the response time including network connectivity, the type of information that is being collected, and the availability and performance of the service.

### Configuring Response Time

To configure response time, do the following:

- 1 **For each threshold, select an option from the Select a comparison method dropdown list as illustrated below:**



The screenshot shows a configuration panel titled "Response time" with a dropdown arrow. It contains two rows: "Warning" and "Critical". Each row has a dropdown menu labeled "Select a comparison method" and a text input field followed by "ms".

- 2 **Enter a Warning threshold, in milliseconds.**

For information on configuring Warning thresholds, see “Configuring Warning and Critical Thresholds” on page 121.

- 3 **Enter a Critical threshold, in milliseconds.**

For information on configuring Warning thresholds, see “Configuring Warning and Critical Thresholds” on page 121.



If you select a comparison method, you must enter a value in the corresponding field for the threshold.


## Monitor Timing Settings

Monitor timing settings determine:

- Whether or not the monitor is active.
- The length of time, in seconds, to wait before determining that a monitor has timed out.
- The interval, in minutes, at which the monitor will perform a service check.
- The interval, in minutes, at which the monitor will recheck the status of a service.
- The maximum number of times that the monitor will recheck a service.



The following image illustrates the **Monitor Timing Settings** area of the monitor template:



| Timing Settings                      |                  |
|--------------------------------------|------------------|
| <input checked="" type="checkbox"/>  | Monitored        |
| <input type="text" value="60"/> sec. | Timeout          |
| <input type="text" value="10"/> min. | Check Interval   |
| <input type="text" value="1"/> min.  | Recheck Interval |
| <input type="text" value="3"/>       | Max Rechecks     |

The monitor timing settings enable you to set up a master service monitor that you can apply to multiple systems. You can do this when setting up a deployment where you may want to apply a service monitor to a large number of entities, or want to apply a very similar service monitor and then make further customizations to it and its children.

## Timing Settings Options

The following options are available in the **Timing Settings** area:

- **Monitored**  
Turns a monitor on or off. The **Monitored** setting is on by default.
- **Timeout**  
How long a monitor runs before [up.time](#) issues an error message. A timeout occurs when the monitoring station has not received a status from the named service monitor after a period of time has passed. When a service monitor does not return data, the status of the monitor changes to Unknown. When a service monitor times out, an error message appears on the **Global Scan** panel.
- **Check Interval**  
How frequently the monitor checks the status of an entity. The minimum check interval is one minute, and the default is 10 minutes. There is no maximum check interval.
- **Re-Check Interval**  
The amount of time between checks. A recheck should occur when a monitor has gone from an OK to a Warning, Critical, or Unknown status.

The duration for rechecks should be shorter than the regular check interval. The minimum recheck interval is one minute.



Rechecks continue to run as they are needed until the maximum number of rechecks has occurred.

- **Max Rechecks**

The maximum number of times that [up.time](#) rechecks a service. Once the specified number of rechecks is completed, the last state that was checked is reported. If the last status was not OK, [up.time](#) generates an alert.

### Adding Monitor Timing Settings Information

To add monitor timing settings information, do the following:

- 1 Select the Monitored check box to activate the service monitor.**



[up.time](#) does not send alerts if the service monitor is not activated.

- 2 Complete the following settings:**

- Timeout.



Ensure that the **Timeout** duration that you define is longer than the Response Time that you defined in the service monitor's Advanced view.

- Check Interval.
- Recheck Interval.
- Max Rechecks.

## Monitor Alert Settings

The monitor alert settings enable you to turn alert notifications on or off based the status of a service monitor. The following options are available in this area:

- **Notification**  
Determines if notifications, regardless of status or interval, should be issued for this monitor.
- **Alert Interval**  
The frequency, in minutes, at which alerts are issued. The default is 120 minutes.
- **Alert on Critical**  
Sends an alert when a monitor reaches a Critical status threshold.
- **Alert on Warning**  
Sends an alert when a monitor reaches a Warning status threshold.
- **Alert on Recovery**  
Sends an alert when a monitor recovers from a Warning or Critical status.
- **Alert on Unknown**  
Sends an alert if any metric or time value for a monitor returns a status of Unknown.

## Adding Monitor Alert Settings Information

To add monitor alert settings information, do the following:

- 1 **Click the Notification check box to turn on alert notifications.**



If you do not click the **Notification** check box, none of the remaining boxes in monitor alert settings template are active.

- 2 **Enter an amount of time, in minutes, in the Alert Interval field**

The alert interval is the frequency at which an alert is repeated if a monitor does not have an OK status.

### 3 Click one or more of the following checkboxes:

- Alert on Critical.
- Alert on Warning.
- Alert on Recovery.
- Alert on Unknown.

## Monitoring Period Settings

The monitoring period settings determine the time periods at which [up.time](#) sends alerts. For more information on alerts, see “Understanding Alerts” on page 13 and “Alert Profiles” on page 129.

To set the monitoring period, do the following:

### 1 Select one of the following options from the Monitoring Period dropdown list to specify when alerts can be sent:

- 24x7
- 9 am to 5 pm weekdays
- 5 pm to 7:30 am weekdays and all weekend until Monday morning
- 12am to 12:30am Monday

## Getting Additional Help

If you need more information about certain fields on the monitor template, hold your mouse over the inverted chevron ( ▼ ) beside the name of the field. A tool tip that describes the field will be displayed.

## Cloning Service Monitors

Cloning a service monitor makes a copy of the service monitor and all of its parameters. Cloning a service monitor is useful if, for example, you want to use similar monitors for several servers in your environment.

To clone service monitors, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Service Instances** subpanel, click the **Clone** icon ( )


## Testing Service Monitors

You can test that a service monitor is functioning and collecting data properly to ensure that the configuration is correct. If the configuration is not correct, then you can immediately fix any configuration errors before they become a problem.

To test a service monitor, do the following:

- 1 **On the up.time toolbar, click the Services tab.**
- 2 **In the navigation menu, click View Service Instances.**  
A list of available service monitors appears in the sub panel.
- 3 **Click the name of the service monitor that you want to test.**
- 4 **Click the Test Service Instance button.**


A pop-up window appears, containing the status of the monitor and a message related to the status. The following image illustrates such a message:



- 5 **When finished, click the Close Window button.**

## 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:



For more information on alert profiles, see “Understanding Alert Profiles” on page 15. For an example of an alert flow, see “Understanding the Alert Flow” on page 14.

## Creating Alert Profiles

To create alert profiles, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **Add New 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

Sends the alert via SMS to the mobile phones of the members of a notification group.

This alert option relies on a script or batch file. You enter the name and path to the script in the **Script Path** field, for example `/usr/local/uptime4/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.

## 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. For more information on variables, see “Custom Alert Format Variables” on page 450.




To define a custom alert format, do the following:

- 1 **Define an alert profile, as described on page 129.**
- 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 page 450) 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.**

## Viewing Alert Profiles

To view alert profiles, do the following:

- 1 On the **up.time** tool bar, click **Services**.
- 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 **Services**.
- 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 129.

## Applying Alert Profiles to Applications

You can apply an alert profile to an application to generate an alert if the state of the application changes from OK to Warning or Critical. For more information on applications, see “Working with Applications” on page 90.

To apply alert profiles to applications, do the following:

- 1 In the **Global Scan** or **My Enterprise** panels, click the name of the application to which you want to apply an alert profile.

The **Service Instance** page appears.

- 2 Click the **Services** tab and then click **Manage Services**.

- 3 In the subpanel, click **Edit Alert Profiles**.

The **Edit Alert Profiles** window appears.

- 4 Select one or more alert profiles from the **Available Alert Profiles** list and then click **Add**.

## Action Profiles

Action profiles are templates that tell [up.time](#) how to deal with a problem that it encounters on a system that is being monitored. The action that [up.time](#) takes could be to record the problem in a log file or to execute a script that attempts to correct the problem. For more information on action profiles, see “Understanding Action Profiles” on page 16.

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 **Services**.
- 2 In the Tree panel, click **Add New 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 In the **Log File** field, enter the name and path to a log file on the monitoring station to which error information will be written.
- 7 In the **Recovery Script** field, enter the name and path to a script that will restart an application, process, or service; or reboot a server.

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/uptime4/recover.sh" "22/3/2006 12:01:05"
"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.

**8 If you are setting up an action profile for a Windows server, you can also 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.
- 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

- up.time Agent Password

Enter the password that is required to access the agent that is running on the system that is being monitored. For information on setting the agent password, see “Changing the Agent Password on Windows” on page 449.

- Use SSL

Select this option if [up.time](#) will securely communicate with the host using SSL (Secure Sockets Layer).

**9 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.

- 10 If the action profile will work with the [up.time](#) Splunk monitors, 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.

Click the **Use SSL** option to securely access the Splunk server using SSL.

For more information on the Splunk monitors, see “Splunk Query” on page 196 and “Live Splunk Listener” on page 198.

- 11 Click **Save**.

## Viewing Action Profiles

To view action profiles, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **View Action Profiles**.

The **Action 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 works, click the **Test Action Profile** button.

A popup window appears, and the 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 **Services**.
- 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 134.

## Applying Action Profiles to Applications

You can apply an action profile to an application. For more information on applications, see “**Working with Applications**” on page 90.

To apply action profiles to applications, do the following:

- 1 In the **Global Scan** or **My Enterprise** panels, click the name of the application to which you want to apply an action profile.

The **Service Instance** page appears.

- 2 Click the **Services** tab and then click **Manage Services**.
- 3 In the subpanel, click **Edit Action Profiles**.

The **Edit Action Profiles** window appears.

- 4 Select one or more action profiles from the **Available Action Profiles** list and then click **Add**.



## Service Groups

Service groups are monitor templates that enable you to simultaneously apply a common service check to one or more hosts that you are monitoring. Defining and using service groups can simplify the setup and maintenance common service checks that you want to perform across multiple hosts. When adding a host to [up.time](#), you assign a service group to it instead of manually adding service checks.

For more information, see “Understanding Service Groups” on page 24.

### Creating Service Groups

To create service groups, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the Tree panel, click **Create New Service Group**.  
The **Add Service Group** window appears.
- 3 Enter a descriptive name for this group in the **Name of Service Group** field.
- 4 Optionally, enter a description of the group in the **Description** field.
- 5 Click **Continue**.
- 6 On the second **Add Service Group** screen, select one of the following options from the **Available Services** dropdown list.

- All

View all of the services that are available.

- The name of a host

If you are monitoring large number of systems, this option enables you to filter the services based on the hosts that you have added to [up.time](#).

- 7 Select one or more services from the list, and then click **Add**.

**8 Select one of the following options from the Available Hosts dropdown list:**

- All

View all of the hosts that have been added to [up.time](#).

- The name of a group

If you have grouped your hosts, use this option enables you to filter the hosts based on the groups that you have added to [up.time](#). The names of the hosts in the group appear below the dropdown list.

If you have hosts that are not members of a specific group, select **My Enterprise** from the dropdown list to view the ungrouped hosts. If you have not created groups, the dropdown list is not available and a list of hosts appears in the list.

See “Working with Groups” on page 73 for more information about grouping hosts.

**9 Select one or more hosts from the list, and then click Add.**

**10 Click Finish.**

## Editing Service Groups

To edit service groups, do the following:

**1 On the [up.time](#) tool bar, click Services.**

**2 In the Tree panel, click View Service Groups.**

**3 Click the Edit icon (  ) beside the name of the service group that you want to edit.**

**4 To change the name and description of the group, do the following:**

- Enter a new name in the **Name of Service Group** field.
- Enter a new description of the Service Group in the **Description** field.
- Click **Save**.

- 5 To edit the services in the service group, do the following:
  - Add services by clicking on one or more services in the **Available Master Services** list, and then clicking **Add**.
  - Remove services by clicking on one or more services in the **Selected Master Services** list, and then clicking **Remove**.
  - Click **Save**.
- 6 To edit the hosts in the group, do the following:
  - Add systems by clicking on one or more systems in the **Available Hosts** list, and then clicking **Add**.
  - Remove systems by clicking on one or more systems in the **Selected Hosts** list, and then clicking **Remove**.
  - Click **Save**.

## Changing Host Checks

Host checks determine whether or not a system that is being monitored is available and functioning properly. If a host check determines that a host is unavailable, then all service checks are temporarily disabled.

The available host checks are:

- Ping check  
This host check uses the ping utility to determine whether or not the server is accessible. This is the default host check.
- [up.time](#) agent check  
This host check communicates with the [up.time](#) agent installed on a system to determine whether or not the system is functioning.
- Any service monitors that you have configured for a system.

### Change a Host Check

To change a host check, do the following:

- 1 **On the [up.time](#) tool bar, click Services.**
- 2 **In the Tree panel click Host Check.**  
A list of the servers and their assigned host checks appears in the subpanel.
- 3 **Click the Edit icon (  ) beside the name of the server whose host check you want to change.**  
A list of the available host checks appears in a new window.
- 4 **Select a host check, and then click Save.**

## The Platform Performance Gatherer

The Platform Performance Gatherer is a host check that collects basic performance metrics – for example, CPU performance and disk statistics – from a system in order to determine whether or not that system is functioning. You can edit the following monitor settings for the Platform Performance Gatherer from the **Info & Rescan** subpanel.

### Editing the Platform Performance Gatherer

To edit the Platform Performance Gatherer settings, the following:

- 1 In the **Global Scan** or **My Enterprise** panels, click the name of a server.
- 2 Click the **Info** tab, and then click **Info & Rescan**.
- 3 Click the **Edit Performance Monitor** link that is beside the **Monitoring Interval** setting, as shown below:

|                     |                                                  |
|---------------------|--------------------------------------------------|
| Monitoring Interval | 5 min - <a href="#">Edit Performance Monitor</a> |
|---------------------|--------------------------------------------------|

The **Edit Service Monitor** window appears.

- 4 **Edit the settings for the Platform Performance Gatherer.**

While you can edit any setting, the settings that you are most likely to change are:

- Port Number

The number of the port on which the Platform Performance Gatherer is collecting data from a host.

For most systems, this setting is labelled **Agent Port Number**. For systems running Net-SNMP this setting is labelled **SNMP Port**, and for Novell NRM (version 6.5) systems this setting is labelled **Novell NRM Port Number**.

- User Name and Password


For Novell NRM systems, the user name and password that are required to access the system.

- **Username**  
The name that is required to connect to the instance of Net-SNMP v3.
- **Authentication Password**  
The password that is required to connect to the instance of Net-SNMP v3.
- **Authentication Method**  
The method by which encrypted information travelling between the Net-SNMP instance and [up.time](#) will be authenticated.
- **Privacy Password**  
The password that will be used to encrypt information travelling between the instance of Net-SNMP v3 and [up.time](#).
- **Privacy Type**  
The method by which information travelling between the instance of Net-SNMP v3 and [up.time](#) will be encrypted.
- **Use SSL (HTTPS)**  
Select this option if the Platform Performance Gatherer will securely communicate with the host using SSL (Secure Sockets Layer).
- **Check Interval**  
The frequency, in minutes, at which the host will be checked.  
  
If the Check Interval is longer than the Alert Interval, the following message appears:  
  
Warning: The alert interval is less than the check interval. up.time will only send alerts after performing checks

**5 Click Save.**

## Topological Dependencies

In large deployments, a single system or node can act as the gateway to other entities or entity groups. For example, up.time might need to go through a router – configured as a node in up.time – to monitor one or more systems that are behind the node. This situation is illustrated below:



### Systems being monitored

If the router fails, then up.time generates alerts for the systems behind the routers because the service monitors cannot communicate with those systems.

Topological dependencies create parent-child relationships between systems. Both entities and entity groups can be dependent on a parent system or node.

A service monitor can determine that systems which are dependent on a specific system or node that is experiencing a problem will be unavailable until the problem is resolved. Alerts will not be generated. However, the checks for the dependent systems will continue to be scheduled.



If a topological parent is down, a descriptive message appears in the **Global Scan** panel for entities and services that are children of that parent.

## Adding Topological Dependencies

To add topological dependencies, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **Add Topological Dependency**.  
The **Add Topological Dependency** window appears.
- 3 Select a system from the **Select a host to create dependencies for dropdown list**.

This host acts as the parent for the dependent systems or nodes. If [up.time](#) cannot communicate with the host, then the service monitors that check the dependent systems or nodes will not run host checks.

- 4 Click **Continue**.
- 5 Select one or more systems or nodes from the **Available Dependent Hosts dropdown list**.

These systems or nodes will be the dependents of the host system that you specified in step 3.

- 6 Optionally, select one or more entity groups from the **Available Dependent Groups dropdown list**.

These groups will be the dependents of the host system that you specified in step 3.

- 7 Click **Finish**.

## Viewing Topological Dependencies

To view topological dependencies, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **View Topological Dependencies**.

The subpanel displays the following dependency information:

- Name of the parent.
- The number of dependent hosts.
- The number of dependent groups (if any).



## Scheduling Maintenance

Scheduled maintenance is a period during which the monitoring station does not monitor a host or service. You can schedule maintenance if, for example, you back up a system at a specific time each day or week, or if a system must be taken down for an upgrade. When a host or service is scheduled for maintenance, the monitoring station assumes that the host or service cannot be contacted but does not issue an alert.

If maintenance is not scheduled, then during those periods [up.time](#) will notify you that the system or service is unavailable when systems or services are not online.

### Creating Scheduled Maintenance Profiles

You can schedule maintenance using *profiles*. A scheduled maintenance profile is a template that enables you to define maintenance periods, and then assign the profile to multiple systems. A profile is a recurring event – for example, a backup cycle that occurs every Monday between 3 a.m. and 5 a.m.

To create scheduled maintenance profiles, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **View Maintenance Profiles**.
- 3 In the **Services** subpanel, click the **Add Maintenance Profile** tab.
- 4 Enter a descriptive name for the profile in the **Profile Name** field.
- 5 In the **Time Range** section, select the start and end times for the scheduled maintenance period from the dropdown lists.

For example, to schedule maintenance to run from 1:30 a.m. to 3:00 a.m., select 01 and 30 from the **Start** dropdown lists and then select 03 and 00 from the **End** dropdown lists, as shown below:



| Start |    | End |    |
|-------|----|-----|----|
| 01    | 30 | 03  | 00 |

### 6 Select one of the following options:

- One Time

This option creates a maintenance profile that is used once and then discarded. Select the month, day, and year on which this profile will apply from the dropdown lists.

- Weekly

Select one or more of the days of the week for which this scheduled maintenance profile applies.

- Monthly

Do one of the following:

- Select the **Day** option, and then select the day (from 1 to 31) for which this scheduled maintenance report applied.

For example, if you select 3 from the dropdown list, the profile will take effect on the third day of every month.

- Select the second option, and then:

- Select first, second, third, fourth, or last from the first dropdown list.
- Select a day of the week on which the scheduled maintenance profile will start from the second dropdown list.

For example, if you select second and Tuesday from the dropdown lists, the profile will take effect on the second Tuesday of every month.

- Yearly

Select a month from the first dropdown list, and then select a day from the second dropdown list.

### 7 Click **Save**.

## Viewing Scheduled Maintenance Profiles

You can view scheduled maintenance profiles to ensure that they meet your needs and that they are applied to the appropriate hosts and services.

To view scheduled maintenance profiles, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **View Maintenance Profiles**.
- 3 In the **Services** subpanel, click the name of the maintenance profile that you want to view.

The scheduled maintenance profile appears in the **Services** subpanel, and contains the following information:

- The name of the profile.
- The time period over which the profile is applied to a system or service.
- The names of the systems and services, if any, to which the profile has been applied.

## Scheduling Maintenance for a Host

To schedule maintenance for a host, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **Host Maintenance Windows**.
- 3 Click the **Assign Maintenance to Host** tab in the subpanel.
- 4 In the **Host Maintenance** window, select the maintenance profile to use from the **Maintenance profile** dropdown list.

If you have not created a maintenance profile, the message No profiles exist appears in the dropdown list.

- 5 Select one or more systems from the **Available Host** list.

The hosts that you select will be the hosts to which the maintenance profile applies.

- 6 Click **Add**, and then click **Save**.

## Scheduling Maintenance for a Service

To schedule maintenance for a service, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the **Tree** panel, click **Service Maintenance Windows**.
- 3 Click the **Assign Maintenance to Service** tab in the subpanel.
- 4 In the **Service Maintenance** window, select a profile from the **Maintenance profile** dropdown list.

If you have not created a maintenance profile, the message No profiles exist appears in the dropdown list.

- 5 Optionally, from the dropdown list above the **Available Service** list, select a system that contains the services for which you want to schedule maintenance.
- 6 From the **Available Service** list, select one or more services for which you want to schedule maintenance.
- 7 Click **Add**, and then click **Save**.

# CHAPTER 9

## Agent Monitors

.....

The agent monitors track the performance and health of the following:

|                                   |            |
|-----------------------------------|------------|
| <i>File System Capacity .....</i> | <i>153</i> |
| <i>Performance Check .....</i>    | <i>155</i> |
| <i>Process Count Check .....</i>  | <i>159</i> |

# Overview

Agent monitors are service monitors that require an agent to be installed on the system being monitored. An agent is software that collects performance information from the system and transmits that information to the monitoring station. Using the information gathered by an agent, [up.time](#) can alert users to changes in an environment based on defined thresholds.



For information on installing agents, see “Installing Agents” on page 43.

## File System Capacity

The File System Capacity monitor checks the amount of total and used space, in kilobytes, on a disk. This monitor then compares the capacity to the specified warning and critical thresholds. 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 individual file systems (for example, /var, /export, /usr) on all of the disks on the server.

On UNIX and Linux systems, you can configure the monitor to check all of the mount points on a system, or just specific mount points.



This monitor does not check floppy drives, tapes drives, or CD-ROM drives.

## Configuring File System Capacity Monitors

To configure File System Capacity monitors, do the following:

**1 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

**2 Complete the following fields:**

- Global Warning Threshold (Mandatory)

Enter the percentage of the file system that must be used for [up.time](#) to generate a warning.

- Global Critical Threshold (Mandatory)

Enter the percentage of the file system that must be used for [up.time](#) to generate a critical alert.

**3 Optionally, to exclude specific mount points on the disk from the capacity calculations enter the names of the mount points in any or all of the five the Exclude Pattern fields.**

For example, you can enter D: (for Windows) or /usr (for Solaris, Linux, or AIX) to ignore that drive or directory. To, for example, ignore all mount points that start with /u enter /u\*.

**4 Optionally, you can set thresholds for specific mount points by entering the following information in any or all of the five Mount Point fields:**

- The name of the mount point, for example /opt. The thresholds that you set for each mount point will be calculated separately from the thresholds that you specified in step 2.
- The **Warning** threshold, which is percentage of space used on the mount point that when exceeded generates a warning.
- The **Critical** threshold, which is the percentage of space used on the mount point that when exceeded generates a critical alert.

**5 Specify values for the Warning and Critical Response Time thresholds.**

For more information, see “Response Time” on page 121.

To save the data from the thresholds for graphing or reporting, click the **Save for Graphing** checkbox beside each of the Response Time metrics.

**6 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**7 Click Finish.**



## Performance Check

The Performance Check monitor gathers the following metrics:

- The percentage of CPU time – user, system, waiting for IO, or total – that is being used. This percentage is averaged over the number of seconds that you specify.
- The percentage of swap space that is available.
- CPU usage (reported by the `ps` utility), averaged over the number of minutes that you specify.
- The number of network collisions per second, inbound errors per second, and outbound errors per second.
- The number of network retransmits, averaged over the number of seconds that you specify.

## Configuring Performance Check Monitors

To configure Performance Check monitors, do the following:

### 1 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.



This monitor does not have separate Standard and Advanced views.

### 2 In the CPU Check area, do the following:

- Select one of the following options from the **CPU Value** dropdown list:
  - User

Time that the CPU spends processing application threads or threads that support tasks which are specific to applications.

- System

Time that the kernel spends processing system calls. If all 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.

- Waiting on I/O

Time that a runnable process requires to perform an I/O operation.

- Total

The total of all CPU time that is being used.

- Enter values, expressed as percentages, in the **CPU Warning Threshold** and **CPU Critical Threshold** fields.
- Enter the time period, in minutes, over which [up.time](#) should check CPU processes in the **CPU Time Interval** field.

**3 In the Swap Check area, enter values, expressed as percentages, in the Used Swap Warning Threshold and Used Swap Critical Threshold fields.**

When the percentage of available swap space exceeds these thresholds, [up.time](#) issues an alert.

**4 In the Process Check area, complete the following fields:**

- Process Name

The name of process that you want this monitor to check. This monitor uses the `ps` utility on UNIX to collect information about active processes. For example, to check the status of the email process enter `sendmail` in this field.

- Enter values, expressed as percentages, in the **Process Warning Threshold** and **Process Critical Threshold** fields.
- Enter the time period, in minutes, at which [up.time](#) will check the process in the **Process Check Time Interval** field.

5 In the **Network Check** area, do the following:

- Select one of the following options from the **Network Value** dropdown list:
  - Collisions

The simultaneous presence of signals from two nodes on the network, which can occur when two nodes start transmitting over a network at the same time. During a collision, both packets involved in a collision are broken into fragments and must be retransmitted.
  - In Errors

Data packets that were received but could not be decoded because either their headers or trailers were not available.
  - Out Errors

Data packets that could not be sent due to problems transmitting the packets or formatting the packets for transmission.
- Enter values, expressed as percentages, in the **Network Warning Threshold** and **Network Critical Threshold** fields.

6 In the **Network Retransmit Check** section, complete the following fields:

- Network Retransmits Warning Threshold

The number of retransmits per second that must be exceeded for [up.time](#) to issue a Warning alert.
- Network Retransmits Critical Threshold

The number of retransmits per second that must be exceeded for [up.time](#) to issue a Critical alert.
- Network Retransmits Time Interval

The time interval, in minutes, at which [up.time](#) checks retransmits.

7 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).

- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**8 Click Finish.**

## Process Count Check

The Process Count monitor measures the number of identical processes that are running on a system. If there is more than one instance of a process running, the check returns an OK status. If the process is not running, the check returns a Critical status.

### Configuring Process Count Check Monitors

To configure Process Count Check monitors, do the following:

- 1 **In the Process Count Check monitor template, select Standard or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- **Process Name (Mandatory)**

The exact name of the process that you want to monitor.

The name is the absolute name of the process, without its path or any parameters. For example, on UNIX systems the process `/usr/bin/vmstat -p` is checked as `vmstat`.

- **Process Occurrences**

Enter the number of process occurrences for which you want to set Warning and Critical thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- **Response Time (Advanced View)**

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 To save the data from the thresholds for graphing or reporting, click the **Save for Graphing** checkbox beside each of the metrics that you selected in step 3.
- 5 Complete the following settings:
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click **Finish**.

# CHAPTER 10

## Microsoft Windows Monitors

.....

The Microsoft Windows monitors track the performance and health of following:

|                                        |            |
|----------------------------------------|------------|
| <i>Windows Event Log Scanner.....</i>  | <i>162</i> |
| <i>Windows Service Check.....</i>      | <i>165</i> |
| <i>Windows File Shares (SMB) .....</i> | <i>168</i> |
| <i>Active Directory .....</i>          | <i>170</i> |

## Windows Event Log Scanner

The Windows Event Log Scanner determines if entries exist in a Windows log file. This monitor searches a specified number of lines in a log file that is returned by an agent for a text string. The monitor then compares those lines to ones in an existing temporary log file. Only the new lines are checked for events. The monitor assumes that existing lines in the temporary file have already been checked.

### Configuring Windows Event Log Scanner Monitors

To configure Windows Event Log Scanner monitors, do the following:

- 1 In the Windows Event Log Scanner monitor template, select Standard or Advanced View.**

- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:**

- Event Log Type (Mandatory)

Choose one of the following types of event log to search:

- Application

A log that records events generated by programs running on the server.

- System

A log that records the activity of various components of the operating system.

- Security

A log that records events such as login attempts and attempts to access files.

- Number of Lines

The number of lines in the log file that [up.time](#) will scan for the criteria specified in the monitor template. The default is 100.



- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the length of time a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

To save the data from the thresholds for graphing or reporting, click the **Save for Graphing** checkbox beside the Response Time metrics.

- Match event ID with

A number that identifies the type of event.

- Match event type with

The type of event to search for, which can be one of the following:

- Information

Describes the successful completion of a task.

- Warning

Indicates that a problem may occur in the future.

- Error

A problem, which may involve the loss of data or system integrity, has occurred.

- Success Audit

Found in the Security log, this describes the successful completion of an audited security event.

- Failure Audit

Found in the Security log, this describes the failure of an audited security event.

- Match computer name with

The name of the computer on which the event occurred.

- Match category with

The way in which the application, system component, or application module that triggered the event classifies the event. For example: System Event (in the Security Log); or Installation, CI Service, or wrapper (in the Application and System logs).

- Match source with  
The application, system component, or application module that triggered the event.
- Search description for  
Enter the string for which you want to search in the event log, for example:  
`The WMI Performance Adapter service entered the running state`  
The string is evaluated as a regular expression.

**4 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**5 Click Finish.**

## Windows Service Check

The Windows Service Check monitor alerts you to changes in the status of Windows services. A service is a process that extends the features of Windows by providing support to other programs. The default installation of Windows provides a core set of services and configurations that suits most needs.

There are approximately 100 services in the Windows Server family of operating systems. You can add services that you develop, or by installing third-party applications on a system.

Every Windows service has one of the following states, which control how the services are launched or prevented from launching:

- **Disabled**  
Services that are installed but not currently running.
- **Set to manual**  
Services that are installed but will start only when another service or application needs its functions.
- **Set to automatic**  
Services that are started by the operating system after device drivers are loaded at boot time.

## Configuring Windows Service Check Monitors

To configure Windows Service Check monitors, do the following:

- 1 In the Windows Service Check monitor template, select Standard or Advanced View.**
- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

### 3 Complete the following fields:

- Service Name (Mandatory)

You can find the name of all available Windows services, their states, and their status in a service property window by doing the following:

- On the Windows desktop, right click on **My Computer** and select **Manage**.
- Click **Services and Applications**, and then click **Services**.
- Double-click on the name of the service that you want to review.



If you enter the name of a service that does not exist, or mistype the name, the monitor changes the status of the service to Critical.

- Service Status (Mandatory)

Select a comparison method from the **Comparison Method** dropdown list, and then select one of the following:

- Stopped

The service is stopped.

- Start Pending

The service is stopped or paused while waiting for another process or condition to be satisfied before starting.

- Stop Pending

The service is running while waiting for another process or condition to be satisfied before stopping.

- Running

The service is running.

- Continue Pending

The service is waiting for another process or condition to be satisfied before continuing to run the service.

- Pause Pending

The service is running while waiting for another process or condition to be satisfied before pausing the service.

- Paused

The service is paused.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **To save the data from the thresholds for graphing or reporting, click the *Save for Graphing* checkbox beside each of the metrics that you selected in step 3.**

- 5 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 **Click *Finish*.**

## Windows File Shares (SMB)

The Windows File Shares (SMB) monitor can check the availability of file shares on a Windows server. If a file share is not available, the status of this monitor becomes critical and [up.time](#) sends an alert.



The Windows File Shares (SMB) monitor does not check the availability of hidden shares.

## Configuring Windows File Shares (SMB) Monitors

To configure Windows File Shares (SMB) monitors, do the following:

- 1 In the Windows File Shares (SMB) monitor template, select Standard View or Advanced View.**

- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:**

- Username

The user name that is required to login to the file share.

- Password

The password that is required to log in to the file share.

- Shares

The names of file shares that you want to monitor on a host system. Specify the name of the file share – for example `Main`.

To specify multiple file shares, add a comma between the names – for example, `Main, home`.

To check all of the file shares on a system, leave this field blank.

- Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 **Click the Save for Graphing checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**
- 5 **Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 **Click Finish.**

## Active Directory

Active Directory is a distributed network management service that is included in the Microsoft Windows Server 2003 and Windows 2000 Server operating systems. Active Directory provides a centralized location for all of the information about the services and resources within your network. Using this information, you can easily manage information about users, network devices, and any other resources that you might find useful to maintain.

The Active Directory monitor can check for any settings or information in your Active Directory. The monitor can start the check from any location within your Active Directory structure.

The Active Directory monitor attempts to match information that you have specified with information available in your Active Directory. If the monitor finds the information, the service monitor returns a status of OK. Otherwise, the monitor returns a Critical error and [up.time](#) generates an alert.

## Configuring Active Directory Monitors

To configure Active Directory monitors, do the following:

**1 In the Active Directory monitor template, select Standard View or Advanced View.**

**2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

**3 Complete the following fields:**

- Port

The number of the port number on which the Active Directory server is listening.

- Password

The password that is required to log in to the Active Directory server.



- Base

The location in the Active Directory from which you want the monitor to begin searching for information.

- Bind

The Bind string, which associates user account properties and Active Directory account attributes. This string gives you access to the Base location of your Active Directory structure.

The format of the Bind string must match the Base location of your Active Directory structure. Depending on your network security model, you will need domain controller administration privileges to bind to the locations on which you want to match information.

- Attribute

The attribute or information for which you want to search in your Active Directory.

An Active Directory entry consists of a set of attributes. Each attribute has a type – which describes the kind of information contained in the attribute – and one or more values, which contain the actual data. For example, the entry `jsmith@inter.net` has the Attribute value `jsmith@inter.net`. The Attribute type is e-mail.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Optionally, click the **Save for Graphing** checkbox beside the **Response Time** option to save the data for a metric to the **DataStore**, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**

# CHAPTER 11

## Application Monitors

.....

The application monitors track the performance and health of following:

|                                   |     |
|-----------------------------------|-----|
| <i>Uptime Agent</i> .....         | 174 |
| <i>Exchange</i> .....             | 176 |
| <i>IIS</i> .....                  | 179 |
| <i>WebLogic</i> .....             | 182 |
| <i>WebSphere</i> .....            | 187 |
| <i>ESX v3 Workload</i> .....      | 193 |
| <i>Splunk Query</i> .....         | 196 |
| <i>Live Splunk Listener</i> ..... | 198 |

## Uptime Agent

The Uptime Agent monitor determines whether or not an agent is running on a system that you are monitoring.

### Configuring Uptime Agent Monitors

To configure Uptime Agent monitors, do the following:

- 1 In the Uptime Agent monitor template, select Standard or Advanced View.**
- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.
- 3 Complete the following options by clicking the checkbox beside each option, then specifying a warning and critical threshold.**

If the thresholds that you set are exceeded, then [up.time](#) generates an alert. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

  - Major (Advanced View)

The major version number of the agent. For more information, see “Understanding Major and Minor Versions” on page 17.
  - Platform (Advanced View)

The operating system on which the agent is installed and running.
  - Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the length of time a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 To save the data from the thresholds for graphing or reporting, click the Save for Graphing checkbox beside each of the metrics that you selected in step 3.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**

## Exchange

The Exchange monitor identifies when certain performance counters for Microsoft Exchange servers have exceeded user-defined thresholds. These thresholds can be, for example, an inordinately high number of inbound connections or a rapidly-growing message queue. Whenever a threshold exceeds a warning or critical amount, [up.time](#) generates an alert.



The Exchange monitor only works with Microsoft Exchange versions 2000 and 2003.

## Configuring Exchange Monitors

To configure Exchange monitors, do the following:

**1 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

**2 Complete the following settings by clicking the checkbox beside each option, and then specifying a warning and critical threshold.**

If the thresholds that you set are exceeded, then [up.time](#) generates an alert. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Web Mail Sends Per Second

The maximum number of messages that can be sent from the Exchange server each second.

- Web Mail Auths Per Second

The maximum number of authorization requests that can be sent to the Exchange server each second.

- SMTP Bytes Sent Per Second

The total number of bytes sent per second by the Exchange SMTP server.

- **SMTP Bytes Received Per Second**  
The total number of bytes received per second by the Exchange SMTP server.
  - **SMTP Bytes Total Per Second**  
The total number of bytes of information passing through the Exchange SMTP server each second.
  - **SMTP Local Queue Length**  
The number of messages in the SMTP queue that are scheduled for local delivery.
  - **SMTP Messages Per Second**  
The maximum number of messages per second that are allowed by the SMTP server.
  - **SMTP Inbound Connections**  
The number of incoming connections that the SMTP server allows.
  - **SMTP Outbound Connections**  
The number of outbound connections that the server allows to all remote domains.
  - **SMTP Connection Errors Per Second**  
The number of number of connection errors that occur per second.
  - **Response Time (Advanced View)**  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 3 To save the data from the thresholds for graphing or reporting, click the Save for Graphing checkbox beside each of the metrics that you selected in step 2.**

### **4 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

### **5 Click Finish.**



## IIS

The IIS (Internet Information Server) service monitor checks the performance of an IIS Web server, based on thresholds that you set against common IIS performance counters. You can use this monitor to determine whether or not IIS is running on a defined port, and according to the thresholds you have set on common performance counters.

### Configuring IIS Monitors

To configure IIS monitors, do the following:

**1 In the IIS monitor template, select Standard or Advanced View.**

**2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

**3 Complete the following settings by clicking the checkbox beside each option, and then specifying a warning and critical threshold.**

If the thresholds that you set are exceeded, then [up.time](#) generates an alert. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Bytes Sent / Sec.

The number of bytes that are sent by the server each second.

- Bytes Received / Sec.

The number of bytes that are received by the server each second.

- Anonymous Users / Sec.

The rate, in seconds, at which users have made anonymous requests to the IIS server.

- Non-anonymous Users / Sec.

The rate, in seconds, at which registered users have made non anonymous requests to the IIS service.



IS 6.0 treats both an anonymous and a non-anonymous user request as a new user.

- Current Connections

The number of active connections to the IIS server.

- Connection Attempts / Sec.

The number of connection attempts that have been made, per second, since the IIS server was started.

- Logon Attempts / Sec.

The number of attempts, per second, that are being made to log on to the server.

- Get Requests / Sec.

The rate, in seconds, at which HTTP requests using the GET method have been made to the server.

- Post Requests / Sec.

The rate, in seconds, at which HTTP requests using the POST method have been made to the server.

- CGI Requests / Sec.

The rate, in seconds, at which the server is processing simultaneous CGI (Common Gateway Interface) requests.

- ISAPI Requests / Sec.

The rate, in seconds, at which the server is processing ISAPI extension requests.

ISAPI enables programmers to develop Web applications that are tightly integrated with IIS. ISAPI can also provide security functions to Windows servers and database connections through IIS.

- Not Found Errors / Sec.  
The maximum number of 404 file not found errors – indicating that the requested document cannot be found on the server – that can occur each second.
  - Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds for the length of time a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 **To save the data from the thresholds for graphing or reporting, click the **Save for Graphing** checkbox beside each of the metrics that you selected in step 3.**
- 5 **Complete the following settings:**
- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 **Click Finish.**

## WebLogic

The WebLogic monitor collects data that enables you to determine whether or not there is a performance problem or a failure on a WebLogic application server. Using the data that the WebLogic monitor collects, you can determine the root cause of the issue by generating a report (see “Reports for J2EE Applications” on page 361 for more information).

The WebLogic monitor collects the following metrics from a WebLogic server:

| Variables        | Metrics                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Connection Pools | <ul style="list-style-type: none"><li>• FailuresToReconnectCount<br/>The number of times that the connection pool failed to reconnect to a data store.</li><li>• ConnectionDelayTime<br/>The average time that was required to connect to a connection pool.</li><li>• ActiveConnectionsCurrentCount<br/>The current number of active connections in a JDBC connection pool.</li><li>• ActiveConnectionsHighCount<br/>The highest number of active connections in a JDBC connection pool.</li><li>• LeakedConnectionsCount<br/>The total number of connections that have been checked out of, but not returned to, the connection pool.</li><li>• CurrCapacity<br/>The current number of database connections in the JDBC connection pool.</li><li>• NumAvailable<br/>The number of available sessions in the session pool that are not currently being used.</li><li>• WaitingForConnectionCurrentCount<br/>The current number of requests that are waiting for a connection to the connection pool.</li></ul> |

| Variables | Metrics                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Per EJB   | <ul style="list-style-type: none"> <li>• AccessTotalCount<br/>The total number of times an attempt was made to get an EJB instance from the free pool.</li> <li>• BeansInCurrentUseCount<br/>The number of EJB instances in the free pool which are currently in use.</li> <li>• CachedBeansCurrentCount<br/>The total number of EJBs that are in the execution cache.</li> <li>• ActivationCount<br/>The number of EJBs that have been activated.</li> </ul>                                                                                                                                                                                            |
| Other     | <ul style="list-style-type: none"> <li>• HeapSizeCurrent<br/>The amount of memory, in bytes, that is in the WebLogic server's JVM heap.</li> <li>• HeapFreeCurrent<br/>The current amount of free memory, in bytes, that is in the WebLogic server's JVM heap.</li> <li>• OpenSocketsCurrentCount<br/>The current number sockets on the server that are open and receiving requests.</li> <li>• AcceptBacklog<br/>The number of requests that are waiting for a TCP connection.</li> <li>• ExecuteThreadCurrentIdleCount<br/>The number of threads in the server's execution queue that are idle or which are not being used to process data.</li> </ul> |

| Variables | Metrics                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | <ul style="list-style-type: none"><li>• PendingRequestCurrentCount<br/>The number of pending requests that are in the server's execution queue.</li><li>• TransactionCommittedTotalCount<br/>The total number of transactions that have been processed by the WebLogic server.</li><li>• TransactionRolledBackTotalCount<br/>The total number of transactions that have been rolled back.</li><li>• InvocationTotalCount<br/>The total number of times that a servlet running on the WebLogic server was invoked.</li></ul> |

Before you can use the WebLogic monitor, you must deploy the weblogic.jar file on the [up.time](#) monitoring station.

## Deploying the weblogic.jar File

In order for [up.time](#) to collect information from a WebLogic server, the file weblogic.jar must be deployed on the monitoring station.

To deploy the weblogic.jar file, do the following:

- 1 Locate the weblogic.jar file on the WebLogic server.**

The file is located in the lib folder in the directory in which WebLogic is installed. For example, on Windows the default folder is:

C:\bea\weblogicXX\server\lib

Where XX is the version of WebLogic. [up.time](#) 4 supports WebLogic versions 8.1 and 9.2.

- 2 Copy the file to the externaljar directory on the monitoring station. For example, on Windows copy the file to:**

C:\Program Files\uptime software\uptime4\externaljar\

## Configuring WebLogic Monitors

To configure WebLogic monitors, do the following:

- 1 **In the WebLogic monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- WebLogic Port

The number of the port number on which the WebLogic server is listening. The default is 7001.

- Username

The user name that is required to log into the WebLogic server.

- Password

The password that is required to log in to the WebLogic server.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Optionally, click the Save for Graphing checkbox beside the Response Time option to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**



## WebSphere

WebSphere is a software platform that provides firms with an environment for developing and deploying Web services and E-Commerce applications. Since WebSphere large and complex, it can be difficult to pinpoint the source of a problem, especially when that problem is intermittent.

The [up.time](#) WebSphere monitor collects data that you can use to generate a report, which will give you a historical view of problems that occur on a WebSphere server. See “WebSphere Report” on page 361 for more information.

The WebSphere monitor enables you to collect data so that you can:

- Determine whether or not the server can cope with its load.
- Determine the cause of problems with the server.
- Collect and retain data for later graphing and reporting.

The WebSphere monitor collects the following counters from a WebSphere Application Server:

| Variable         | Counters                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Connection pools | <ul style="list-style-type: none"> <li>•PoolSize<br/>The size of the connection pool to the data source.</li> <li>•FreePoolSize<br/>The number of free connections in the pool.</li> <li>•PercentUsed<br/>The percentage of the connection pool that is currently in use.</li> <li>•WaitTime<br/>The average time, in milliseconds, that a connection is used. The average time is the difference between the time at which the connection is allocated and the time at which it is returned.</li> </ul> |

| Variable | Counters                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | <ul style="list-style-type: none"><li>•CreateCount<br/>The total number of connections that were created.</li><li>•CloseCount<br/>The total number of connections that were closed.</li><li>•WaitingThreadCount<br/>The number of threads that are currently waiting for a connection.</li><li>•UseTime<br/>The average time, in milliseconds, that a connection is used. The average use time is the difference between the time at which the connection is allocated and that time at which it is returned.</li></ul>                                                                      |
| Per EJB  | <ul style="list-style-type: none"><li>•CreateCount<br/>The number of times that the Enterprise JavaBeans that are running on the server were created.</li><li>•RemoveCount<br/>The number of times that the EJBs were removed.</li><li>•PassivateCount<br/>The number of times that EJBs were removed from the cache. Note that passivation preserves the state of the EJBs on the disk</li><li>•MethodCallCount<br/>The total number of method calls that were made to the EJBs.</li><li>•MethodResponseTime<br/>The average response time, in milliseconds, on the bean methods.</li></ul> |

| Variable             | Counters                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Java Virtual Machine | <ul style="list-style-type: none"> <li>•cpuUsage<br/>The percent of CPU resources that were used since the last query.</li> <li>•HeapSize<br/>The total amount of memory that is available for the JVM.</li> <li>•UsedMemory<br/>The amount of memory that is being used by the JVM.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Other                | <ul style="list-style-type: none"> <li>•ActiveCount<br/>The number of global transactions which are concurrently active.</li> <li>•CommittedCount<br/>The total number of global transactions that have been committed.</li> <li>•RolledBackCount<br/>The total number of global transactions that have been rolled back.</li> <li>•LiveCount<br/>The number of servlet sessions that are currently cached in memory.</li> <li>•PoolSize<br/>The average number of threads in the servlet connection thread pool.</li> <li>•TimeSinceLastActivated<br/>The difference, in milliseconds, between the previous and current access time stamps of a servlet session. This counter does not include session time out values.</li> </ul> |

Before [up.time](#) can start collecting performance data from a WebSphere server, you must deploy the WebSphere performance servlet.

## Deploying the WebSphere Performance Servlet

The WebSphere performance servlet uses WebSphere's Performance Monitor Interface (PMI) infrastructure to retrieve performance information from a WebSphere Application Server. The information that the servlet collects is saved to an XML file.

By default, the PMI is enabled on the WebSphere server and is set to collect the performance metrics that [up.time](#) supports. Before [up.time](#) can begin collecting information from a WebSphere server, you must deploy the performance servlet in the WebSphere directory that contains your Web application.



The following steps must be completed for each Web application server that you want to monitor with [up.time](#).

To deploy the performance servlet do the following:

**1 On the WebSphere server, locate the following file:**

```
install_root/perfServletApp.ear
```

Where `install_root` is the directory under which WebSphere is installed.

**2 Copy the file `perfServletApp.ear` to the directory in which your Web application is installed. For example:**

```
install_root/installedApps/<cell_name>/
DefaultApplication.ear/DefaultApplication.war/WEB-INF/
classes
```

Where:

- `install_root` is the directory under which WebSphere is installed.
- `<cell_name>` is the name of the WebSphere node under which your Web application is installed.

## Deploying the Performance Servlet on WebSphere 6

If you are using WebSphere Application Server version 6, you will need to change two settings in the WebSphere management console to avoid an Access Denied error when [up.time](#) attempts to connect to the performance servlet to collect metrics.

To make the changes, do the following:

- 1 **In the WebSphere management console, modify the following settings:**
  - Under Security - Secure administration, applications, and infrastructure - turn Application Security on.
  - Under Enterprise Applications - perfServletApp - Security role to user/group mapping - turn Everyone off.
- 2 **Restart the server. up.time should now be able to connect to the servlet and gather performance metrics.**

## Configuring WebSphere Monitors

To configure a WebSphere monitor, do the following:

- 1 **On the WebSphere monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- WebSphere Port

The number of the port number on which WebSphere is listening. The default is 9080.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Optionally, click the Save for Graphing checkbox beside the Response Time option to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**

## ESX v3 Workload

The ESX v3 Workload monitor collects a set of metrics from all of the instances that are running on a ESX v3 server over a specified time period. The monitor then compares the highest values returned by the instances and then compares them to the thresholds that you set. If the values exceed the thresholds, [up.time](#) issues an alert. The monitor does not pinpoint the specific instance(s) that have exceeded the defined thresholds.

For example, you are monitoring an ESX v3 server that is running three instances. You configured the ESX v3 Workload monitor to collect data samples every 10 minutes, and to issue a warning when memory usage exceeds 300 MB. The three instances are using the following amounts of memory: 110 MB, 227 MB, and 315 MB. The ESX v3 Workload monitor focuses on the value of 315 MB and, since it exceeds the warning threshold, issues an alert.

## Configuring ESX v3 Workload Monitors

To configure an ESX v3 Workload monitor, do the following:

- 1 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 2 Complete the following fields:**

- Time Interval

The amount of time, in minutes, at which the monitor will collect data samples from the ESX server.

- CPU Warning Threshold

The amount of processor power, measured in megahertz (MHz), that the

- **Network Bandwidth Warning Threshold**  
The amount of network traffic in and out of the server, measured in megabits per second (Mbit/s), that must be exceeded before [up.time](#) issues a warning.
- **Network Bandwidth Critical Threshold**  
The amount of network traffic in and out of the server, measured in megabits per second (Mbit/s), that must be exceeded before [up.time](#) issues a critical alert.
- **Disk Usage Warning Threshold**  
The amount of data being written to the server's hard disk, measured in kilobytes per second (kB/s), that must be exceeded before [up.time](#) issues a warning.
- **Disk Usage Critical Threshold**  
The amount of data being written to the server's hard disk, measured in kilobytes per second (kB/s), that must be exceeded before [up.time](#) issues a critical alert.
- **Memory Usage Warning Threshold**  
The amount of overall system memory, measured in megabytes (MB), that must be exceeded before [up.time](#) issues a warning.
- **Memory Usage Critical Threshold**  
The amount of overall system memory, measured in megabytes (MB), that must be exceeded before [up.time](#) issues a critical alert.
- **Percent Ready Warning Threshold**  
The percentage of time that one or more instances running on an ESX v3 server is ready to run, but cannot run because it cannot access the processor on the ESX v3 server. If the value returned from the server exceeds this threshold, then [up.time](#) issues a warning.
- **Percent Ready Critical Threshold**  
The percentage of time that one or more instances running on an ESX v3 server is ready to run, but cannot run because it cannot access the processor on the ESX 3 server. If the value returned from the server exceeds this threshold, then [up.time](#) issues a critical alert.



- Percent Used Warning Threshold

The percentage of CPU time that an instance running on an ESX v3 server is using. If the value returned from the server exceeds this threshold, then **up.time** issues a warning.

- Percent Used Critical Threshold

The percentage of CPU time that an instance running on an ESX v3 server is using. If the value returned from the server exceeds this threshold, then **up.time** issues a critical alert.

For more information about setting thresholds, see “Configuring Warning and Critical Thresholds” on page 121.

### 3 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

### 4 Click Finish.

## Splunk Query

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 data center.

When you click the Splunk icon ( **splunk**➤) beside the names of services that are in WARN or CRIT states in the **My Portal** panel, you will be taken to your Splunk search page.

You can use the Splunk Query monitor to perform Splunk queries on log files to pinpoint an error condition.



Before you can use a Splunk Query monitor, you must add some settings specific to Splunk to the file `uptime.conf`. See “Splunk Settings” on page 468 for more information.

## Configuring Splunk Query Monitors

To configure a Splunk Query monitor, do the following:

### 1 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

### 2 Complete the following fields:

- Splunk query

The Splunk query string that you want to use to search log file for an error condition. For example, entering the following query string:

```
host::mailServer sendmail error hoursago::2
```

Will search log files, that were generated for the system named mailServer, for the word sendmail and error that were logged within the last two hours.

You can enter any Splunk query string in this field. For more information on the syntax of Splunk queries, see the Splunk user manual.

- Result count of splunk query

Enables **up.time** to alert you when the number of results that match your Splunk query exceeds the warning and critical thresholds that you set.

For example, you can configure the monitor to issue a Warning alert when five or more Splunk queries matching your query are returned, and a Critical alert when 10 or more results for your query are returned

| Result count of splunk query ▾ |                               |            | <input checked="" type="checkbox"/> |
|--------------------------------|-------------------------------|------------|-------------------------------------|
| Warning                        | is greater than or equal to ▾ | 5 results  |                                     |
| Critical                       | is greater than or equal to ▾ | 10 results |                                     |

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

**3 To save the data from the thresholds for graphing or reporting, click the Save for Graphing checkbox beside any of the options listed in step 2.**

**4 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**5 Click Finish.**

## Live Splunk Listener

Live Splunks are scheduled searches of Splunk queries that are saved on the Splunk server. A Live Splunk automatically runs a search, can initiate an alert, and can perform actions based on that alert. You can, for example, set up a Live Splunk to search for all critical error conditions.

The Live Splunk Listener monitor enables you to capture the information generated by a Live Splunk. This monitor is very similar to the External Check monitor (see page 286), and uses the script `liveSplunkHandler.py` (which is bundled with [up.time](#)) to return Live Splunk information to the monitoring station.

The `liveSplunkHandler.py` script takes the following options:

- `--message`  
A message that will be returned to the [up.time](#) monitoring station. For example, if the Live Splunk is configured to search for warning conditions, you can enter the message "Changed to WARN".

- `--status`

The script can return the following status codes:

- 0 – OK  
The services are functioning properly.
- 1 – Warning  
There is a potential problem with one of more of the services being monitored.
- 2 – Critical  
There is a critical problem with one or more of the services being monitored.
- 3 – Unknown  
There is an error in the configuration of the monitor itself, or [up.time](#) cannot execute the service check.

- `--monitorName`

The name of the [up.time](#) monitor to which the information from the Live Splunk will be directed.

The following is an example of the script with all of its options specified:

```
liveSplunkCommand.py --message="sendmail has some traffic
going through new command!" --status=2 --monitorName="Live
Splunk"
```

[up.time](#) captures the output from the script, which appears in the service status section of the **Global Scan** panel (see “Understanding the Status of Services” on page 25). The [up.time](#) monitoring framework picks up any error codes and triggers the appropriate monitoring action.

## Before You Begin

Before you can monitor Live Splunks, you must:

- Edit the script `liveSplunkHandler.py` to point to the [up.time](#) monitoring station. Do the following to edit the script:
  - Navigate to the scripts folder on the monitoring station.**
  - Open the file `liveSplunkHandler.py` in a text editor.**
  - Find the following entry in the file:**

```
Specify the up.time server and port
by setting the following two variables
host = "localhost"
port = "9996"
```
  - Change the values for host and port to the host name and port of the monitoring station.**
  - Save the file and exit the text editor.**
- Copy the script `liveSplunkHandler.py` from the scripts folder on the monitoring station to the folder `/data/splunk/bin/scripts` on the Splunk server.
- Configure a Live Splunk. For information on configuring Live Splunks, see the Splunk user manual.

When setting up your Live Splunk, select the **Run the shell script option** on the configuration page. Then, enter the path to `liveSplunkHandler.py`, along with the script options, in the field:

☒ Run the shell script  
`/data/splunk/bin/scripts/liveSplunkHandler.pl --host="dev-latest" --port=9996 --message="failed login (windows)" --status=1 -`

## Configuring Live Splunk Listener Monitors

To configure a Live Splunk Listener monitor, do the following:

### 1 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

### 2 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

### 3 Click Finish.

# CHAPTER 12

## Database Monitors

.....

The database monitors track the performance and health of following:

|                                            |     |
|--------------------------------------------|-----|
| <i>MySQL (Advanced Metrics)</i> .....      | 202 |
| <i>MySQL (Basic Checks)</i> .....          | 209 |
| <i>Oracle (Advanced Metrics)</i> .....     | 211 |
| <i>Oracle (Basic Checks)</i> .....         | 214 |
| <i>Oracle Tablespace Check</i> .....       | 217 |
| <i>SQL Server (Basic Checks)</i> .....     | 220 |
| <i>SQL Server (Advanced Metrics)</i> ..... | 224 |
| <i>SQL Server Tablespace Check</i> .....   | 228 |
| <i>Sybase</i> .....                        | 233 |

## MySQL (Advanced Metrics)

The MySQL (Advanced Metrics) monitor checks the performance of MySQL databases and instances that are running on a system against the thresholds that you define. If MySQL is not responding, the database can process queries but the results will demonstrate behaviour that alerts you to a problem.

The MySQL (Advanced Metrics) monitor can:

- Determine whether or not a MySQL instance is running on your system.
- Check whether or not MySQL is listening on a specific port.
- Check performance values to determine the efficiency of a MySQL instance.

## Configuring MySQL (Advanced Metrics) Monitors

To configure MySQL (Advanced Metrics) monitors, do the following:

**1 Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

**2 Complete the following settings by entering the appropriate Warning and Critical thresholds.**

If the thresholds that you set are exceeded, then [up.time](#) generates an alert. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- **MySQL Port**

The number of the port on which the MySQL instance is listening. The default is 3306.

- **Username**

The user name that is required to log into the MySQL instance.

- **Password**

The password that is required to log into the MySQL instance.



- Uptime

The number of seconds that MySQL has been running.

- Questions

The number of queries that have been sent to the database.

- Slow Queries

The number of queries that take longer than `long_query_time` to complete.

When started with the `--log-slow-queries[=file_name]` option, MySQL writes a log file containing all SQL statements that took more than the `long_query_time` to execute. The time taken to acquire the initial table locks is not counted as execution time.

If the `file_name` value is not specified, the information is written to a file with the name of the host machine along with the suffix `-slow.log`. If a filename is given, but not as an absolute path name, the file is written to the default MySQL data directory.

You can use the `--log-queries-not-using-indexes` option to log queries that do not use indexes to the slow query log.

Queries handled by the query cache are not added to the slow query log, nor are queries that would not benefit from the presence of an index because a database table has no rows or just one row.

- Open Tables

The number of database tables that are opened independently by each concurrent thread.

Multiple clients can simultaneously issue queries for a given table. Each table is opened independently by each concurrent thread to ensure that multiple client threads do not have different states on the same table.

For each concurrent thread, the table must be opened twice if two threads access the same table or if a thread accesses the table twice in the same query. Each concurrent open requires an entry in the table cache. The first time any table is opened, it takes file descriptors for the data file and the index file. Each additional use of the table takes only a descriptor for the data file. The index file descriptor is shared among all threads.

The cache of open tables should be at the level specified by `table_cache` entries. The default value is 64. MySQL may temporarily open more tables to execute queries.

Unused tables are closed and removed from the table cache when:

- The cache is full and a thread tries to open a table that is not in the cache.
- The cache contains more than `table_cache` entries and a thread is no longer using a table.
- A table flushing operation occurs. This happens when someone issues a `FLUSH TABLES` statement, or executes either the `mysqladmin flush-tables` or `mysqladmin refresh` commands.

When the table cache fills up:

- The server locates a cache entry to release tables that are not currently in use, in least-recently used order.
- If a new table needs to be opened, but the cache is full and no tables can be released, the cache is temporarily extended as necessary.

When the cache is in a temporarily extended state and a table goes from a used to an unused state, the table is closed and released from the cache.

- **QPSA**

The average number of queries, per second, that must be exceeded before [up.time](#) generates an alert.

- **Bytes Received**

The number of bytes received by the server.

- **Bytes Sent**

The number of bytes sent by the server to all clients.

- **Delayed Insert Threads**

Select a comparison method for the Warning and Critical Thresholds. Then, enter the number of delayed insert threads that must be exceeded before [up.time](#) sends an alert.

The DELAYED option for the INSERT statement is a MySQL extension to standard SQL that you can use with clients that cannot wait for the INSERT statement to complete.

When a client uses the INSERT DELAYED statement, the row is immediately queued to be inserted when the table is not in use by any other thread. INSERT DELAYED also bundles inserts from multiple clients and writes them in one block.

The DELAYED option has the following constraints:

- It only works with MEMORY tables.
- You should only use INSERT DELAYED for INSERT statements that specify value lists. The server ignores DELAYED for INSERT DELAYED ... SELECT statements.
- The server ignores DELAYED for INSERT DELAYED ... ON DUPLICATE UPDATE statements.
- You cannot use LAST\_INSERT\_ID() to get the AUTO\_INCREMENT value the statement might generate because the statement returns immediately before the rows are inserted.
- DELAYED rows are not visible to SELECT statements until they actually have been inserted.

- Delayed Errors

The number of delayed insert threads that had an error.

- Max Used Connections

The maximum number of connections that have been in simultaneous use since the server was started.

- Open Files

The number of open files that must be exceeded before up.time generates an alert.

- Open Streams

The number of open data streams that must be exceeded before up.time generates an alert.

- **Table Locks Immediate**  
The number of times that a table lock is acquired immediately. For more information on table locks, see “SQL Server Locks” on page 475.
- **Table Locks Waited**  
The number of table locks waited that must be exceeded before [up.time](#) generates an alert. For more information on table locks, see “SQL Server Locks” on page 475.
- **Threads Cached**  
The number of threads in the thread cache that must be exceeded before [up.time](#) generates an alert.
- **Threads Connected**  
The maximum number of clients that can be connected to the database at any one time.
- **Threads Running**  
The number of threads that are running, which can be used to determine whether or not the database is becoming overloaded.  
  
If the database is overloaded, the monitor will report an increased number of running queries. However, you can have values that exceed this limit for very short times.
- **QCache Queries in Cache**  
The number of queries in the query cache (QCache) that must be exceeded before [up.time](#) generates an alert.
- **QCache Inserts**  
The number of queries added to the query cache.  
  
You should compare the value of the `qcache_hits` to the total number of select queries to determine the current hit rate. You can increase or decrease `query_cache_size` to find the value which provides optimal performance.
- **QCache Hits**  
The number of hits to the query cache (`qcache_hits`) to determine the number of query results taken directly from the cache instead of

executing them. When this number is exceeded, up.time generates an alert.

This metric shows the number of query results taken directly from the query cache instead of executing them. You should compare the value of QCache Hits to the total number of your SELECT queries to determine the current hit rate. Then, you can increase or decrease the query\_cache\_size to find the value which provides optimal performance.

- QCache Lowmem Prunes

The number of QCache\_lowmem\_prunes that can be deleted from the cache because of low memory.

This variable counts the number of queries that have been removed from the cache to free up memory for caching new queries. The query cache removes the least-recently used queries from the cache.

- QCache Not Cached

The maximum number of queries that are not cached.

- QCache Free Memory

The amount of free memory for the query cache.

- QCache Free Blocks

The number of free memory blocks in query cache.

- QCache Table Blocks

The amount of query cache memory fragmentation.

- Response Time

Enter the Warning and Critical Response Time thresholds for the overall time required to perform a service check. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 3 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.
- 4 Complete the following settings:
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 5 Click **Finish**.

## MySQL (Basic Checks)

The MySQL (Basic Checks) monitor:

- Determines whether or not a host that is running a MySQL database is available.
- Determines whether or not you can log into a MySQL database.
- Evaluates a response based on a script that is executed against a database or database instance.

## Configuring MySQL (Basic Checks) Monitors

To configure MySQL (Basic Checks) monitors, do the following

- 1 **In the MySQL (Basic Checks) monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**



If you enter a value in the **SID** field, [up.time](#) can capture the port value from the SID of the Oracle instance.

- Port Check (Optional)

Select this option to open a socket connection that determines whether or not the database is listening on the defined port.

- Username

The user name that is required to login to the MySQL database.

- Password

The password that is required to login to the MySQL database.

- Database

The name of the MySQL database instance.

- Script File

Enter the full path on the monitoring station to the script that this monitor will run against the database.

- Script

Select this option and then type or copy the script that you want [up.time](#) to against the database into this text box. Use this option if you do not have access to the file system on the monitoring station or if your script is short or will not regularly change.

- Match

Enter a string that you want to match against the return value from the script.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Click the [Save for Graphing](#) checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

- 5 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 **Click [Finish](#).**



## Oracle (Advanced Metrics)

The Oracle (Advanced Metrics) monitor captures a number of performance tuning metrics for your Oracle database. Some Oracle metrics are for tuning devices for long-term performance gains, rather than avoiding outages. This applies to following probes: Buffer Cache, Data Dictionary Cache, Disk Sort Ratio, Library Cache, and Redo Log. You should schedule the monitor to gather data less frequently – perhaps every hour or every two days.

### Configuring Oracle (Advanced Metrics) Monitors

To configure Oracle (Advanced Metrics) monitors, do the following:

- 1 In the Oracle (Advanced Metrics) monitor template, select **Standard View** or **Advanced View**.

- 2 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:

- Username

The user name that is required to login to the database.

- Password

The password that is required to login to the database.

- SID

The Oracle System Identifier (SID) that identifies this Oracle instance. The SID defaults to the database name. The SID is included in the `CONNECT DATA` paths of the connect descriptors in the `tnsnames.ora` file, and in the definition of the TNS listener in the `listener.ora` file.



If you do not complete the **Username** and **Password** fields up.time will attempt to connect to the database. If connection fails, the database returns a SQL exception error.

- Buffer Cache Hits Ratio

Enter the Warning and Critical thresholds for buffer cache hits that are completed without accessing disk I/O. To gather as much application data as possible, you should enter a high buffer cache hits ratio.

An Oracle database maintains its own buffer cache inside the system global area for each instance. A properly-sized buffer cache can yield a cache hit ratio over 90%. If a buffer cache is too small, the cache hit ratio will be small and the database uses more physical disk I/O. If a buffer cache is too large, then parts of the buffer cache will waste memory resources.

- Data Dictionary Cache Hits Ratio

Enter the Warning and Critical thresholds for data dictionary cache hits that are completed without accessing disk I/O.

The data dictionary cache tables provide information about all of the objects stored in your dictionary – for example tablespaces, files, users, rollback segments, constraints, synonyms. A hit ratio approaching 100% is ideal.

- Library Cache Hits Ratio

Enter the Warning and Critical thresholds for the rate at which library cache pin misses occur.

A pin miss occurs when an session executes a statement that has already been parsed, but which is no longer in the shared pool.

- Redo Log Space Request Ratio

Enter the Warning and Critical thresholds for the number of redo log space requests per minute that have been made since the server was started.

- Disk Sort Rate

Enter the Warning and Critical thresholds for the rate of Oracle sorts that are too large to be completed in memory and which are sorted using a temporary segment.

- **Active Sessions**  
Enter the Warning and Critical thresholds for the number of active sessions based on the value of `V$PARAMETER.PROCESSES` in the file `init.ora`.
  - **Oracle Blocking Sessions**  
Enter the Warning and Critical thresholds for the number of sessions that are preventing other sessions from committing changes to the Oracle database.
  - **Oracle Idle Sessions**  
Enter the Warning and Critical thresholds for the number of Oracle sessions that are idle, as determined by the Time Idle value that you specify. Only the sessions that have been idle for the duration (measured by the Time Idle value), in seconds, are considered idle.
  - **Response Time (Advanced View)**  
Enter the Warning and Critical Response Time thresholds for the length of time a service check needs to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 **Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**
  - 5 **Complete the following settings:**
    - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
    - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
    - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
    - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
    - Action Profile settings (see “Action Profiles” on page 134 for more information).
  - 6 **Click **Finish**.**

## Oracle (Basic Checks)

The Oracle (Basic Checks) monitor:

- Determines whether or not a host running an Oracle database is available.
- Determines whether or not an Oracle service is running on a system.
- Determines whether or not you can log into an Oracle database.
- Evaluates a response based on a script that you have executed against a database or database instance.



Use the Oracle Tablespace Check monitor (see “Oracle Tablespace Check” on page 217) to check Oracle tablespaces.

## Configuring Oracle (Basic Checks) Monitors

To configure Oracle (Basic Checks) monitors, do the following

**1 In the Oracle (Basic Checks) monitor template, select Standard View or Advanced View.**

**2 Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

**3 Complete the following fields:**

- Port

The number of the port on which the Oracle service is listening.



If you enter a value in the **SID** field, [up.time](#) can capture the port value from the SID of the Oracle instance.

- Port Check (Optional)

Select this option to open a socket connection that determines whether or not the database is listening on the defined port.

- Username  
The user name that is required to login to the Oracle database.
- Password  
The password that is required to login to the Oracle database.
- SID  
The Oracle System Identifier (SID) that identifies the Oracle instance.  
The SID defaults to the database name.  
  
If you enter a value in this field, up.time can capture the number of the port on which Oracle is listening.
- Script File  
Click the **Script File** check box and then enter the full path on the monitoring station to the script that this monitor will run against the database.



If you configured your database to allow logins with a user name and password and you specify the script file but no login information, the script will fail. The script will run properly if you have configured your database to allow logins without a user name and password.

- Script  
Select this option and then type or copy the script that you want up.time to against the database into this text box. Use this option if you do not have access to the file system on the monitoring station or if your script is short or will not regularly change.
- Match  
Enter a string that you want to match against the return value from the script.
- Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.
- 5 Complete the following settings:
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click **Finish**.

## Oracle Tablespace Check

The Oracle Tablespace Check monitors the size (as a percentage) of individual tablespaces within Oracle database instances. The Oracle Tablespace Check alerts you when a tablespace in your instance exceeds the defined thresholds.

Each database is logically divided into one or more tablespaces. One or more data files are explicitly created for each tablespace to physically store the data in a tablespace. The combined size of the data files in a tablespace is the total storage capacity of the tablespace. For example:

| ID# | TABLESPACE_NAME | Total Bytes | Bytes Free  | % Free |
|-----|-----------------|-------------|-------------|--------|
| 6   | INDX            | 56,623,104  | 56,614,912  | 99.99  |
| 8   | OEM_REPOSITORY  | 31,465,472  | 3,473,408   | 11.04  |
| 3   | RBS             | 104,857,600 | 75,489,280  | 71.99  |
| 9   | SUPPORT         | 52,428,800  | 52,420,608  | 99.98  |
| 1   | SYSTEM          | 56,623,104  | 2,850,816   | 5.03   |
| 4   | TEMP            | 71,303,168  | 71,294,976  | 99.99  |
| 2   | TOOLS           | 8,388,608   | 8,380,416   | 99.90  |
| 7   | UNANET          | 314,572,800 | 300,539,904 | 95.54  |
| 5   | USERS           | 20,971,520  | 20,963,328  | 99.96  |

In the above table, the SYSTEM tablespace is over 95% full. If you set the Warning threshold to 90%, and the Critical threshold to 95%, the Oracle Tablespace Check returns a status of Critical.



Use the Oracle (Basic Checks) monitor to determine the availability of Oracle databases, the performance of services, and the matched response of scripts. For more information, see “Sybase” on page 233.

## Configuring Oracle Tablespace Check Monitors

To configure Oracle Tablespace Check monitors, do the following:

- 1 In the Oracle Tablespace Check monitor template, select Standard View or Advanced View.**
- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

### 3 Complete the following fields:

- Port  
The number of the port on which the Oracle service is listening. The default is 1521.
- Username  
The user name that is required to login to the Oracle database.
- Password  
The password that is required to login to the Oracle database.
- SID  
The Oracle System Identifier (SID) that identifies the Oracle instance. The SID defaults to the database name.  
  
The SID is a unique name for an Oracle instance to switch between Oracle databases. The SID is included in the `CONNECT DATA` paths of the connect descriptors in the `tnsnames.ora` file. As well, the SID is in the definition of the TNS listener in the `listener.ora` file.



If you do not complete the **Username**, **Password**, **SID** fields **up.time** will attempt to connect to the database. If connection fails, the database returns a SQL exception error.

- Full Warning Threshold (Mandatory)  
Enter a value that will change the status of the Oracle Tablespace Check from OK to Warning.  
  
The warning threshold should be a percentage of the maximum file size, against which the monitor will check data files and log files.
- Full Critical Threshold (Mandatory)  
Enter a value that will change the status of the Oracle Tablespace Check from OK to Warning.  
  
The critical threshold should be a percentage of the maximum file size, against which the monitor will check data files and log files.



- Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds for the length of time that a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 **Click the Save for Graphing checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**
- 5 **Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 **Click Finish.**

## SQL Server (Basic Checks)

The SQL Server (Basic Checks) monitor compares the performance of SQL Server databases and instances running on a system to the thresholds that you define. The SQL Server (Basic Checks) monitor:

- Determines whether or not SQL Server is running on your system.
- Checks whether or not SQL Server is listening on a specific port.
- Determines whether or not SQL Server can process queries.
- Checks for values in base and computed tables.

You can use regular expressions to identify a wide range of responses and to detect problems after they occur. You can also run scripts through [up.time](#) to alert you when a database component that is being monitored is not performing as required.



To properly configure this monitor, you should have a strong knowledge of regular expressions, Transact-SQL, and SQL Server.

## Configuring SQL Server (Basic Checks) Monitors

To configure SQL Server monitors, do the following:

- 1 In the SQL Server (Basic Checks) monitor template, select Standard View or Advanced View.**

- 2 Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:**

- SQL Server Port

The number of the port on which SQL Server is listening.

SQL Server uses Static Port Allocation or Dynamic Port Allocation ports. For more information, see “SQL Server Ports” on page 470.

- Port Check (Optional)

Select this option to open a socket connection that determines whether or not the database is listening on the defined port.

You should perform a port check because SQL Server can communicate statically on a defined or default port, or communicate dynamically on a port assigned by the operating system.

- Username

The user name that is required to log into the SQL Server database.

- Password

The password that is required to log into the SQL Server database.

- Instance

The name of the SQL server instance to which you want to connect.

You can install multiple versions of Microsoft SQL Server on one computer. When installing a new version of SQL Server 2000, or maintaining an existing installation, you can specify it as:

- A default instance of SQL Server

This instance is identified by the network name of the computer on which it is running. SQL Server version 6.5 or SQL Server version 7.0 servers can operate as default instances. However, a computer can have only one version functioning as the default instance at one time.

- A named instance of SQL Server

This instance is identified by the network name of the computer plus an instance name, in the format `<computername>\<instancename>`.

Most applications must use SQL Server 2000 client components to connect to a named instance. However, you can use the SQL Server version 7.0 Client Network Utility to configure a server alias name that the SQL Server version 7.0 client components can use to connect to a named instance.

A computer can concurrently run any number of named instances of SQL Server. An instance name cannot exceed 16 characters.

- Database

The name of the SQL Server database that you want to monitor.

**up.time** views each database along the path /<system>/<instance>/<database>.

Each instance of SQL Server has four system databases – master, model, tempdb, and msdb – and one or more user databases. Depending on their permissions, users can access some or all of the databases in an instance.

A connection to an instance is associated with a particular database on the server, called the *current database*. You can switch from one database to another using the Transact-SQL `USE database_name` statement.

**up.time** gathers information from all of the databases in all instances on a system and aggregates this information in the metrics it returns to you. Unless you must identify a particular database on your system – for example, you have applied a name to the default instance – you should leave the **Database** field blank.

- Script File

Click the **Script File** check box and then enter the full path on the monitoring station to the script that this monitor will run against the database.



If you configured your database to allow logins with a user name and password and you specify the script file but no login information, the script will fail and an error message appears in the **Global Scan** panel. The script will run if you have configured your database to allow logins without a user name and password.

- Script

Click the **Script** checkbox and then type or copy the script that you want **up.time** to against the database into this text box. Use this option if you do not have access to the file system on the monitoring station or if your script is short or will not regularly change.

- Match

The value to match the script results against, which can be either a string or a regular expression. For more information, see “Comparison Methods” on page 120. For example, you can enter the following in the **Match** text box:

`^[OK]+`

Where:

- `^` means start the match at the beginning of the line.
- `[OK]` is the pattern to match.
- `+` is the pattern to match anywhere on the line.

The value that your script returns can be a string that you can match to. If you match to the value you checked for, the status of the service monitor is OK. Otherwise, the status of the service monitor is Critical.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.

- 5 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 Click **Finish**.

## SQL Server (Advanced Metrics)

SQL Server (Advanced Metrics) monitor collects information on the availability and performance of individual SQL Server databases.

You only need to configure one SQL Server (Advanced Metrics) monitor for each system. You can, however, create multiple SQL Server (Advanced Metrics) monitors for a system if you need to separately capture different SQL Server performance metrics. See the section Using Multiple SQL Server (Advanced Metrics) Monitors for more information.

For example, a host in your environment has:

- An [up.time](#) agent installed.
- Two database instances.
- Four databases.

The SQL Server (Advanced Metrics) monitor can capture performance information from all four databases. It can also aggregate the information to present a single performance value for each metric.

### Using Multiple SQL Server (Advanced Metrics) Monitors

You can create several SQL Server (Advanced Metrics) monitors for a system if you must separately capture different SQL Server performance metrics. For example, the SQL Server (Advanced Metrics) monitor provides metrics for SQL Server locks including lock requests, waits, and averages. For information about locks, see “SQL Server Locks” on page 475.

Lock requests do not always provide meaningful information. When you compare the length of waits with the number of lock requests, the length of the lock waits should be much lower than requests. If the lengths of waits and requests are about the same, then there is a performance problem. When the average lock wait time is high, there is a problem with SQL Server.

## Configuring SQL Server (Advanced Metrics) Monitors

To configure SQL Server (Advanced Metrics) monitors, do the following:

- 1 In the SQL Server (Advanced Metrics) monitor template, select Standard or Advanced View.**
- 2 Complete the monitor information fields.**  
To learn how to configure monitor information fields, see “Monitor Identification” on page 117.
- 3 In the Instance field, type the name of the SQL server instance to which you want to connect.**
- 4 Complete the following options by clicking the checkbox beside each option, then specifying a warning and critical threshold.**

If the thresholds that you set are exceeded, then [up.time](#) generates an alert. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Lock Wait / Sec.

The amount of time, in seconds, to wait for a database lock. For more information about locks, see “SQL Server Locks” on page 475.

- Lock Requests / Sec.

The number of new database locks and lock conversions that are requested from the lock manager every second. For more information about locks, see “SQL Server Locks” on page 475.

- Average Lock Wait Time

The average time, in milliseconds, that you must wait for database locks to clear before [up.time](#) sends an alert.

- User Connections

The number of user connections that are allowed before [up.time](#) sends an alert.

For example, a single host is running two databases. There are five users logged on to the first database and three users logged on to the second database. The total number of user connections is eight.

- Transactions / Sec.

In the Warning and Critical threshold fields, enter the number of transactions started for the databases across the host per second.

- Data File(s) Size / KB

The cumulative size of all the files in all of the databases on the host system.

This metric is returned from the SQL Server Database object. The Database object provides such information about the database as the amount of free log space available or the number of active transactions in the database. There can be multiple instances of this object.

- Total Latch Wait Time (ms)

The total time, in milliseconds, that it takes to complete the latch requests that were waiting over the last second.

- Latch Waits / Sec.

The number of latch requests that were not immediately granted, and which waited before being granted.

- Average Latch Wait Time (ms)

The average time, in milliseconds, that latch requests had to wait before being granted.

- Maximum Workspace Memory (KB)

The maximum amount of memory, in kilobytes, that the server has available to execute such processes as sort, bulk copy, hash, and index creation.

This metric is returned by the SQL Server Memory Manager object, which monitors overall server memory usage. By monitoring overall server memory usage, you can determine whether or not:

- Bottlenecks exist due to a lack of available physical memory for storing frequently accessed data in cache. If so, SQL Server must retrieve the data from the disk.
- You can improve query performance by adding more memory or by making more memory available to the data cache or to SQL Server internal structures.



- **Connection Memory (KB)**  
The total amount of dynamic memory, in kilobytes, that the server is using to maintain connections.
  - **SQL Cache Memory (KB)**  
The amount of memory, in kilobytes, that the server is using for the dynamic SQL cache.
  - **Total Server Memory (KB)**  
The total amount of committed memory from the buffer pool, in kilobytes, that the server is using.
  - **Response Time (Advanced View)**  
Enter the Warning and Critical Response Time thresholds. If the amount of time taken to perform a check exceeds the defined thresholds, it could indicate a problem that requires investigation.
- 5 To save the data from the thresholds for graphing or reporting, click the **Save for Graphing** checkbox beside each of the metrics that you selected in step 3.**
- 6 Complete the following settings:**
- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 7 Click Finish.**

## SQL Server Tablespace Check

The SQL Server Tablespace Check monitor evaluates the size of data files within SQL Server databases. [up.time](#) gathers information from all the databases across all instances on a system and aggregates this information in the metrics that it returns.

This monitor also reports whether or not any of the data files in a filegroup or any log file in any database in the instance exceeds warning and critical thresholds. If warning or critical thresholds are exceeded, [up.time](#) generates an alert.

### Structure of a SQL Server Database

Each SQL Server database consists of at least two files:

- A primary data file, with the extension `.mdf`.
- A log file, with the extension `.ldf`.

There are also secondary data files, with the extension `.ndf`. A database can have only one primary data file, zero or more secondary data files, and one or more log files. Each database file can only be used by one database.

In a database, data files store persistent data. For ease of management, you can group one or more data files into logical tablespaces. The SQL Server equivalent of an Oracle tablespace is the *filegroup*. SQL Server filegroups come under and are associated with the individual databases. The SQL Server data hierarchy is:

Instance / Database / FileGroup / Data file

Each data file can be a member of only one filegroup, but the log files are managed separately from one another. There are three types of filegroups:

- Primary
- User-defined
- Default


When you configure your SQL Server databases, you can the maximum size of data files to prevent disk drives from running out of space. If you do not

specify the size of data files, the database assumes that the size is unlimited.



**up.time** measures the size of data files and log files as a percentage of their maximum size. If a data file has an infinite maximum size, the percent of maximum datafile size must be near zero. You should always specify the maximum size of each data file.

The following diagram illustrates six data files in three filegroups in three databases across two instances of a system.



If you set SQL Server Instance\_B with a Critical threshold of 90% and a Warning threshold of 70%, the SQL Server Tablespace Check monitor watches the size of all data files in that instance. The monitor sends an alert if any of the files reaches or exceeds the defined thresholds.

## Configuring SQL Server Tablespace Check Monitors

To configure SQL Server Tablespace Check monitors, do the following:

- 1 In the SQL Server Tablespace Check monitor template, select **Standard View** or **Advanced View**.

### 2 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

### 3 Complete the following fields:

- SQL Server Port

The number of the port on which the SQL Server is listening.

SQL Server can use static or dynamic ports. For information about SQL ports and how to determine and configure port allocation, see “Configuring a Static Port” on page 472 and “Configuring a Dynamic Port” on page 473.

- Username

The user name that is required to login to the SQL Server database.

When a user connects through a Windows user account, SQL Server re-validates the account name and password by contacting a Windows domain controller to determine the network user name. SQL Server then verifies the credentials of the users, and then permits or denies login access.

- Password

The password that is required to login to the SQL Server database.

When a user connects with a specified login name and password from a non-trusted connection, SQL Server determines if a SQL Server login account has been set up and if the specified password matches the one previously recorded. If SQL Server does not find a login account, authentication fails and the user receives an error message.

SQL Server authentication is provided for backward compatibility because applications written for SQL Server version 7.0 or earlier may require the use of SQL Server logins and passwords.



If you do not complete the **Username** and **Password** fields **up.time** will attempt to connect to the database. If the connection attempt fails, the database returns a SQL exception error.

SQL Server can use one of the following authentication modes:

- Windows Authentication Mode

Enables users to connect to a SQL Server instance through a Windows user account.

- Mixed Mode

Enables users who to connect to a SQL Server instance through a Windows account to use either Windows authentication or SQL Server authentication.

- Instance

The SQL Server instance name. This is usually the default instance.

You can install multiple instances of SQL Server on one computer. An instance can be:

- The default instance

This instance is identified by the network name of the computer on which it is running. Applications using client software from earlier versions of SQL Server can connect to a default instance. SQL Server version 6.5 or 7.0 servers can operate as default instances. A computer can have only one version functioning as the default instance at a time.

- A named instance of SQL Server

This instance is identified by the network name of the computer plus an instance name, in the format `<computername>\<instancename>`.

Most applications must use SQL Server client components to connect to a named instance. However, you can use the SQL Server version 7.0 Client Network Utility to configure a server alias name that the version 7.0 client components can use to connect to a named instance of SQL Server. A computer can concurrently run any number of named instances of SQL Server. A named instance can run at the same time as an existing installation of SQL Server version 6.5 or SQL Server version 7.0. The instance name cannot exceed 16 characters.

A new instance name must begin with a letter, an ampersand (&), or an underscore (\_), and can contain numbers, letters, or other characters. Do

not use SQL Server sysnames and reserved names as instance names. For example, default is a reserved name and should not be used as an instance name.

You can have multiple instances of SQL Server installed on one computer. Each instance operates independently from the other instances, and applications can connect to any of the instances.

- **Full Warning Threshold**

Enter a percentage of the maximum file size you want to set as your warning threshold.

- **Critical Warning Threshold**

Enter a percentage of the maximum file size you want to set as your critical threshold.

- **Response Time (Advanced View)**

Enter the Warning and Critical Response Time thresholds for the length of time a service check takes. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

**4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click **Finish**.**

## Sybase

The Sybase monitor:

- Determines if the database is responding on the standard port.
- Sends Sybase/Transact-SQL scripts to the database for processing.

The Transact-SQL scripts can be very basic SQL statements, such as:

```
sphelp_db sampled1; exit (select 1);
```

The scripts can also be more complex statements that involve functions and other data processing.

## Configuring Sybase Monitors

To configure Sybase monitors, do the following:

- 1 **In the Sybase monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- Port

The number of the port number on which the database is listening. The default is 5000.

- Port Check (Optional)

Select this option to open a socket connection that determines whether or not the database is listening on the defined port.

- Username

The user name that is required to login to the database.

- Password

The password that is required to login to the database.

- Database

The name of the Sybase database to which you want to connect.

- Script

Click the **Script** checkbox and then type or copy the script that you want **up.time** to againt the database into this text box. Use this option if you do not have access to the file system on the monitoring station or if your script is short or will not regularly change.

- Script File

Click the **Script File** check box and then enter the full path on the monitoring station to the script that this monitor will run against the database.



If you configured your database to allow logins without a user name and password and you specify the script file but no login information, the script will fail. The script will run if you have configured your database to allow logins without a user name and password.

- Match (Regular Expression)

Enter a regular expression that you want to match against the string returned from the database. If the string matches, the status is OK. Otherwise, the status is Critical.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

**4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).



- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**



# CHAPTER 13

## Network Service Monitors

.....

The network service monitors track the health and performance of the following:

|                                     |     |
|-------------------------------------|-----|
| <i>DNS</i> .....                    | 238 |
| <i>FTP</i> .....                    | 241 |
| <i>HTTP (Web Services)</i> .....    | 243 |
| <i>IMAP (Email Retrieval)</i> ..... | 247 |
| <i>LDAP</i> .....                   | 249 |
| <i>NFS</i> .....                    | 253 |
| <i>NIS/YP</i> .....                 | 255 |
| <i>NNTP (Network News)</i> .....    | 257 |
| <i>Ping</i> .....                   | 261 |
| <i>POP (Email Retrieval)</i> .....  | 263 |
| <i>SSH (Secure Shell)</i> .....     | 265 |
| <i>SMTP (Email Delivery)</i> .....  | 267 |
| <i>SNMP</i> .....                   | 269 |
| <i>TCP</i> .....                    | 277 |

## DNS

DNS (Domain Name Server) is a distributed database that links various host names to specific Internet addresses. The DNS monitor determines the IP addresses of external and internal host names by matching a virtual host name to an expected IP address. If a match is made, the status of the service monitor is OK.

You can, for example, use the DNS monitor to:

- Ensure that your audience can access your Web site or portal by making sure that a selected address can be resolved.
- Identify instances in your network environment where resources have had their IP addresses changed, and now the resource is no longer available.

To collect performance information, the DNS monitor:

- Opens a UDP socket to a DNS server.
- Creates a query packet.
- Sends the query packet.
- Waits for a response.
- Parses the answers.

The DNS monitor does not check for the NS or MX records, which return names and not IP addresses. Non-authoritative answers as well as authoritative responses are used.

## Before You Begin

Before configuring the DNS monitor, determine the IP address for the host that you want to monitor. For internal hosts, you can use the `ipconfig` command from the command line.

The `ipconfig` command returns information similar to the following:

```
Connection-specific DNS Suffix . : uptimesoftware.com
IP Address : 10.1.1.42
Subnet Mask : 255.255.255.0
Default Gateway : 10.1.1.1
```

For external hosts, you can use the `nslookup` command from the command line as follows:

```
nslookup <host name>
```

The `nslookup` command returns information about the host, similar to the following:

```
Server: filter.uptimesoftware.com
Address: 10.1.1.100

Name: uptimesoftware.com
Addresses: 217.160.226.70, 10.1.1.95,
 192.168.23.1, 192.168.190.1
```

## Configuring DNS Monitors

To configure DNS monitors do the following:

- 1 In the DNS monitor template, select **Standard View** or **Advanced View**.

- 2 Complete the monitor information fields.

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:

- Hostname to Lookup

The host name that the monitor will check. The host name can be a Web site address, a server name, or a cluster name.

For example, for a Web site enter `www.uptimesoftware.com` in this field.

- Port

The number of the port on which the DNS server is listening. The default is 53.

- IP Address (Advanced View)

The IP address for which you want to check. If this address is not returned, the status of the service monitor becomes Critical.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the amount of time required to complete a service check. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Click the *Save for Graphing* checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

- 5 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 **Click *Finish*.**

## FTP

The FTP monitor can determine:

- Whether or not an FTP server is listening or is available on a specified port.
- The response time of an FTP server.

The FTP monitor tries to open an FTP connection to the server. If the response takes longer than the defined thresholds, [up.time](#) generates an alert.

## Configuring FTP Monitors

To configure FTP monitors, do the following:

- 1 In the FTP monitor template, select Standard View or Advanced View.**

- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:**

- Port

The number of the port number on which the FTP server is listening.  
The default is 21.

- Server Response (Advanced View)

Enter the Warning and Critical time thresholds required to receive a ready response from the FTP server. A server ready response can look like the following:

```
220 filter FTP server (Version wu-2.6.2(1) Mon Dec 3
15:29:55 EST 2005) ready
```

For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds for the length of time that the service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**
- 5 Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click Finish.**



## HTTP (Web Services)

The HTTP (Web Services) monitor simulates the steps that you take to access a Web site. Using this monitor, you can determine whether or not:

- You can access a Web site using HTTP.
- You can log on to a Web site.
- A Web site is running according to your expectations.

You can determine this by examining the values that are returned from the Web server.

The HTTP (Web Services) monitor relies on a Universal Resource Identifier (URI), which defines a specific file location on a Web server. This monitor can test for application calls, database responses, or any other information that a URI can return.

## Configuring HTTP (Web Services) Monitors

To configure HTTP (Web Services) monitors, do the following:

- 1 **In the HTTP (Web Services) monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- URI

The URI of the Web page that you want to monitor. For example, `/login.php`.

- Text to Look For (Optional)

Enter the text that you want the monitor to search for in the response from the server.

This monitor parses the text from the server and, using the threshold values you enter, determines if the entire Web page returned by the server is within acceptable parameters.

For example, if a Web page is returned then the monitor parses the entire page for the text that you input to match against. If you want to ensure that a particular page is returned, you could enter <TITLE>Expected Page</TITLE>, where Expected Page is the title of the Web page. The monitor generates an alert if this page is not matched.

- Authentication (Advanced View)

The user ID and password, in the form `userid:password`. For example:

```
jlamport:bluefrog5
```

- Virtual Host (Advanced View)

The unique domain name that resolves to the IP address of the domain that you want to monitor. A virtual host has its own domain name, but has the same IP address as other domain names hosted by the Web server.

- Server Response (Advanced View)

Enter a string to match against the response from the server. For example, `HTTP/1.1 200 OK` or `HTTP 404 - File not found`. Then, set the Warning and Critical comparison methods. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Follow Re-Direct Actions (Advanced View)

Select an action that enables you to specify whether or not you want to be redirected to another Web address:

- OK

Return an OK status for any re-direction.

- Follow

Follow any re-direction.

- Warning

Return a Warning status for any re-direction.

- Critical

Return a Critical status for any re-direction.

- POST String (Advanced View)

The URL-encoded POST string to be sent to the server. This string simulates what a Web browser sends to a Web server CGI script or binary. You can use the POST string to, for example, simulate logging into a Web application.

For example, if you define the POST string as `userid=bob&sku=123456`, the page to request would be `/cgi-bin/sku_lookup`. The text `SKU count` is the expected response. If the SKU lookup is not successful or if the response from the application server is not fast enough, then **up.time** generates an alert.

- Set-Cookie String (Advanced View)

Enter a cookie string, which can take the following form:

Set-Cookie: name=value; expires=date; path=pathname;  
domain=domainname; secure

Where:

- name is a name by which you can later reference the cookie.
- value is a regular string to be stored as a cookie. The string should be encoded using URL-style %xx encoding, which converts all reserved and unsafe characters – such as tildes (~) and spaces – to their ASCII equivalents.

For example, using %xx encoding, the URL `http://www.mydomain.com/~jdoe/index.html` becomes `http://www.mydomain.com/%7ejdoe/index.html`.

The name=value pair is the only required attribute of the **Set-Cookie String** field.

- expires is an optional attribute that specifies the expiration date and time for the cookie.
- HTTP Header Settings (Advanced View)

The HTTP header settings for the response.

The HTTP header settings define the syntax and semantics of all standard HTTP/1.1 header fields. For entity header fields, both sender and recipient refer to either the client or the server, depending on who sends and who receives the entity.

- Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 Click the Save for Graphing checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**
- 5 Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click Finish.**

## IMAP (Email Retrieval)

The IMAP (Email Retrieval) monitor determines whether or not:

- An IMAP server is listening on a defined port.
- An IMAP server is running on a defined system or on a group of systems.
- A particular version of IMAP is running.

### Configuring IMAP (Email Retrieval) Monitors

To configure IMAP (Email Retrieval) monitors, do the following:

- 1 **In the IMAP (Email Retrieval) monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- Port

The number of the port on which IMAP is listening.

The default is 143. If you are applying a monitor to a service group, ensure that all of the systems use the defined port. Otherwise, create a monitor for each IMAP instance that listens on a different port.

For information on service groups, see “Service Groups” on page 139.

- Server Response (Advanced View)

Select a comparison method, and then enter the Warning and Critical thresholds for the server response. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

The server response is the same for Windows, UNIX, and Linux. For example, an expected response is:

```
+OK [CAPABILITY IMAP4REV1 LOGIN-REFERRALS STARTTLS
AUTH=LOGIN] filter IMAP4rev1 2002.336 at Thu, 2 Jun 2005
10:55:02 -400 (EDT)
```

If IMAP is not available, then the following is an expected response:

```
BAD Null command
```

By making string comparisons on the returned values to the monitor, you can check:

- The version of IMAP that is running to support your network routing.
- The system on which IMAP is or is not running.
- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the length of time a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Click the *Save for Graphing* checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

- 5 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 **Click *Finish*.**

## LDAP

LDAP (Lightweight Directory Access Protocol) is a protocol that organizes directory hierarchies and enables communication with directory servers. Individuals in an organization can use LDAP to search for information, files, or devices in a network.

The LDAP monitor can check for any settings or information in your LDAP directory. The monitor can start the check from any location within your LDAP directory structure.

The LDAP monitor attempts to match information that you have specified with information available in your LDAP directory. If the monitor finds the information, the service monitor returns a status of OK. Otherwise, the monitor returns a Critical error and up.time generates an alert.



If you do not specify any parameters, then this monitor only validates that an LDAP server is listening on the specified port.

### Before You Begin

To configure the LDAP monitor, you should understand how an LDAP directory works, and know how LDAP is configured in your environment. You can use the following tools to determine the Base, Bind, and Attribute values of the LDAP directory for which you want to search:

- At the Windows command line, use `ntdsutil.exe` to retrieve information.
- Use one of the many freely-available LDAP browsing and editing tools.
- Use your own network documentation and determine whether or not the proper configurations have been maintained.

## Configuring LDAP Monitors

To configure LDAP monitors, do the following:

- 1 In the LDAP monitor template, select **Standard View** or **Advanced View**.

- 2 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:

- Port

The number of the port number on which the LDAP server is listening. The default is 389.


- Password

The password that is required to log in to the LDAP server.

- Base

The location in the LDAP directory from which you want the monitor to begin searching for information.

The following diagram shows a simple LDAP directory structure:





Using this directory structure, you can check your LDAP structure for your European employees by selecting the following as your base:

```
dc=ldap,dc=uptime,ou=employees,ou=Europe
```

- Bind

The Bind string, which associates user account properties and LDAP account attributes. This string gives you access to the Base location of your LDAP directory structure.

The format of the Bind string must match the Base location of your LDAP directory structure. For example, if you are checking for information found below the European employees directory, you can use the following Bind string:

```
cn=ldapadmin,dc=ldap,dc=uptime,dc=com
```

Depending on your network security model, you will need domain controller administration privileges to bind to the locations on which you want to match information.

- Attribute

The attribute or information for which you want to search in your LDAP directory.

An LDAP entry consists of a set of attributes. Each attribute has a type – which describes the kind of information contained in the attribute – and one or more values, which contain the actual data. For example, the entry `jsmith@inter.net` has the Attribute value `jsmith@inter.net`. The Attribute type is `e-mail`.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**

## NFS

NFS (Network File System) enables UNIX and Linux systems to share directories across a network. The NFS monitor can determine the performance of your NFS (Network File System) server and its ability to communicate with NFS clients by measuring the available NFS mounts.

This monitor runs the `showmount -e` command to extract the number of NFS file systems that are exported. If the `showmount` command fails, then [up.time](#) generates an alert.

## Configuring NFS Monitors

To configure NFS monitors, do the following:

- 1 **In the NFS monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- Mounts (Advanced View)

Select a comparison method, and then enter the Warning and Critical Mount thresholds for the number of mounts on which NFS is loaded. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the length of time a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.
- 5 Complete the following settings:
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click **Finish**.

## NIS/YP

NIS/YP (Network Information Services/Yellow Pages) is a distributed database system that enables you to configure multiple hosts from a central location as well as store and maintain common configuration information in that location. You can then propagate the information to all of the nodes in a network. The collection of network information is referred to as the *NIS namespace*.

The NIS/YP monitor performs a lookup on the domain, table, and key, and enables you to:

- Check that a Network Information Service (NIS) server for a given domain is responding.
- Request a specific key from a NIS table. This is useful if the contents of the NIS maps are often rebuilt.

## Configuring NIS/YP Monitors

To configure NIS/YP monitors, do the following:

- 1 **In the NIS/YP monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following NIS/YP monitor settings:**

- YP/NIS Domain

The domain of the NIS service. For example, `uptimesoftware.com`. NIS administration databases that contain name service information are called *maps*. A domain is a collection of systems that share a common set of NIS maps.

- YP/NIS Table

The name of the NIS/YP table that contains the values for which you want to search.

- Key

Enter a value you want to search for in the NIS table. For example, the key is `jsmith` in the following string returned from a NIS table:

```
jsmith:LLZDusFe5Da3s:20080:100:Jim Smith:
/export/home/jsmith:/bin/sh
```

- Lookup (Advanced View)

The Lookup value associated with the value in the **Key** field. For example, the following is returned from the `passwd` table of a NIS database based on the key `jsmith`:

```
jsmith:LLZDusFe5Da3s:20080:100:Jim Smith:
/export/home/jsmith:/bin/sh
```

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the length of time that a service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Click the `Save for Graphing` checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

- 5 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 **Click `Finish`.**

## NNTP (Network News)

NNTP is a protocol for distributing, searching, retrieving, and posting of messages and news articles from USENET (a global collection of online discussion groups). NNTP stores content in a central database, enabling subscribers to select only the messages and articles that they want to read.

The NNTP (Network News) monitor measures the performance of your NNTP server. It can also determine the server status in terms of the following:

- Command Implementation.
- Response Category.
- Response Codes.

### Command Implementation

Status reports from the server indicate the response to the last command that was received from the client. Status response lines begin with a three-digit numeric code, which is used to distinguish between all responses.

The first digit of the response broadly indicates the success, failure, or progress of the previous command.

- 1xx – An informative message.
- 2xx – The command is OK.
- 3xx – The command OK to this point, but the rest of it will be sent.
- 4xx – The command was correct, but could not be carried out.
- 5xx – The command is not implemented, or it is incorrect, or a serious program error has occurred.

## Response Category

The next digit in the status response code indicates the function response category.

- x0x – Connection, setup, and miscellaneous messages.
- x1x – Newsgroup selection.
- x2x – Article selection.
- x3x – Distribution functions.
- x4x – Posting.
- x8x – Nonstandard extensions.
- x9x – Debugging output.

## Response Codes

The following is a list of general response codes that may be sent by an NNTP server. These are not specific to any one command, but may be returned as the result of a connection, a failure, or an unusual condition.

- 100 – Help text.
- 190 through 199 – Debugging output.
- 200 – The server is ready and posting is allowed.
- 201 – The server is ready, but no posting is allowed.
- 400 – Service has been discontinued.
- 500 – The command is not recognized.
- 501 – A command syntax error occurred.
- 502 – An access restriction or permission is denied.
- 503 – A program fault occurred and the command was not executed.

You can ignore 1xx codes. Code 200 or 201 is sent upon initial connection to the NNTP server, depending upon the posting permission. Code 400 is sent when the NNTP server discontinues service – for example, by request of the operator. The 5xx codes indicate that the command could not be performed for some unusual reason.



## Configuring NNTP (Network News) Monitors

To configure NNTP (Network News) monitors, do the following:

- 1 In the NNTP (Network News) monitor template, select **Standard View** or **Advanced View**.

- 2 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:

- Port

The number of the port on which the NNTP server is listening. The default is 119.

- Server Response (Advanced View)

The server response according to the value that you want to measure.

For information on command implementation, see “Command Implementation” on page 257.

For information on response categories, see “Response Category” on page 258.

For information on general response, see “Response Codes” on page 258.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**

## Ping

The Ping monitor determines whether or not you can communicate with other IP addresses or domain names. The Ping monitor can check:

- Whether or not you can reach a specified system.
- The amount of time required to bounce a packet off of another site.

You will receive a response if the connections are good and the target system is running. If you have successfully pinged a system in the past, but you cannot get a response, there is a problem either with the network or with the system. If it takes a long time for a ping to return, the network or system may be extremely busy.

The ping program sends a small packet of information containing 64 bytes – 56 bytes of data and eight bytes of protocol reader information. The computer that sent the packet listens for a reply from the specified IP address. The ping program then evaluates this reply, and up.time captures the report that the program displays.

## Configuring Ping Monitors

To configure Ping monitors, do the following:

- 1 **In the Ping monitor template, select *Standard View* or *Advanced View*.**

- 2 **Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- Number to send

The number of packets to send to an IP address or domain name.

This value determines the number of times the ping command attempts to contact a server.

- Average Round Trip Time

Enter the Warning and Critical thresholds for the average round trip time for the number of packets sent by the ping command. The round trip time is in milliseconds.

This value is a good indicator of ping performance because a variety of factors – including different packet paths to and from the server – can affect the round trip time of a packet.

- Percent Loss (Advanced View)

Enter the Warning and Critical thresholds for the number of packets that did not returned a reply. For example, if four packets were sent and only two are returned, the percent loss is 50%.

- Response Time (Advanced View)

Enter the Warning and Critical Response thresholds for the length of time the service check takes to complete. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

**4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**6 Click Finish.**

## POP (Email Retrieval)

The POP (Email Retrieval) service monitor checks the status of POP2 servers (which requires SMTP to send messages) and POP3 servers.

Use the POP (Email Retrieval) monitor to determine whether or not:

- A POP server is listening on a defined port.
- A POP server is running on a defined system.
- A POP server is running on a group of systems.
- A particular version of POP is running.

## Configuring POP (Email Retrieval) Monitors

To configure POP (Email Retrieval) monitors, do the following:

- 1 **In the POP (Email Retrieval) monitor template, select Standard View or Advanced View.**

- 2 **Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 **Complete the following fields:**

- Expected Server Response (Advanced View)

Enter the response from the server, as a string, that determines whether or not a connection is made to the POP service. Then, set the Warning and Critical thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

The expected server response is the same for Windows, Solaris, and Linux. For example, if the POP service is available then the following is an expected response:

```
+OK POP3 <server name> v2002.81 server ready
```

If the POP service is not available, the following is an expected response:

```
-ERR Null command
```

- Response Time (Advanced View)  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 Click the Save for Graphing checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**
- 5 Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click Finish.**

## SSH (Secure Shell)

- SSH Server Version (Advanced View)

The version of the SSH server that you want to monitor. This is the string immediately following the major and minor version numbers of SSH. In the following example the SSH server version is `SUN_SSH1.0`:

```
SSH_2.0_SUN_SSH1.0
```

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds for the overall time required to perform a service check. For more information, “Configuring Warning and Critical Thresholds” on page 121.

- 4 **Click the `Save for Graphing` checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

- 5 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 6 **Click `Finish`.**



## SMTP (Email Delivery)

The SMTP monitor tests a mail server for the standard mail response header. If the mail server does not respond within the specified thresholds, up.time generates an alert.

### Configuring SMTP (Email Delivery) Monitors

To configure SMTP (Email Delivery) Monitors, do the following:

- 1 In the **SMTP (Mail Delivery) monitor template**, select **Standard View** or **Advanced View**.

- 2 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:

- Port

The number of the port on which the SMTP server is listening. The default is 25.

- Expected Server Response (Advanced View)

Enter the Warning and Critical thresholds for the amount of time that is required to send and receive a ready response from the SMTP server.

For example, the following response reveals the ready status of the SMTP server:

```
220 mail.yourdomain.com ESMTP
Sendmail 8.12.10+SUN/8.12.8;
Tue, 14 Dec 2005 13:25:15: -0400 <EDT>
```

For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

- 4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.
- 5 Complete the following settings:
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
  - Action Profile settings (see “Action Profiles” on page 134 for more information).
- 6 Click **Finish**.

## SNMP

Simple Network Management Protocol (SNMP) is a widely-used protocol that monitors the health of computer and network equipment. The SNMP monitor enables you to query SNMP devices or systems for a given object identifier (OID) of an SNMP Management Information Base (MIB). A MIB is a listing that defines variables needed by the SNMP protocol to monitor and control network equipment.

The OIDs identify the managed variables in a system. Each OID is represented by a set of numbers separated by periods – for example, .1.3.6.1.2.1.1.1.0. The period at the start of an OID indicates that the name of the OID begins at the root of its associated MIB. However, each object is also assigned a unique name – for example, sysObjectID – that makes it easier to identify that object.

The SNMP monitor enables you to compare the response to a specific pattern. If the device is protected by a community password, you can specify the password in the monitor parameters. The default OID that you specify should be the Enterprise identification string.

## Net-SNMP

The **up.time** SNMP monitor also supports Net-SNMP, which is a suite of command line and graphical applications that:

- Request information from SNMP agents.
- Set information on SNMP agents.
- Generate and handle SNMP traps.

To take advantage of the Net-SNMP features, you must:

- Install and configure the Net-SNMP application suite on your server. Visit <http://net-snmp.sourceforge.net> for more information:
- Have a Net-SNMP agent already installed on the host or hosts that you want to monitor. The Net-SNMP HOST-RESOURCES-MIB (used to gather performance statistics from a host) must also be enabled. See the Net-SNMP documentation for details.
- Add a Net-SNMP entity to **up.time**. For more information, see “Adding Systems or Network Devices” on page 78.

## SNMP MIB Browser

The SNMP monitor uses the SNMP MIB Browser to load OIDs from MIBs on your system or on a server. The first step in setting up an SNMP monitor is to use the [up.time](#) SNMP MIB Browser applet to:

- Load MIBs into [up.time](#).
- Select managed objects.

## Supported Versions of SNMP

The [up.time](#) SNMP monitor works with the following versions of SNMP:

- v1  
The first implementation of the SNMP protocol.
- v2  
The second implementation of the SNMP protocol, which contains additional protocol operations as well as improved security and data authentication.
- v3  
The latest implementation of the SNMP protocol, which adds security and privacy features that are missing in versions 1 and 2 of the protocol.

## Using the SNMP MIB Browser

The SNMP MIB Browser is a Java applet that enables you to locate MIBs and their OIDs (object identifiers) on your local file system or your network. Use the SNMP MIB Browser for:

- Loading MIBs from a File or a Server
- Adding OIDs
- Deleting OIDs



The MIB Browser requires version 1.5 of the Java Web browser plugin. [up.time](#) will install the newer Java plugin if it detects that your computer has version 1.4.2 or earlier of the plugin installed.

## Loading MIBs from a File or a Server

You can load MIBs and their associated OIDs into [up.time](#) from your computer or from a server. Once you have loaded the MIBs, you can select the OIDs that you want monitored by the SNMP service monitor.

To load MIBs from a file or a server, do the following:

- 1 From the [up.time](#) tool bar, select **Services**.
- 2 In the **Tree** panel, click **Add New Service Instance**.
- 3 In the **Add New Service Instance** window, click **List agentless up.time monitors**, then click **SNMP**, and then click **Continue**.

The SNMP MIB browser applet appears.



If a Java security warning dialog box appears while the applet is loading, click **Always** or **Accept** (depending on your Web browser) to close the dialog box.

- 4 In the **SNMP MIB Browser**, click one of the following options:
  - Load MIB from File
  - Load MIB from Server
- 5 In the window that appears, do one of the following:
  - If you are loading a MIB from your computer, navigate to the directory containing the MIB or OID. Select the MIB, and then click **Open**.
  - If you are loading a MIB from a server, select the MIB from the list that appears, and then click **Load Selected MIB**.

The MIB appears in the MIB selection tree. You can select any OID within the MIB to monitor with the SNMP service monitor.

### Adding OIDs

Once a MIB is loaded into the MIB selection tree, you can add the OIDs in the MIB to the SNMP monitor.

To add OIDs, do the following:

- 1 **Navigate the MIB directory tree to find the OID that you want to add.**

- 2 **Double click the OID.**

The OID appears in the **Selected OIDs** panel.

- 3 **Click Next.**

The **Add SNMP Service Monitor** window appears. See “Configuring SNMP Monitors” on page 273 for information on setting up the SNMP monitor.

### Manually Adding OIDs

If you know the OID that you want to add, you can add it without navigating the MIB tree.

To add OIDs manually, do the following:

- 1 **Type the name of the OID in the Add OID Manually field.**

- 2 **Click Add OID Manually.**

- 3 **Click Next.**

The **Add SNMP Service Monitor** window appears. See “Configuring SNMP Monitors” on page 273 for information on setting up the SNMP monitor.

### Deleting OIDs

After adding several OIDs, there may be OIDs that you no longer want to monitor. You can use the SNMP MIB browser to delete the unwanted OIDs.

To delete OIDs from the **Selected OIDs** panel, do the following:

- 1 **Select the OID you want to remove in the Selected OIDs panel.**

- 2 **Click Delete Selection.**

## Configuring SNMP Monitors

To configure SNMP monitors, do the following:

- 1 In the **SNMP monitor template**, select **Standard View** or **Advanced View**.
- 2 Select the version number of an **SNMP** implementation from the **SNMP Version** dropdown list.
- 3 In the **v1/v2 Community** field, enter the community string.

The community string acts like a user ID or password, giving you access to a device via SNMP. Common communities are **public** (enables you to retrieve read-only information from the device) and **private** (enables you to access all information on the device).

- 4 Enter the number of the port on which **SNMP** is listening in the **SNMP Port** field.
- 5 If you selected **v3** from the **SNMP Version** dropdown list, complete the following settings:

- v3 Username

The user name that is required to connect to an SNMP instance that is using version 3 of SNMP.

- v3 Authentication Method

If the server uses version 3 of SNMP, select one of the following options from the list. The option that you select determines how encrypted information travelling between the SNMP instance and **up.time** will be authenticated:

- MD5

A widely-used method for creating digital signatures that are used to authenticate and verify the integrity of data.

- SHA

A secure method of creating digital signatures. SHA is considered the successor of MD5 and is widely used with network and Internet data transfer protocols.



Ensure that the authentication method you select in **up.time** matches the method that is used by the system you want to monitor.

- v3 Auth Password

The password that is required to connect to an SNMP instance that is using version 3 of SNMP.

- v3 Privacy Method

If the server uses version 3 of SNMP, select one of the following options from the list. The option that you select determines how information travelling between the SNMP instance and **up.time** will be encrypted:

- DES

An older method used to encrypt information. DES is considered weak compared to more modern encryption methods.

- AES

The successor to DES, which is used with a variety of software that require encryption including SSL servers.



Ensure that the privacy method that you select in **up.time** matches the method that is used by the system you want to monitor.

- v3 Privacy Password

The password that will be used to encrypt information travelling between an SNMP instance that is using version 3 of SNMP and **up.time**.




## 6 Complete the following fields:

- Warning and Critical Thresholds

Enter the Warning and Critical thresholds for each OID that you added using the SNMP MIB Browser. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

Each OID has one or more settings associated with it, as shown in the following image:



The screenshot shows a configuration interface for SNMP traps. It is divided into two main sections: **snmpTrapOID** (with OID .1.3.6.1.6.3.1.1.4.1) and **snmpTrapEnterprise** (with OID .1.3.6.1.6.3.1.1.4.3). Each section contains two rows: **Warning** and **Critical**. For each row, there is a **Comparison Method** dropdown menu and a text input field for the threshold value.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

## 7 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.

## 8 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).

- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information.)
- Action Profile settings (see “Action Profiles” on page 134 for more information).

**9 Click Finish.**

## TCP

The TCP monitor can determine whether or not a service or application is listening on a specific port. This monitor can also execute commands against an application or a service listening on a port and evaluate the result.

By extending the TCP monitor to evaluate the returned string based on a command over a network using TCP, you can test and monitor for a wide variety of responses.

For example, to have [up.time](#) generate an alert if the file `Weekly_Report` was changed in your source code control system, you can send the string:

```
get -e Weekly_Report1
```

and set the critical threshold value to `1.2`, where `1.1` represents no changes and `1.2` or greater represents one or more changes to the document.

## Configuring TCP Monitors

To configure TCP monitors, do the following:

- 1 In the TCP monitor template, select **Standard View** or **Advanced View**.

- 2 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:

- Port

The number of the port on which the service or application that you want to monitor is listening.



To check whether or not an application is listening on a port, leave the remaining TCP service monitor settings blank.

- **String to Send**  
The string that contains the command to which the service or application can respond.
- **Use SSL**  
Select this option if your connection uses SSL (Secure Sockets Layer) for security.
- **String to Receive**  
The string that is returned by the specified port and host. The string is the response to the command that was specified in the String to Send field.
- **Response Time (Advanced View)**  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

**4 Click the **Save for Graphing** checkbox to save the data for a metric to the DataStore, which can be used to generate a report or graph.**

**5 Complete the following settings:**

- **Timing Settings** (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- **Alert Settings** (see “Monitor Alert Settings” on page 125 for more information).
- **Monitoring Period settings** (see “Monitor Timing Settings” on page 122 for more information).
- **Alert Profile settings** (see “Alert Profiles” on page 129 for more information.)
- **Action Profile settings** (see “Action Profiles” on page 134 for more information).

**6 Click **Finish**.**

# CHAPTER 14

## Advanced Monitors

.....

You can configure advanced monitors to collect performance information. Advanced monitors are described in the following sections:

|                                        |            |
|----------------------------------------|------------|
| <i>Overview .....</i>                  | <i>280</i> |
| <i>Custom Monitors .....</i>           | <i>282</i> |
| <i>Custom with Retained Data .....</i> | <i>284</i> |
| <i>External Check .....</i>            | <i>286</i> |

## Overview

The standard [up.time](#) service monitors may not be adequate to monitor all of the systems, applications and proprietary devices in your environment. In some cases, you may need to capture metrics that are unique to your environment. To do this, you can define advanced service monitors.

Advanced monitors can be simple scripts that run service checks on a host. You can write a shell script, or use a higher-level scripting language like Perl, Python, or Ruby. Or, the advanced monitors can be binary programs that interact with more sophisticated applications. On top of that, advanced monitors do not require an agent to be installed on the system that you are monitoring.

Regardless of how you develop your advanced monitor scripts or programs, those scripts or programs should return the following codes:

- 0 – OK  
The services are functioning properly.
- 1 – Warning  
There is a potential problem with one of more of the services being monitored.
- 2 – Critical  
There is a critical problem with one or more of the services being monitored.
- 3 – Unknown  
There is an error in the configuration of the monitor itself, or [up.time](#) cannot execute the service check.

[up.time](#) captures the output from the script or program, usually from standard output (`stdout`). The output appears in the service status section of the **Global Scan** panel (see “Understanding the Status of Services” on page 25). The [up.time](#) monitoring framework picks up any error codes and triggers the appropriate monitoring action.

If you have already written scripts or programs for other monitoring tools, you can re-use those scripts or programs with [up.time](#). You simply point your advanced monitor to where your scripts or programs are located and [up.time](#) will run them.

The uptime user account on the [up.time](#) monitoring station must be able to execute the script or program that you use.



Contact uptime software Client Care for help with creating advanced monitor scripts.

## Before You Begin

When creating a script or an executable for an advanced monitor, you should ensure that:

- The necessary interpreter for the scripting language that you are using is installed on the monitoring station.
- You have determined the arguments that the script or program requires, and the parameters that you want your script or program to return.
- You use forward slashes when specifying directory paths in your scripts, regardless of the operating system. For example, use C:/ on Windows, or /opt on Solaris or Linux.

Many of the fields that you use to define an advanced monitor are the same as those used with agent and agentless monitors. You can find more information about those fields in the following sections:

- To learn how to access the custom monitor definition window, see “Using Agentless Monitors” on page 114.
- For a description of monitor identification information fields, see “Monitor Identification” on page 117.
- For a description of monitor timing settings, see “Monitor Timing Settings” on page 122.
- For a description of alert settings, see “Monitor Alert Settings” on page 125.
- For a description of alert profiles, see “Alert Profiles” on page 129.
- For a description of action profile, see “Action Profiles” on page 134.

## Custom Monitors

A Custom monitor runs a script that captures information which is related to a situation that may be unique to your environment. When the script is run, the system being monitored returns a single line of information to standard output (`stdout`). The script reads `stdout`, which may contain an error or return value. This error or return value is then displayed in the [up.time](#) monitoring station.

As well, you can specify that the monitor writes the data that the script returns to the [up.time](#) DataStore. You can use the retained data to later generate a Service Metrics report (see “Service Monitor Metrics Report” on page 329) or a Service Metrics graph (see “Viewing System and Service Information” on page 57).

## Configuring Custom Monitors

To configure Custom monitors, do the following:

- 1 In the Custom monitor template, select Standard View or Advanced View.**

- 2 Complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete following fields:**

- Script Name

The name of, and path to, the script or program on the monitoring station that will collect metrics.



The uptime user account on the [up.time](#) monitoring station must be able to execute the script or program that you use. Ensure that the permissions for the uptime user account are set correctly.

- Arguments (Optional)

Specify any arguments that are required by the script or program.



- **Output (Optional)**

Specify a comparison method to override the settings of an alert profile, or to return only the most severe errors.

Do this by selecting an option from the **Comparison Method** dropdown lists beside the **Warning** and **Critical** fields. Then, enter a value in the field. For example, to return only unknown errors you can select **Exactly Matches** from the dropdown list, and type UNKNOWN in the field.

For more information on comparison methods, see “Comparison Methods” on page 120.
  - **Response Time (Advanced View)**

Optionally, enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.
- 4 **Click the **Save for Graphing** option to save the output in the DataStore. You can later use the retained data to generate a report or a graph.**
  - 5 **Complete the following settings:**
    - Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
    - Monitoring Period settings (see “Monitor Timing Settings” on page 122 Monitor Timing Settings for more information).
    - Alert Profile settings (see “Alert Profiles” on page 129 for more information).
    - Action Profile settings (see “Action Profiles” on page 134 for more information).
  - 6 **Click **Finish**.**

## Custom with Retained Data

Custom monitors with Retained Data return the following information:

- Up to 10 values that you can save and evaluate.
- A return status of 0 to 3. See “Overview” on page 280 for more information.

As well, you can specify that the monitor writes any returned data to the [up.time](#) DataStore. You can use the retained data to later generate a Service Metrics report (see “Service Monitor Metrics Report” on page 329) or a Service Metrics graph (see “Viewing System and Service Information” on page 57).

## Configuring Custom Monitors with Retained Data

To configure Custom monitors with Retained Data, do the following:

- 1 In the Custom with Retained Data monitor template, select Standard View or Advanced View.**

- 2 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 3 Complete the following fields:**

- Script Name

The name of, and path to, the script or program on the monitoring station that will collect metrics on the system.



The script or program that you specify must be executable by the [uptime](#) user account on the [up.time](#) monitoring station. Ensure that the permissions are set correctly.

- Arguments (Optional)

Specify any arguments required by the script or program.

- Variable 1 to Variable 10 (Optional)

Specify up to 10 variables that your custom script will return to the [up.time](#) monitoring station. If you click the **Save for Graphing** checkbox, these variables will be saved to the DataStore.

- Response Time (Advanced View)

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 121.

#### 4 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 122Monitor Timing Settings for more information)
- Alert Profile settings (see “Alert Profiles” on page 129 for more information)
- Action Profile settings (see “Action Profiles” on page 134 for more information)

#### 5 Click Finish.

## External Check

The External Check monitor captures asynchronous events. [up.time](#) does not actively monitor these events by polling or initiating service checks. Instead, External Check monitors rely on an external event to generate the information that the monitors capture. External Check monitors enable you to determine when to collect service data for the event that you specify.

After you define an External Check monitor, the monitor runs a Perl script named `extevent.pl`. The script `extevent.pl` is included with [up.time](#), in the `scripts` subfolder. When it is run, the script connects to the port on which the server is listening. It then triggers the application on the server that generates the external event that is sent to [up.time](#).

This script `extevent.pl` has the following command line syntax:

```
extevent.pl --host=Hostname --port=PortNumber
--status=StatusNumber --message=message
--monitorName=name
```

Where:

- **host**  
The host name of the server that is running [up.time](#).
- **port**  
The [up.time](#) port on the server (usually 9996).
- **status**  
The status of the service being monitored. See “Overview” on page 280 for more information.
- **message**  
A human readable diagnostic message.
- **monitorName**  
The name of the service monitor to which the output will be returned.



Before using an External Check monitor, contact uptime software Client Care for assistance. You will need specific instructions for configuring this monitor depending on the nature of the applications that will be generating asynchronous events for [up.time](#).

## Configuring External Check Monitors

To configure External Check monitors, do the following:

- 1 **In the External Check monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 117.

- 2 **Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 124 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 125 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 122 for more information).
- Alert Profile settings (see “Alert Profiles” on page 129 for more information).
- Action Profile settings (see “Action Profiles” on page 134 for more information).

- 3 **Click Finish.**



# CHAPTER 15

## Configuring Users

.....

This chapter describes the [up.time](#) user management functions in the following sections:

|                                               |            |
|-----------------------------------------------|------------|
| <i>Working with User Roles .....</i>          | <i>290</i> |
| <i>Working with Users .....</i>               | <i>293</i> |
| <i>Working with User Groups .....</i>         | <i>297</i> |
| <i>Working with Notification Groups .....</i> | <i>300</i> |

## Working with User Roles

User roles define the following:

- What a user will see in the [up.time](#) Web interface.
- The items that a user can add, view, edit, or delete in the [up.time](#) Web interface.

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:**
  - User
  - Entities
  - Services
  - Entity Groups
  - Action Profiles
  - Alert Profiles



**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 13 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 307 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. The assigned permissions are denoted by a green check mark (✓), as shown below:

| Permission                | View    | Add | Edit | Delete |
|---------------------------|---------|-----|------|--------|
| <b>Users</b>              | ✓       | -   | -    | -      |
| <b>Entities</b>           | ✓       | -   | -    | -      |
| <b>Services</b>           | ✓       | -   | -    | -      |
| <b>Entity Groups</b>      | ✓       | -   | -    | -      |
| <b>Action Profiles</b>    | ✓       | -   | -    | -      |
| <b>Alert Profiles</b>     | ✓       | -   | -    | -      |
|                           |         |     |      |        |
| Permission                | Allowed |     |      |        |
| <b>Administrator</b>      | -       |     |      |        |
| <b>Acknowledge Alerts</b> | -       |     |      |        |
| <b>Save Reports</b>       | ✓       |     |      |        |

## **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 290.**

## 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.
- 3 Enter the password, which will be used to log into [up.time](#), in the **Password** field.
- 4 Re-enter the password 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
  - 5:30pm to 7:30am weekdays and all weekend till Monday morning

- 12:00am to 12:30am Monday

- 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,

**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 the services being monitored.

- Alert On Warning

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

- 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 290.

**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 297.

**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 293.

## 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:

- 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 Select the users to add to the group in the **Available Users** list, and then click **Add**.



You must always add the user admin to a new user group.

- 4 Optionally, select one of the systems or elements from the **Available Elements** list, and then click **Add**.
- 5 Optionally, select one of the groups from the **Available Element Groups** list, and then click **Add**.
- 6 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**.
- 2 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.



- 3 Edit the information as described in the section “Adding User Groups” on page 298.

## 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**.

## 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 129 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 129.
- 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, and then click Add.**
- 6 Select one or more users to add to the group from the Available Users list, and then click Add.**
- 7 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 **Users** subpanel.
- 3 **Click the name of the notification group that you want to view.**  
The details of the group appear in the **Users** 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 name of 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 300.**



# CHAPTER 16

## Understanding Report Options

.....

This chapter is an overview of the options available for generating reports in [up.time](#), and contains the following sections:

|                                 |            |
|---------------------------------|------------|
| <i>Overview .....</i>           | <i>304</i> |
| <i>Generating Reports.....</i>  | <i>305</i> |
| <i>Scheduling Reports .....</i> | <i>310</i> |
| <i>The Report Log.....</i>      | <i>313</i> |

### 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 your 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 317.

## 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 310.



If you do not receive a scheduled report, check the Report Log (see “The Report Log” on page 313) 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 an **up.time** user group.

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:

Email address: `jsmith@mydomain.com, rjones@mydomain.com, admin@int`

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 310 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/uptime4/`. 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:

- 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 date and time for the report to run.**

For more information on scheduling reports, see “Scheduling Reports” on page 310.

- 6 **Click Save Report.**

## Viewing Saved Reports

You can quickly view any reports that were generated on the 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 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 mount 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:



A screenshot of a web interface showing a checkbox labeled "Scheduled Report (Run at" which is checked. To its right are two dropdown menus separated by a colon. The first dropdown menu shows the number "15" and the second shows "30".

- 5 Select one of the following options:**


- Daily



A screenshot of a web interface showing two radio button options. The first option, "Every 1 day(s)", is selected with a green dot. The second option, "Every Weekday", is unselected with a blue dot. The number "1" is in a dropdown menu next to "Every".

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, and then:
  - 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:

| Pending Reports                                                                                                   |                    |            |                              |
|-------------------------------------------------------------------------------------------------------------------|--------------------|------------|------------------------------|
| Report Name                                                                                                       | Report Description | Scheduled? | Next Run Time                |
|  Emailing report in PDF format |                    | ✗ No       | Mon Oct 17 15:54:48 EDT 2005 |
|  Emailing report in PDF format |                    | ✗ No       | Mon Oct 17 15:55:29 EDT 2005 |

- Running Reports

Reports that are being run. This section contains the same information as the **Pending Reports** section, as illustrated below:

| Running Reports               |                    |            |                              |
|-------------------------------|--------------------|------------|------------------------------|
| Report Name                   | Report Description | Scheduled? | Next Run Time                |
| Emailing report in PDF format |                    | ✗ No       | Tue Nov 01 09:36:11 EST 2005 |
| Workload - Linux              |                    | ✗ No       | Mon Oct 31 14:16:14 EST 2005 |

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:

| Completed Reports                                                                                               |                              |                              |          |                                      | Remove Completed Reports |
|-----------------------------------------------------------------------------------------------------------------|------------------------------|------------------------------|----------|--------------------------------------|--------------------------|
| Report Name                                                                                                     | Started                      | Ended                        | Status   | Status Message                       |                          |
|  Emailing report in PDF format | Mon Oct 17 14:58:02 EDT 2005 | Mon Oct 17 14:58:40 EDT 2005 | finished | Sending to: [REDACTED]<br>Email sent |                          |
|  Emailing report in PDF format | Mon Oct 17 15:52:39 EDT 2005 | Mon Oct 17 15:53:16 EDT 2005 | finished | Sending to: [REDACTED]<br>Email sent |                          |
|  Emailing report in PDF format | Mon Oct 17 15:53:39 EDT 2005 | Mon Oct 17 15:54:09 EDT 2005 | finished | Sending to: [REDACTED]<br>Email sent |                          |

## 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**.



# CHAPTER 17

## Using Reports

.....

This chapter describes the reporting features of [up.time](#) in the following sections:

|                                                   |            |
|---------------------------------------------------|------------|
| <i>Reports for Performance and Analysis .....</i> | <i>318</i> |
| <i>Reports for Capacity Planning .....</i>        | <i>332</i> |
| <i>Reports for Availability .....</i>             | <i>357</i> |
| <i>Reports for J2EE Applications .....</i>        | <i>361</i> |
| <i>Reports for Virtual Environments .....</i>     | <i>368</i> |

## 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 27.

### 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.
- The average queue length.
- The number of reads and writes per second.
- The number of blocks being accessed per second.
- The average wait time, in seconds.
- The 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

- 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:



The screenshot shows a configuration window titled "Daily Hours". Inside, there is a text label "Include data samples between these hours only :". Below this label are two dropdown menus. The first dropdown is labeled "Start:" and the second is labeled "End:". Both dropdown menus currently display "N/A".

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 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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.



- 6 To generate reports for one or more views, select the groups from the **List of Views** area.  
See “Working with Views” on page 93 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 305 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 310 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. You must:

- 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 27.

- 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 93 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.
- 7 Click the checkboxes in the **Select** column to select the variables on which you want to report as shown below:

| Current Retained Service Metrics |                      |                          |                 |       |           |
|----------------------------------|----------------------|--------------------------|-----------------|-------|-----------|
| Instance Name                    | Instance Description | Select                   | Variable        | Units | Data Type |
| FS Capacity - Filter             |                      |                          |                 |       |           |
|                                  |                      | <input type="checkbox"/> | Response time ▾ | ms    | integer   |

- 8 Optionally, select one of the following:
  - Show all non-ranged metrics on one chart  
This option combines all of the variables that you 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** du.click the

## 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:

- Service Monitor Outages 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

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

- 3 Optionally, in the **Exclude file system names like** field enter either the name of a file system or a regular expression that [up.time](#) will use to ignore certain file systems when generating the report.

For example, if you want to exclude the /boot file system from the report, enter /boot in the field.

- 4 **Optionally, enter a value in the Exclude filesystems over % full field.**

This value is expressed as a percentage. The report displays the information for file systems whose used disk space is less than the amount you enter in this field. For example, if you set this field to 45, the report only displays file systems whose percentage used values are less than or equal to 45%.

- 5 **Click the Show totals for each system only checkbox to report only on the total amount by which all file systems on all disks drives have grown, rather than displaying amounts for each file system.**

- 6 **If you want to generate reports for systems in specific groups, select the groups from the List of Groups area.**

- 7 **To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

- 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. For more information, see page 89.

- 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 93 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 305 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 310 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) stalls. 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 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 27.

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.

The average number of mutex stalls that exceed the value in this field 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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: 2007-05-01 00:00:00 to 2007-05-23 13:06:36 between the time range: 00:00 to 23:59 |                |              |
| Multi CPU systems with an Average SMTX over 20.0 are highlighted                              |                |              |
| System Name                                                                                   | Number of CPUs | Average SMTX |
| Opteron                                                                                       | 2              | 38.33        |
| Filter                                                                                        | 4              | 17.29        |
| vmh-t1k-a                                                                                     | 24             | 3.58         |

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 392 for more information)
- Run Queue Length (see page 390 for more information)
- Run Queue Occupancy (see page 390 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. The amount of data moving through each interface is measured in megabytes.

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 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.

However, the following systems store that information 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.

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 27.

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:



The screenshot shows a configuration window titled "Daily Hours". Inside, there is a label "Include data samples between these hours only :". Below this label are two dropdown menus. The first is labeled "Start:" and the second is labeled "End:". Both dropdown menus currently display "N/A".

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 93 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 305 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 307 and “Scheduling Reports” on page 310 for more information.

## Using the Network Bandwidth Report

The following is an example of a Network Bandwidth report:

| Network Bandwidth Report                                                                      |           |             |              |
|-----------------------------------------------------------------------------------------------|-----------|-------------|--------------|
| Date Range: 2007-05-13 00:00:00 to 2007-05-19 23:59:59 between the time range: 00:00 to 23:59 |           |             |              |
| Bytes per packet: 750                                                                         |           |             |              |
| Hostname                                                                                      | Interface | Total MB In | Total MB Out |
| Filter                                                                                        | hme0      | 44,669.64   | 27,317.53    |
|                                                                                               | hme1      | 0.00        | 0.00         |
|                                                                                               | hme2      | 0.00        | 0.00         |
| Brightmail                                                                                    | netif0    | 1,276.92    | 1,276.92     |
|                                                                                               | netif1    | 6,183.59    | 1,013.10     |
| McKay                                                                                         | eth0      | 406.95      | 231.75       |
| FogBugz                                                                                       | netif0    | 8.38        | 8.38         |
|                                                                                               | netif1    | 1,616.91    | 1,495.04     |

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 407 for more information) to get a better idea of why network I/O is so high on the system.

## Disk I/O Bandwidth Report

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

The report contains the following information:

- The display name of the system in [up.time](#).
- The names of each disk on the system.
- Where applicable, the name of the file system on the disk.
- The total amount of data, measured in megabytes, that is being read from and written to the disk.

## Using Regular Expressions

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

|                                                               |                                                              |
|---------------------------------------------------------------|--------------------------------------------------------------|
| <b>Exclude Disks:</b><br><code>sda</code>                     | <b>Exceptions:</b><br><code>sdb</code>                       |
| Note: you can enter regular expressions into these fields.    |                                                              |
| <b>Exclude File Systems:</b><br><code>/u[0-9]<br/>/usr</code> | <b>Exceptions:</b><br><code>/proc<br/>/home<br/>/ex.*</code> |
| Note: you can enter regular expressions into these fields.    |                                                              |

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

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

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

## Generating a Disk I/O Bandwidth Report

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


- 1 In the **Reports Tree** panel, click **Disk I/O Bandwidth**.
- 2 In the **Date and Time Range** area, select the dates and times on which to report.

For more information, see “Understanding Dates and Times” on page 27.



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 346 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 346 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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: 2007-05-01 00:00:00 to 2007-05-23 13:55:56 between the time range: 00:00 to 23:59 |           |                |               |
| Output displayed in megabytes, 512 bytes per block                                            |           |                |               |
| Hostname                                                                                      | Disk Name | Filesystem     | I/O Total     |
| Brightmail                                                                                    | 0         | C:             | 13,857,837.70 |
|                                                                                               | 1         | D:             | 4,186,010.56  |
| FogBugz                                                                                       | 0         | C:             | 4,492,642.14  |
|                                                                                               | 1         | E:             | 208,868.71    |
| WebLogic Server                                                                               | 0         | C:             | 7,716,554.50  |
| devprod                                                                                       | dad0      | /, /data, /var | 104.82        |

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 409 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 determines the number of jobs that are currently ready to run. If the length of the run queue is appreciably greater than the number of available processors, 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 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 27.

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 Enter the length of the CPU run queue threshold in the Threshold field.**

The threshold is amount that you enter in the field multiplied by the number of CPUs on the 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 93 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 305 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 307 and “Scheduling Reports” on page 310 for more information.

### Using the CPU Run Queue Threshold Report

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



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 388 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 enables you to pinpoint systems that are spending considerable amounts of time waiting for disk access operations to complete.

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 percentile column for each disk or file system. For example, a percentile of 50% is the median service time value. A percentile of 100% indicates the highest service time value.

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 27.

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:



The screenshot shows a window titled "Daily Hours". Inside, there is a label "Include data samples between these hours only:". Below this label are two dropdown menus. The first dropdown is labeled "Start:" and the second is labeled "End:". Both dropdown menus currently display "N/A".

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 346 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 346 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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: 2007-05-13 00:00:00 to 2007-05-19 23:59:59 between the time range: 00:00 to 23:59 |         |            |                             |      |         |                 |
| Sorted by Descending High Service Time                                                        |         |            |                             |      |         |                 |
| Showing Percentile: 95.0, Threshold 20ms                                                      |         |            |                             |      |         |                 |
| System Name                                                                                   | Disk    | Filesystem | Service Time (milliseconds) |      |         |                 |
|                                                                                               |         |            | High                        | Low  | Average | 95th Percentile |
| Filter (filter)                                                                               | md13    | /export    | 85.00                       | 2.00 | 10.50   | 20.00           |
|                                                                                               | md2     |            | 85.00                       | 3.00 | 12.20   | 21.85           |
|                                                                                               | c0t4d0  |            | 81.00                       | 0.00 | 10.20   | 20.00           |
|                                                                                               | md1     |            | 67.00                       | 0.00 | 16.56   | 29.00           |
|                                                                                               | md21    |            | 65.00                       | 0.00 | 15.50   | 28.00           |
|                                                                                               | md11    |            | 63.00                       | 0.00 | 15.51   | 29.00           |
|                                                                                               | c1t13d0 |            | 58.00                       | 2.00 | 16.51   | 27.00           |
|                                                                                               | c1t10d0 |            | 56.00                       | 2.00 | 16.44   | 27.00           |
|                                                                                               | md23    |            | 54.00                       | 3.00 | 10.80   | 21.00           |
|                                                                                               | md0     |            | 31.00                       | 0.00 | 0.72    | 0.00            |
| McKay (mckay)                                                                                 | sda3    | /          | 30.00                       | 0.00 | 0.22    | 1.00            |
| Filter (filter)                                                                               | md20    |            | 30.00                       | 0.00 | 0.66    | 0.00            |
|                                                                                               | md10    |            | 28.00                       | 0.00 | 0.67    | 0.00            |
| McKay (mckay)                                                                                 | sda2    |            | 20.00                       | 0.00 | 1.80    | 4.00            |

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 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
- 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 90.

### 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 for more information.

## Service Monitor Availability Report

The Service Monitor Availability report tracks the status of the services associated with the hosts in your environment. This report lists the percentage of time each service was in the following states over the time period that you specify: OK, Warning, Critical, Maintenance, or Unknown. For more information on each status, see “Understanding the Status of Services” on page 25.

## 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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.

### WebSphere Server - Connection Pool - Pool size



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 182 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 27.

- 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 93 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 305 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 307 and “Scheduling Reports” on page 310 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.

WebLogic Server - JVM - Heap Size



WebLogic Server - JVM - Free Memory



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
- VI3 Workload Profile 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*.

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 enables you to pinpoint which VMware instances are using the most system resources.

You can also use the VMware Workload report to determine whether or not you are using a particular VMware server to its optimal capacity. For an example of this, see “Using the VMware Workload Report” on page 370.

### 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 27.

**3 In the Report Options section, select one of the following:**

- Workload - 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 - Memory

The amount of physical memory, in kilobytes, that is being used by a VMware instance.

- Workload - Disk IO

The amount of the disk I/O capacity, in kilobytes per second, that is being used by a VMware instance.

- Workload - Network IO

The amount of the network I/O capacity, in kilobits per second, that is being used by a VMware instance.

**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 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 93 for more information about views.

**7 If you are generating reports for specific systems in your environment, select them from the List of VMware Servers.****8 Select a report generation option. See “Report Generation Options” on page 305 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 310 for more information on configuring a scheduled report.

## Using the VMware Workload Report


The VMware Workload report can be a useful tool for determining whether or not a VMware server is being used to its optimal capacity. For example, you have a VMware server that has four instances running on it. The VMware Workload report returns the following information about the CPU load of the VMware server:



This graph indicates that, on average, the VMware server is using only 20% of its CPU capacity; most of the load is generated by a single instance. The PU on the server can handle up to five 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. You can use the Reports for Virtual Environments (see page 368 for more information) to determine which servers in your environment are good candidates for consolidating on a VMware server.

## VI3 Workload Profile Report

Virtual Infrastructure 3 (VI3; also called VirtualCenter) is a software suite that manages multiple, physical VMware ESX v3 servers. VI3 enables you to manage and monitor virtual servers, as well as allocate resources among virtual machines.

The VI3 Workload Profile report charts the workload of both the server on which VI3 is running and the ESX v3 servers that VI3 is managing. It does this by graphing the key performance counters the [up.time](#) collects from VI3.

### Creating a VI3 Workload Profile Report

To create a VI3 Workload Profile report, do the following:

- 1 In the Reports Tree panel, click **VI3 Workload Profile**.

- 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 27.

- 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 v3 server is ready to run, but cannot run because it cannot access the processor on the ESX v3 server.

- Workload Profile - % Used

The percentage of CPU time that an instance running on an ESX v3 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 93 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 305 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 310 for more information on configuring a scheduled report.

## 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 27.

- 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 93 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 305 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 307 and “**Scheduling Reports**” on page 310 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.



# CHAPTER 18

## Understanding Graphing

.....

This chapter introduces the graphing features of [up.time](#) in the following sections:


|                                     |     |
|-------------------------------------|-----|
| <i>Graphing in up.time</i> .....    | 378 |
| <i>Using the Graph Editor</i> ..... | 380 |

## Graphing in up.time


You can graph performance information to learn about the behaviour 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 385.

**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 Mozilla 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:

- 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:

- 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 385 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  $n^{\text{th}}$  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.

# CHAPTER 19

## Using Graphs

.....

This chapter describes each [up.time](#) graph in the following sections:

|                                                      |            |
|------------------------------------------------------|------------|
| <i>Overview .....</i>                                | <i>386</i> |
| <i>Viewing the Status of a System .....</i>          | <i>387</i> |
| <i>Monitoring CPU Performance .....</i>              | <i>388</i> |
| <i>Multi-CPU Usage .....</i>                         | <i>392</i> |
| <i>Graphing Memory Usage .....</i>                   | <i>395</i> |
| <i>Graphing Processes .....</i>                      | <i>398</i> |
| <i>Graphing TCP Retransmits .....</i>                | <i>400</i> |
| <i>Graphing User Activity .....</i>                  | <i>401</i> |
| <i>Workload Graphs .....</i>                         | <i>402</i> |
| <i>Network Graphs .....</i>                          | <i>407</i> |
| <i>Disk Performance Statistics Graph .....</i>       | <i>409</i> |
| <i>Top 10 Disks Graph .....</i>                      | <i>411</i> |
| <i>File System Capacity Graph .....</i>              | <i>413</i> |
| <i>VXVM Stats Graph .....</i>                        | <i>414</i> |
| <i>Novell NRM Graphs .....</i>                       | <i>416</i> |
| <i>Instance Motion Graphs .....</i>                  | <i>418</i> |
| <i>Displaying Detailed Process Information .....</i> | <i>419</i> |

## 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 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 lists, in a table, the basic hardware and process information for a system. The Quick Snapshot also includes graphs that chart the CPU statistics and free memory of a system over 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 the graphs.

The Quick Snapshot is useful when you want an overview of the status of a system. 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:

| Current CPU/Memory Performance Summary                                                                                                        | Current Top 10 Processes                                                                                                                                                             | Current File System Statistics                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• % User</li> <li>• % Sys</li> <li>• % Wait IO</li> <li>• % Mem Free</li> <li>• % Swap Used</li> </ul> | <ul style="list-style-type: none"> <li>• Process</li> <li>• PID</li> <li>• PPID</li> <li>• MemSize</li> <li>• RSS</li> <li>• % CPU</li> <li>• % Mem</li> <li>• Start Time</li> </ul> | <ul style="list-style-type: none"> <li>• Filesystem</li> <li>• Available</li> <li>• Used</li> <li>• Free</li> <li>• % Full</li> </ul> |

## Viewing a Quick Snapshot

To view a Quick Snapshot, do the following:

- 1 In the **Global Scan** panel, click the name of the system whose information you want to graph.
- 2 In the **Tree** panel, click the **Graphing** tab and then click **Quick Snapshot**.

## 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 391.

### 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 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 27.
- 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 27.

## 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:

- Free
- 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 397.

### Free

This graph charts the amount of free memory available on a system. Free memory is the amount of physical memory available to the operating system, system library files, and applications.

### Cache Hit Rate

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

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

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

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

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

## Paging Statistics

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

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

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

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

## Free Swap

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

Microsoft Windows writes data to the Windows Page File when it needs additional memory. The Windows Page File can range in size from 20 million bytes to over 200 million bytes. The `\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:
  - Free
  - 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 27.

- 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 399.

[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 419.
- “Workload Graphs” on page 402.

### 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

## 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 27.
- 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 27.
- 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 27.

- 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 403.

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



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 27.

- 5 Click one of the following metrics:

- CPU %
- Memory Size
- RSS



You can only graph one metric at a time.

- 6 Select one or more of the available users, groups, or processes from the list.



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

- 7 Click **Add**.
- 8 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 indiv815.2(tds visuat)-4.5e LPARs4( bei)-4

- 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 27.

- 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 27.

- 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

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a Network Graph” on page 408.

### 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.

## 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
- 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 27.
- 5 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 27.

- 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 27.

- 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 27.

- 5 Select one or more file systems from the list.**

If you are generating a graph for a Windows system, you will only be able to generate a graph for the C:\ drive.

- 6 Click Generate Graph.**

## VXVM Stats Graph

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

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

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

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

In the **Info & Rescan** panel, verify that the entry **Has a Logical Volume Manager?** is set to **Yes**. If it is, then ensure that you can connect to the host from the monitoring station. See “Viewing System and Service Information” on page 57 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 27.

- 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.

For more information about Novell NRM systems, see “Novell NRM Systems” on page 440.

## 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 27.
- 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 27.

- 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 CPU resources that the process has been consuming over the past few seconds.

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 27.

- 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:

Choose one ▼

| Acidburn - Process Information - displaying latest sample dated 2005-09-06 18:05:24 |       |       |        |        |             |          |       |           |                  |                     |
|-------------------------------------------------------------------------------------|-------|-------|--------|--------|-------------|----------|-------|-----------|------------------|---------------------|
| Process                                                                             | PID   | PPID  | UID    | GID    | Memory Used | RSS      | CPU % | User Time | User System Time | Start Time          |
| mysqld                                                                              | 12079 | 12070 | uptime | uptime | 20.34 MB    | 14.13 MB | 2     | 7h 33m    | 0s               | 2005-09-01 01:00:00 |
| dtfile                                                                              | 429   | 426   | root   | other  | 8.31 MB     | 536 KB   | 0     | 59s       | 9m 54s           | 2005-03-28 10:55:00 |
| dtfile                                                                              | 455   | 429   | root   | other  | 8.31 MB     | 232 KB   | 0     | 11s       | 0s               | 2005-03-28 10:55:00 |
| dthelpview                                                                          | 430   | 426   | root   | other  | 8.05 MB     | 0 KB     | 0     | 0s        | 0s               | 2005-03-28 10:55:00 |
| dtlogin                                                                             | 314   | 1     | root   | root   | 4.89 MB     | 240 KB   | 0     | 0s        | 2m 29s           | 2005-03-28 10:54:00 |
| dtlogin                                                                             | 349   | 314   | root   | root   | 6.65 MB     | 0 KB     | 0     | 0s        | 0s               | 2005-03-28 10:54:00 |
| dtscreen                                                                            | 604   | 603   | root   | other  | 4.06 MB     | 256 KB   | 0     | 45s       | 0s               | 2005-03-28 11:12:00 |
| dtsession                                                                           | 426   | 413   | root   | other  | 9.13 MB     | 1.88 MB  | 0     | 7m 55s    | 0s               | 2005-03-28 10:54:00 |
| dtwm                                                                                | 428   | 426   | root   | other  | 8.70 MB     | 1.42 MB  | 0     | 19s       | 1s               | 2005-03-28 10:54:00 |
| fsflush                                                                             | 3     | 0     | root   | root   | 0 KB        | 0 KB     | 0     | 5h 39m    | 0s               | 2005-03-28 10:53:00 |
| httpd                                                                               | 20022 | 1     | root   | other  | 5.13 MB     | 1.21 MB  | 0     | 19s       | 41m 29s          | 2005-05-31 12:40:00 |
| httpd                                                                               | 20025 | 20022 | uptime | nobody | 71.56 MB    | 21.60 MB | 0     | 2h 58m    | 34m 50s          | 2005-05-31 12:40:00 |

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



# CHAPTER 20

## Using the Config Panel

.....

This chapter describes the **Config Panel** in the following sections:

|                                    |            |
|------------------------------------|------------|
| <i>Archive Policy .....</i>        | <i>424</i> |
| <i>Mail Servers .....</i>          | <i>427</i> |
| <i>License Information .....</i>   | <i>428</i> |
| <i>Monitoring Periods .....</i>    | <i>429</i> |
| <i>Problem Reporting .....</i>     | <i>431</i> |
| <i>up.time Configuration .....</i> | <i>432</i> |

## Archive Policy

You can selectively archive data in the DataStore by configuring an archive policy in the **Config** panel. Each month, **up.time** checks the data in the DataStore to determine whether or not the data is older than the archive period that you specify. If **up.time** discovers data that is older than the archive period, it writes the data to a set of XML files. The XML files are then compressed and moved into a directory named **archives** under the root installation directory.

For example, if you installed **up.time** in the default location, the path to the archived data will be:

- On Solaris or Linux:  
`/usr/local/uptime4/archives`
- On Windows:  
`C:\Program Files\uptime software\uptime4\archives`

The names of the XML files have the following format:

`<table_name>_<date>.xml.gz`

Where `<table_name>` is the name of DataStore table that was archived, and `<date>` is the date on which the archive was created. The following is an example of an archive containing data from the `performance_cpu` table:

`performance_cpu_2006-09-13.xml.gz`

You can then back up the archived files to a tape, a CD, or any other form of offline storage. When needed, you can bring the data back online by importing the archived files into the DataStore.



The archiving procedure that is described in this section only works with MySQL databases.

## 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 **Specify the number of months worth of data that you want to archive for each of the following categories:**
  - 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 **Click the Enable Archiving option.**
- 5 **Click Set Archive Policy.**
- 6 **Optionally, you can click the Archive Now button to immediately create archives of the data in your DataStore.**

All of the data in the categories listed in step 3 is archived.

## Importing Archived Data into the DataStore

You can import archived data into the DataStore using a command line utility called `restorearchive`, which enables you to import:

- A single archive filed into the DataStore.
- A set of files that include a specific date in their file names.

To import archived data into the DataStore, do the following:

- 1 **At the command line, navigate to the following directory:**
  - On Linux and Solaris: `/usr/local/uptime4/scripts/`.
  - On Windows: `C:\Program Files\uptime software\uptime4\scripts\`.

### 2 Run the command `restorearchive` with one or more of the following options:

- `-f <filename>`  
Imports a single file. 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/uptime4/
archives/ -c /usr/local/uptime4
```



## Mail Servers

You can manually configure the email server that is used to send alerts and reports to [up.time](#) users. The changes you make in the [up.time](#) Web interface will be written to the `uptime.conf` file. For more information about `uptime.conf`, see “Enabling the Windows Messaging Service” on page 448.

To configure mail servers, 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 Mail Server Configuration**.
- 4 Type the name of the mail server in the **SMTP Server** field.

The name of the server should follow this convention:  
`smtp.<domain_name>`

- 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.

The originating email address must be a valid email address.

- 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.
- 9 In the **SMTP Password** field, enter the password that is used to authenticate connections.
- 10 Click **Save**.
- 11 To test the mail server configuration, click the **Test Configuration** button in the sub panel.

The 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.

## 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](mailto: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.**

## Monitoring Periods

Monitoring periods are the times over which a service monitor will be actively monitoring a host. The monitoring periods also apply to 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 **Config**.
- 2 In the **Tree** panel, click **Monitoring Periods**.  
The existing monitoring periods appear in the subpanel.
- 3 Click the **Add Monitoring Period** tab in the subpanel.  
The **Add Monitoring Periods** window appears.
- 4 Type a name in the **Monitoring Period Name** field.
- 5 In the **Time Range** section, select the start and end times for the monitoring period from the dropdown lists.

For example, to specify a monitoring period that runs from 1:30 a.m. to 3:00 am, select 01 and 30 from the **Start** dropdown lists and then select 03 and 00 from the **End** dropdown lists.

### 6 Select one of the following options:

- Weekly

Select one or more of the days of the week for which this monitoring period applies.

- Monthly

Do one of the following:

- Select the **Day** option, and then select the day (from 1 to 31) for which this monitoring period applies.

For example, if you select 3 from the dropdown list, the monitoring period will be the third day of every month.

- Select the second option, and then:

- Select first, second, third, fourth, or last from the first dropdown list.
- Select a day of the week for which the period will apply from the second dropdown list.

For example, if you select second and Tuesday from the dropdown lists, the monitoring period will be in effect on the second Tuesday of every month.

- Yearly

Select a month from the first dropdown list, and then select a day from the second dropdown list.

### 7 Click **Save**.

## 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 folder on the monitoring station and has a file name with the following format:

prYYYYMMDD-HHMMSS.zip

Where:

- YYYYMMDD is the date on which the report was generated, for example: 20061212.
- HHMMSS is the time at which the report was generated, for example: 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 like 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 and then send the archive to uptime software Client Care.

## up.time Configuration

The file `uptime.conf` is a configuration file that contains parameters controlling certain aspects of the behavior of [up.time](#). By editing `uptime.conf` you can, for example, change the port on which your DataStore is listening or turn the audit log on or off.

`uptime.conf` is found in the root `up.time` installation directory. It is a text file that you can modify in any text editor. You can also edit the settings in `uptime.conf` from within the [up.time](#) interface.

To edit `uptime.conf` 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 Find the entry that you want to edit and make the changes.**  
See “The `uptime.conf` File” on page 455 for more information on the contents of the `uptime.conf` file.
- 4 Click Update to save your changes.**

# APPENDIX A

## Additional Information

.....

This appendix contains the following sections:

|                                                            |            |
|------------------------------------------------------------|------------|
| <i>Adding Multiple Systems to up.time.....</i>             | <i>434</i> |
| <i>Novell NRM Systems .....</i>                            | <i>440</i> |
| <i>Enabling SSH on a Hardware Management Console .....</i> | <i>447</i> |
| <i>Enabling the Windows Messaging Service .....</i>        | <i>448</i> |
| <i>Changing the Agent Password on Windows.....</i>         | <i>449</i> |
| <i>Custom Alert Format Variables.....</i>                  | <i>450</i> |
| <i>Archiving and Restoring the DataStore .....</i>         | <i>452</i> |

## Adding Multiple Systems to up.time

It can be time consuming to add large numbers of systems to [up.time](#) using the Web interface. You can, however, add multiple systems to [up.time](#) using the `addsystem` command line tool and a text file.

A text file, called a *hosts file*, contains entries which mirror the fields in the **Add System** window of the [up.time](#) Web interface. These fields contain information about the systems that you want to add.

You can find examples of entries in a hosts file in the section “Examples of Hosts File Entries” on page 438.

In the hosts file:

- The information for each host consists of a name-value pair. Each name-value pair is on a separate line, and is separated by a colon. For example, `Group: Solaris Servers`.
- The information for each host is separated by two percentage signs (%%) on a new line.

## Creating a Hosts File

There are a number of ways in which you can create a hosts file. The simplest way is to use a text editor to type the entries in a file. If you have a large number of systems to add, you can copy and paste an entry, and modify the fields as needed.

If you keep a list of all the systems in your environment in a spreadsheet, you can save the list as a text file or a comma separated values (.csv) file. Then, you can write a script that can manipulate the text or .csv file into the proper format.

## Fields in the Hosts File

The following table explains the fields that you can include in the hosts file. The fields that are needed to add a system will vary depending on the type of system that you want to add. For example, to add an agent system you only need to include the Host Name, Type, and Port fields. See “Working with Systems” on page 76 for more information.



| Field         | Description                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Host Name     | The name or the IP address of the system that you want to add to <a href="#">up.time</a> .                                                                                                                                                                                                                                                                                               |
| Display Name  | The name for the system that will appear in the <a href="#">up.time</a> Web interface.                                                                                                                                                                                                                                                                                                   |
| Description   | A short description of the system. This field is optional.                                                                                                                                                                                                                                                                                                                               |
| Type          | The type of system, which can be one of the following: <ul style="list-style-type: none"> <li>•Agent</li> <li>•Node</li> <li>•Novell NRM</li> <li>•Net-SNMP v1/2</li> <li>•Net-SNMP v3</li> <li>•Virtual Node</li> <li>•pSeries LPAR Server (HMC)</li> </ul>                                                                                                                             |
| Service Group | The name of the <a href="#">up.time</a> service group – which enables you to simultaneously apply common service checks to hosts that you are monitoring – to which you want to add the system.<br><br>This field is optional.                                                                                                                                                           |
| Port          | The number of the port on which you will be connecting to the system. Leave this field blank to use the default port for the type of system that you are adding.                                                                                                                                                                                                                         |
| Community     | If you are adding a Net-SNMP system to <a href="#">up.time</a> , specify the read community (which acts like a user ID or password) that gives you access to the system. Valid options are: <ul style="list-style-type: none"> <li>•<b>public</b>, which enables you to retrieve read-only information.</li> <li>•<b>private</b>, which enables you to access all information</li> </ul> |

| Field                 | Description                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HMC Hostname          | The name or the IP address of the Hardware Management Console (HMC) that is being used to manage one or more pSeries LPAR servers in your environment.                                                                                                                                                                                                                      |
| Managed Server        | The unique identifier of a pSeries LPAR server that is managed by an HMC.                                                                                                                                                                                                                                                                                                   |
| Username              | If you are adding a Net-SNMP or Novell NRM system to <a href="#">up.time</a> , specify the user name required to access the system.                                                                                                                                                                                                                                         |
| Password              | If you are adding a Net-SNMP or Novell NRM system to <a href="#">up.time</a> , specify the password required to access the system.                                                                                                                                                                                                                                          |
| Group                 | <p>The name of the entity group – a set of systems that have been combined in a meaningful way – to which you want to add this system.</p> <p>This field is optional.</p>                                                                                                                                                                                                   |
| SSL                   | <p>For agent systems, use this field to determine whether or not <a href="#">up.time</a> will securely communicate with an agent installed on the system using SSL. Valid options are <code>true</code> and <code>false</code>.</p> <p>This field is optional.</p>                                                                                                          |
| Authentication Method | <p>For Net-SNMP systems, use this field to determine how encrypted information travelling between the Net-SNMP instance and <a href="#">up.time</a> will be authenticated. Valid options are:</p> <ul style="list-style-type: none"><li>•MD5, a widely-used method for creating digital signatures.</li><li>•SHA, a secure method of creating digital signatures.</li></ul> |

| Field            | Description                                                                                                                                                                                                                                                                                                                                        |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Privacy Password | For Net-SNMP systems, the password that will be used to encrypt information travelling between the Net-SNMP instance and <a href="#">up.time</a> .                                                                                                                                                                                                 |
| Privacy Type     | For Net-SNMP systems, how information travelling between <a href="#">up.time</a> and the Net-SNMP instance is encrypted. Valid options are: <ul style="list-style-type: none"><li>•DES, an older method used to encrypt information.</li><li>•AES, the successor to DES, which is used with a variety of software including SSL servers.</li></ul> |
| Pingable         | For nodes, use this field to specify whether or not <a href="#">up.time</a> can contact the node using the ping utility. Valid options are true and false.                                                                                                                                                                                         |

## Adding Multiple Systems to up.time

To add multiple systems to [up.time](#), do the following:

- 1 Copy the hosts file to the directory in which you installed the [up.time](#) monitoring station.**
- 2 At the command line, navigate to the scripts folder.**

For example, if you installed the monitoring station in the default location on a Windows system, navigate to the following folder:

```
C:\Program Files\uptime software\uptime4\scripts\
```

- 3 Enter the following command:**

```
addsystem <path_and_filename>
```

Where `<path_and_filename>` is the name of the text file that contains the list of systems that you want to add to [up.time](#) along with its full path.

The systems listed in the file are added to **up.time**, unless:

- **up.time** cannot connect to the system.
- The system does not exist in your environment.
- The system has already been added to **up.time**.

## Examples of Hosts File Entries

The following table contains sample host file entries for each type of system that you can add to **up.time**:

| Host Type      | Sample Hosts File Entry                                                                                                                                                                  |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Agent          | Host Name: prod-mainSystem<br>Display Name: prod1<br>Description: Main production server<br>Type: Agent<br>Service Group: Production Systems<br>Port:9998<br>Group: Windows 2003 Servers |
| Node           | Host Name: www.myDomain.ca<br>Display Name: Your Domain<br>Description: A Web site<br>Type: Node<br>Group: Web Sites                                                                     |
| Novell NRM     | Host Name: novell01<br>Display Name: dn3<br>Type: Novell NRM<br>SSL: true<br>Port: 546<br>Group: Unix Boxes<br>Group: Novell System                                                      |
| Net-SNMP v1/v2 | Host Name: gateway.mydomain.com<br>Display Name: gatewaySNMP<br>Description: snmp v2<br>Type: Net-SNMP v1/2<br>Read Community: myCo-pub                                                  |

| Host Type    | Sample Hosts File Entry                                                                                                                                                                                                               |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Net-SNMP v3  | Host Name: SNMP-1<br>Display Name: SNMP-1<br>Description: Net-SNMP system<br>Type: Net-SNMP v3<br>Read Community: public<br>Username: myUsername<br>Password: myPassword<br>Privacy Password: myOtherPassword<br>Group: Linux Systems |
| Virtual Node | Host Name: router-Toronto<br>Display Name: Toronto Router<br>Description: Router for Toronto branch<br>Type: Virtual Node<br>Pingable: True<br>Group: Routers                                                                         |
| pSeries LPAR | Host Name: 10.1.2.42<br>Display Name: HMC Managed Server<br>HMC Hostname: 10.1.1.255<br>Type: pSeries LPAR Server (HMC)<br>Managed Server: Server-7610-31C-SN01B030K<br>Username: hscroot<br>Password: hscroot                        |

## Novell NRM Systems

**up.time** collects performance metrics and availability information from version 6.5 of the Novell Remote Manager (NRM) using HTTP or HTTPS. **up.time** extracts performance information from the NRM by reading and parsing XML files.

This section contains the following topics:

- Adding a Novell NRM System to up.time
- NRM Statistics Captured by up.time
- Graphing Statistical Information

### Adding a Novell NRM System to up.time

To add a Novell NRM version 6.5 system to **up.time**, do the following:

**1 On the **up.time** tool bar, click My Enterprise and then click the Add System/Network Device tab.**

**2 Complete the Display name in up.time and Description fields.**

See “Adding Systems or Network Devices” on page 78 for more information.

**3 Select Novell NRM from the Type of System/Device dropdown list.**

**4 Complete the following fields:**

- Host name

The actual name of the machine that **up.time** will be monitoring, or the IP address of the machine.

- Port

The port on which the NRM is listening. The default is 8008 for a port that is not using SSL. The default for a port that is using SSL is 8009.

- Username

The NRM administrator account name. This field is mandatory.

- Password

The NRM administrator password. This field is mandatory.



The password is encrypted and stored in the [up.time DataStore](#).

- 5 If you want to associate this system with a group, select the name of the group from the Group dropdown list.
- 6 Click **Save**.

## NRM Statistics Captured by up.time

[up.time](#) captures the following Novell NRM system (version 6.5) statistics:

- Work To Do Response Time
- Allocated Service Processes
- Available Server Processes
- Abended Thread Count
- CPU Utilization
- Connection Usage
- Available Memory
- DS Thread Usage
- Packet Receive Buffers
- Available Event Control Blocks (ECBs)
- LAN Traffic
- Available Disk Space
- Disk Throughput

Each statistic returns one of the following statuses:

- Good

The statistic is well within the threshold suspect value.

- Suspect  
The statistic is between the threshold good and critical values.
- Bad  
The statistic is greater than the threshold critical value.

## Work To Do Response Time

This statistic enables you to view how processes share the CPU. The response time is the amount of time that a Work To Do process requires to run.

If this statistic returns a value of Suspect, you can check the running threads to determine why there is a delay in the Work To Do threads. If the value is Bad, thread is probably running more than it should or it is hung. You should identify the parent NetWare Loadable Module and then unload and reload it if possible.

## Allocated Service Processes

This statistic enables you to view, as a graph, how the service processes are allocated on your server.

If the service processes are approaching the maximum, increase the value of the Maximum Server Processes Set parameter. If you have only a few available server processes, increase the Minimum Server Processes Set parameter.

If the status is Bad, examine your server by doing the following:

- 1 In Novell NRM, click Profiling / Debugging.**
- 2 Check the information for server process functions.**
- 3 Change the Maximum Server Processes and the Minimum Server Process Set parameters.**

## Available Server Processes

This statistic enables you to view the number of available processes on your server as a graph. The graph charts the processes that are available every five seconds over a 50 second period.



If the status is Suspect or Bad, you should increase the Set parameters for Maximum Server Processes and the Minimum Server Processes settings. If the number of available server processes has not reached the maximum and is not increasing, you should add memory to your server.

## Abended Thread Count

This statistic enables you to view the threads that have ended abnormally (abended) and are suspended. This statistic returns the following statuses:

If the status is Suspect or a Bad, your server has abended and has recovered automatically by suspending the offending thread while leaving the rest of the server processes running. As a result, some of the server's functions were compromised. You must determine which module, driver, or hardware the abended threads belong to, and then take the appropriate action.

## CPU Utilization

This statistic enables you view, as a graph, how busy any given CPU is. [up.time](#) tracks usage on a per CPU basis, collecting data every 30 seconds. The graph displays a 10 second history.

If the status is Suspect or Bad, determine which thread or module is causing the most CPU cycles and take appropriate action, including:

- Unloading and reloading the module.
- Reporting problems to the vendor of the module.
- Loading an updated module.

To determine which thread or module is using the most CPU cycles, do the following:

- 1 In Novell NRM, click **Profile / Debug**.
- 2 Do one of the following:
  - View the Execution Profile Data by Thread data.
  - Click **Profile CPU Execution by NLM**.

## Connection Usage

[up.time](#) monitors connections on a per-server basis. NRM displays only the following metrics:

- The number of connections that are being used.
- The peak number of connections used on this server.

## Available Memory

This statistic enables you to view the amount of memory that is not allocated to any service. Most, if not all, of this memory is used by the file system cache. When available memory gets too low, modules might not be able to load or file system access might become sluggish.

## DS Thread Usage

This statistic enables you view the number of server threads that Novell eDirectory uses. The server thread limit ensures that threads are available for other functions as needed – for example, when large number of users log in at the same time.

eDirectory uses multiple server threads. However, its thread requirements should not cause poor performance because eDirectory cannot use more than its allocated maximum number of threads.

If this statistic returns a Good status, eDirectory is using less than 25% of the available server threads. If it returns a Suspect status, eDirectory is using between 25% and 50% of the available server threads. If the status is Bad, eDirectory is using more than 50% of the available server threads.

## Packet Receive Buffers

This statistic enables you to view the status of Packet Receive Buffers for the server. Packet Receive Buffers transmit and receive packets. You can set the maximum or minimum number of buffers to allocate using the Maximum Packet Receive Buffers or Minimum Packet Receive Buffers SET parameters. The minimum number of buffers is the number of packets that are allocated at when the system is initialized.

If the number of Packet Receive Buffers is increasing, the system will be sluggish. If the number of Packet Receive Buffers reaches the maximum, and no Event Control Blocks (ECBs) are available, the server will become very sluggish and will not recover.

## Available Event Control Blocks (ECBs)

This statistic enables you to view the status of available Event Control Blocks (ECBs). Available ECBs are Packet Receive Buffers that have been created but which are not currently being used.

If the available ECB count is zero, the server will become sluggish until enough ECBs are created to fill the demand. The server will recover as long as the number of Packet Receive Buffers does not increase to the maximum that can be allocated.

## LAN Traffic

This statistic shows whether or not your server can transmit and receive packets. If this statistic returns a Good status, the server is able to accept or transmit packets through the network board. If the status is Bad, the network board is not transmitting or receiving packets.

All servers should be able to transmit or receive packets. If your server is not transmitting, your LAN is not functioning properly. Check the drivers and protocol bindings for the network board on the server. If the drivers and protocol bindings are functioning properly, then the network board is probably faulty. If the network board is functioning, you should perform a diagnostic on your LAN.

## Available Disk Space

This statistic enables you to view the status of the available disk space on all mounted volumes on a server. This statistic returns the following statuses:

## Disk Throughput

This statistic enables you to view the status of amount of the data that is being read from and written to the storage media on this server.

If this statistic returns a Good status, then the storage system is experiencing reads or writes, and there are no pending disk I/Os. If the status is Suspect, the storage system has disk I/Os pending, no reads or writes have occurred, and less than four samples have been taken. If the status is Bad, the storage system has disk I/Os pending, no reads or writes have occurred, and four or more samples have been taken.

## Graphing Statistical Information

To graph statistical information, do the following

- 1** On the [up.time](#) tool bar, click **Global Scan**.
- 2** Click the name of a Novell NRM system in the list of systems.
- 3** Select one of the statistics listed above from the **Tree** panel.
- 4** Select the date range and complete required and optional fields.  
For more information, see “Understanding Dates and Times” on page 27.
- 5** Click **Generate Graph**.

A graph containing statistical information appears in a separate window.

## Enabling SSH on a Hardware Management Console

The Hardware Management Console (HMC) is an interface for managing and configuring pSeries servers that are hosting multiple logical partitions (LPARs). When an HMC is attached to one or more pSeries servers with LPARs, the servers are *managed servers*. In this configuration, the HMC manages all I/O requests from the LPARs. In order to monitor the managed servers and their LPARs, up.time must execute commands on the HMC.

Before up.time can execute commands on the HMC, you must enable SSH (Secure Shell) on the Hardware Management console. Once you have done that, you can remotely log into the HMC using the HMC's user name (default: hmcroot) and password.

To enable SSH on a Hardware Management Console, do the following:

- 1 **Log into the graphical interface on the HMC.**
- 2 **In the Navigation area, select HMC Management and then click HMC Configuration.**
- 3 **Click Enable/Disable Remote Command Execution.**
- 4 **In the window that appears, click Enable SSH.**

This creates the file `authorized_keys2` on the HMC.

- 5 **Create a user on the HMC with one of the following roles:**
  - Super Administrator
  - Service Representative

## 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**.

## Changing the Agent Password on Windows

By default, you can access an up.time agent installed on a Windows server without a password. However, if you want communication between the agent and the up.time monitoring station to be more secure, you should add a password to the agent.

To add or change the agent password on Windows, do the following:

- 1 **Select Start > Run.**
- 2 **Type regedit and then click OK.**
- 3 **In the Windows Registry Editor, find the following entry:**  
 HKEY\_LOCAL\_MACHINE\SOFTWARE\uptime software  
 \up.time NT agent
- 4 **Double click the string CmdsPassword to open the Edit String dialog box.**



- 5 **In the Edit String dialog box, type the new password in the Value data field and then click OK.**

# 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. For more information on creating a custom alert profile, see “Working with Custom Alert Formats” on page 130.

The table below explains the variables used in a custom alert format:

| Variable         | Definition                                                                                                                                                                                                                                                                                  |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$DISPLAYNAME\$  | <p>The name of the entity as it appears in the <a href="#">up.time</a> Web interface.</p> <p>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.</p>                    |
| \$DATETIME\$     | <p>The date and time at which the alert was generated. This appears in the subject line of the message.</p>                                                                                                                                                                                 |
| \$SERVICENAME\$  | <p>The name of the service, along with the name of the host for which the alert was generated.</p> <p>For example, if the alert was generated by the ping check for the server MailHub, then PING-MailHub appears in the alert.</p> <p>This appears in the subject line of the message.</p> |
| \$SERVICESTATE\$ | <p>One of the following:</p> <ul style="list-style-type: none"><li>•OK</li><li>•WARN</li><li>•CRIT</li><li>•MAINT</li><li>•UNKNOWN</li></ul> <p>This appears in the subject line of the message.</p>                                                                                        |
| \$DATE\$         | <p>The date on which the alert was generated.</p>                                                                                                                                                                                                                                           |



| Variable      | Definition                                                                                                                                                                 |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$TIME\$      | The time at which the alert was generated.                                                                                                                                 |
| \$HOSTNAME\$  | The name of the host (as saved in <a href="#">up.time</a> ) for which this alert was generated.                                                                            |
| \$HOSTSTATE\$ | The status of the host, which can be one of the following: <ul style="list-style-type: none"><li>•OK</li><li>•WARN</li><li>•CRIT</li><li>•MAINT</li><li>•UNKNOWN</li></ul> |
| \$TYPE\$      | The type of notification, which can be one of the following: <ul style="list-style-type: none"><li>•Problem</li><li>•Recovery</li></ul>                                    |
| \$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                                        |

## Archiving and Restoring the DataStore

[up.time](#) includes two command line utilities for archiving and restoring the contents of your DataStore. The utilities are:

- `fulldatabasedump`  
Creates a compressed XML file of the contents of your DataStore.
- `fulldatabaseimport`  
Imports the archived data back into your DataStore.

Both utilities work with all of the databases that [up.time](#) supports.

### Archiving the DataStore

To archive your DataStore, do the following:

- 1 Shut down the [up.time](#) Data Collector.**
- 2 Navigate to the scripts folder under the directory where [up.time](#) is installed.**
- 3 Run the following command:**

```
fulldatabasedump
```

Depending on the size of your DataStore, this process can take anywhere from several minutes to several hours.

The utility creates the file `uptimedump_YYYY-MM-DD.xml.gz` – for example `uptimedump_2007-01-02.xml.gz`. This file is saved in [up.time](#)'s root installation directory.

### Restoring the DataStore

To restore your DataStore, do the following:

- 1 Ensure that the DataStore service is running.**

- 2 Use the `resetdb` utility with the `really` option to delete and then recreate the database structure that is used by up.time 4 by running one of the following commands:

- On Windows: `C:\Program Files\uptime software\uptime4\resetdb really`
- On Linux or Solaris: `/usr/local/uptime4/resetdb really`

- 3 Run the following command:

```
fulldatabaseimport path/<filetoimport>.xml.gz
```

Where `path/<filetoimport>.xml.gz` is path to and file name of the archived contents of your DataStore. For example, to import archive that is located in up.time's root installation directory enter the following:

```
fulldatabaseimport uptimedump_2007-01-02.xml.gz
```



# APPENDIX B

## The uptime.conf File

.....

This appendix explains the entries in the file `uptime.conf` in the following sections:

|                                             |            |
|---------------------------------------------|------------|
| <i>Overview .....</i>                       | <i>456</i> |
| <i>Database Settings .....</i>              | <i>457</i> |
| <i>Mail Server Settings .....</i>           | <i>459</i> |
| <i>SNMP Trap Settings .....</i>             | <i>460</i> |
| <i>Basic System Settings .....</i>          | <i>462</i> |
| <i>Global Scan Threshold Settings .....</i> | <i>465</i> |
| <i>Remote Reporting Settings .....</i>      | <i>466</i> |
| <i>RSS Feed Settings .....</i>              | <i>467</i> |
| <i>Splunk Settings .....</i>                | <i>468</i> |

## Overview

The file `uptime.conf` is a configuration file that contains parameters controlling certain aspects of the behavior of [up.time](#). By editing `uptime.conf` you can, for example, change the port on which your DataStore is listening or turn the audit log on or off.

`uptime.conf` is found in the root [up.time](#) installation directory. It is a text file that you can modify in any text editor. You can also edit the settings in `uptime.conf` from within the [up.time](#) interface. See “[up.time Configuration](#)” on page 432 for more information.

The parameters in this file determine the following:

- Database Settings
- Mail Server Settings
- SNMP Trap Settings
- Basic System Settings
- Global Scan Threshold Settings
- Remote Reporting Settings
- RSS Feed Settings



Not all of the setting listed in this section are found in the file `uptime.conf`.

## Database Settings

The database settings determine how [up.time](#) communicates with the DataStore. The following are the database settings in the `uptime.conf` file:

- `dbDriver=`

The database driver that is used to connect the monitoring station to the DataStore. By default, [up.time](#) uses a JDBC (Java Database Connectivity) driver. You can also use an ODBC driver, which enables you to connect to the DataStore with tools like MySQL Query Browser, Microsoft Excel and Crystal Reports.

The supported drivers are:

- `com.mysql.jdbc.Driver` (for MySQL)
- `net.sourceforge.jtds.jdbc.Driver` (for SQL Server)
- `oracle.jdbc.OracleDriver` (for Oracle)

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”.

- `dbType=`

The type of database that is being used to store data from [up.time](#). The default is `mysql`. You can also specify `mssql` and `oracle`.

- `dbHostname=`

The name of the system on which the database is running. The default is `localhost`.

- `dbPort=`

The port on which the database is listening. The default is `3307`.

- `dbName=`

The name of the database. The default is `uptime_v4`.

- `dbUsername=`

The name of the default database user, which is `uptime`. See “Database Users” on page 485 for more information on the database users that are installed with [up.time](#).

## *Database Settings*

- `dbPassword=`

The password for the default database user, which is `uptime`. See “Database Users” on page 485 for more information on the database user passwords.

- `connectionPoolMaximum=`

The maximum number of connections that are allowed to the DataStore. Setting this option to a lower number will help increase the performance of [up.time](#).



## Mail Server Settings

up.time uses a mail server to send alerts and reports to up.time users. The following are the mail server settings in the uptime.conf file:

- **smtpServer=**  
The host name of the system on which up.time is running.
- **smtpPort=**  
The port on which the mail server is listening.
- **smtpSender=**  
The name and email address from which reports and alerts will be sent to users. Set the email address to your domain, for example:  
`uptime4@mail.uptimesoftware.com`
- **smtpHeloString=**  
An string that identifies the domain from which a message is being sent. Change the domain name to the domain from which you will be sending email messages to users of up.time. For example, `uptimesoftware.com`.
- **smtpUser=**  
The user name that is used to authenticate connections. This setting is optional.
- **smtpPassword=**  
The password that is used to authenticate connections. This setting is optional.

## SNMP Trap Settings

When **up.time** detects a problem with a host in your environment, it can send an SNMP trap to a particular host. An SNMP trap is a 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 in order to receive and process the trap.

The type of information included in the trap is determined by the variables for the trap setting. These are appended to the trap OID in the following manner:

`OID.variable_number`

For example, if the trap OID is `1.3.6.1.4.1.20000` and it contains information about the service names, the OID will look like this:

`1.3.6.1.4.1.20000.7`

The following are the SNMP trap settings in the `uptime.conf` file:

- `snmpUptimeTrapOID=1.3.6.1.4.1.24216`  
The SNMP object identifier (OID) that identifies the SNMP trap on the host.
- `snmpUptimeTrapInfo=1`  
Indicates that the trap host will accept SNMP traps.
- `snmpUptimeTrapLongMessage=1`  
Indicates that the trap host accepts detailed trap information.
- `snmpUptimeTrapProblemType=2`  
Indicates that the type of problem is included in the trap sent to the host.
- `snmpUptimeTrapDatetime=3`  
Indicates that the date and time of the trap are included in the information that is sent to the host.
- `snmpUptimeTrapHostname=4`  
Indicates that the name of the host that is being monitored is included in the trap.

- `snmpUptimeTrapHoststate=6`  
Indicates that the state of the host that is being monitored is included in the trap.
- `snmpUptimeTrapServiceName=7`  
Indicates that the name of the service being monitored on the host is included in the trap.
- `snmpUptimeTrapServiceStatus=8`  
Indicates that the status of the service being monitored on the host is included in the trap.
- `snmpUptimeTrapOutput=9`  
Indicates that some output – for example, a status and a description of the problem – is included in the trap.

## Basic System Settings

The basic system settings determine the values for such system parameters as the path in which to look for executables, the performance check interval, and log levels.

The following are the basic system settings in the `uptime.conf` file:

- **externalErdcScriptPath=**

The path in which **up.time** searches for executables – specifically, `icmp.exe` and `lmhostid.exe`. This path is `<install_dir>/bin/`, where `<install_dir>` is the directory in which **up.time** is installed.

For example, if you installed **up.time** in the default location, then on Windows this path will be:

`C:/Program Files/uptime software/uptime4/support/bin/`

On Solaris or Linux, the path will be:

`/usr/local/uptime4/bin`

- **trustedIpAddresses=**

A list of IP addresses from which the core will accept connections. If there are multiple IP addresses, separate them with commas.

- **performanceCheckInterval=**

The amount of time, in seconds, between performance checks. The default is 300. This parameter only applies to hosts that have been recently created.

- **serviceThreads=**

The number of threads that will run service and performance monitors. The default is 100.

- **reportCacheExpiryDays=**

The number of days over which **up.time** keeps reports cached on the system. The default is 30.

- **auditEnabled=**

Enables the audit log, which writes a log of user actions to the file `audit.log`, located in the `logs` subdirectory. Specify yes or no. The default is yes.

- `loggingLevel=`

The type of information that **up.time** writes to a log file. The logging level can be any of the following:

- ALL
- DEBUG
- INFO
- WARN
- ERROR
- FATAL
- OFF

The default is `DEBUG`.

- `httpContext=`

The context of the **up.time** Web server, in the format `http://<hostname>:<port>`. Where:

- `<hostname>` is the host name of the server on which **up.time** is running. For example, `localhost`.
- `<port>` is the port on which the **up.time** Web server is listening for requests. For example, `9999`. You can optionally omit the port number.

- `publishedReportRoot=`

The directory in which reports will be published, usually `<install_dir>/GUI/`, where `<install_dir>` is the directory in which **up.time** is installed. The report files will be saved in a subdirectory published. See “Saving Reports to the File System” on page 307 for more information.

- `versionCheck`

This option checks the current version of **up.time** when the Data Collector is started, once every three days, and whenever the license is updated from within the **up.time** Web interface. This information is displayed in the **My Portal** panel.

By default, this option is set to `true`. To disable it, set `versionCheck` to `false`.

- `default3DActivexGraphs=`

This option determines whether or not ActiveX graphs are displayed in 2D or 3D. The default is `false`. Note that this setting will only affect users who connect to the monitoring station using Internet Explorer.

- `acknowledgedSeparate=`

If set to yes, 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.

## Global Scan Threshold Settings

The Global Scan threshold settings are the default warning and critical alert values that are applied to any hosts that are added to [up.time](#).

The following are the Global Scan threshold settings in the `uptime.conf` file:

- `globalscan.cpu.warn=70`  
A warning alert is issued when CPU usage is at 70% or greater.
- `globalscan.cpu.crit=90`  
A critical warning is issued when CPU usage is at 90% or greater.
- `globalscan.diskbusy.warn=70`  
A warning alert is issued when a disk on the host is busy for 70% or more of the time.
- `globalscan.diskbusy.crit=90`  
A critical alert is issued when a disk on the host is busy for 90% or more of the time.
- `globalscan.diskfull.warn=70`  
A warning alert is issued when 70% or more of the disk space on the host is used.
- `globalscan.diskfull.crit=90`  
A critical alert is issued when 90% or more of the disk space on the host is used.
- `globalscan.swap.warn=70`  
A warning alert is issued when 70% or more of the swap space on a disk is in use.
- `globalscan.swap.crit=90`  
A critical alert is issued when 90% or more of the swap space on a disk is in use.

## Remote Reporting Settings

If you are using a reporting instance (instance of [up.time](#) that only generates and serves reports; see the Knowledge Base article “Setting up a reporting instance” for details), the remote reporting settings enable you to specify the location of the reporting instance and the port on which it is listening.

The following are the reporting instance settings in the `uptime.conf` file:

- `remoteReportServer=localhost`

The name of the server on which the reporting instance is running. The default is `localhost`, which indicated that there is no reporting instance.

- `remoteReportServerPort=9996`

The port on which the reporting instance is listening.



## 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). You simply click the title of the article to open it in your Web browser.

When the following message appears in the My Portal panel:

✗ up.time can't get the latest knowledgebase articles from <http://support.uptimesoftware.com>. [Click here to read what's new.](#)

Then your monitoring station probably accesses the Internet through a proxy server. You can change the RSS feed settings to point to the proxy server rather than directly to the uptime software Web site.

The following are the RSS feed settings in the `uptime.conf` file:

- `rssFeedUrl`  
The URL of the RSS feed. The default is `http://support.uptimesoftware.com/rss/kb.xml`.
- `httpProxyHost`  
The host name of the proxy server that the monitoring station uses to access the Internet.
- `httpProxyPort`  
The port on which of the proxy server that the monitoring station uses to access the Internet is listening.

## 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 data center.

If the settings listed below are in the `uptime.conf`, then 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 into your Splunk search page.

The following settings enable the automatic log in to your Splunk search page:

- `splunk.url=`

The URL of the server on which your Splunk search pages is hosted. For example, `http://webportal:8000`.

- `splunk.username=`

The user name required to log into your Splunk search page.

- `splunk.password=`

The password required to log into your Splunk search page.

- `splunk.soapurl=`

The URL that points to the SOAP management port which Splunk uses to communicate with the `splunkd` daemon. For example, `https://webportal:8089`.

You must set up the port on which the Splunk server listens for requests on the server itself. See the Splunk Admin Manual for more information.

# APPENDIX C

## SQL Server Ports and Locks

.....

This appendix explains SQL Server ports and locks in the following sections:

|                               |            |
|-------------------------------|------------|
| <i>SQL Server Ports .....</i> | <i>470</i> |
| <i>SQL Server Locks .....</i> | <i>475</i> |

## SQL Server Ports

This section contains information about:

- Verifying the Configuration of a Port
- Configuring a Static Port
- Configuring a Dynamic Port
- SQL Server Scripts

### Verifying the Configuration of a Port

To verify the configuration of a port, you must first find the TCP/IP port of your SQL Server instance.

#### Finding the TCP/IP Port

To find the TCP/IP port of a SQL Server instance, do the following:

- 1 Start SQL Query Analyzer, and then connect to the instance of SQL Server.**
- 2 Run the following Transact-SQL statement in SQL Query Analyzer:**

```
Use master
Go
Xp_readerrorlog
```

- 3 In the Results pane, locate the following text:**

SQL server listening on X.X.X.X: Y

Where x.x.x.x is the IP addresss6 c7.X.X:9.61459.1.2048 Tm0.0048 Tc0.003erch3(L r)-15(A)

## Verifying Port Configuration

To verify the port configuration of a SQL Server instance, do the following:

- 1 **Start the Windows Registry Editor.**
- 2 **In Registry Editor, locate the following registry key:**

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\
Microsoft SQL Server\<InstanceName>\
MSSQLServer\SuperSocketNetLib\Tcp
```

The TCPDynamicPorts and TCPPort values appear as follows, depending on your port allocation method:

- **Static Port Allocation**

If you configure an instance of SQL Server to use a static port, and you have not yet restarted the instance of SQL Server, the registry values are set as follows:

TCPDynamicPorts = Last port used

TCPPort = New static port that you set by using the Server Network Utility.

If you configure an instance of SQL Server to use a static port, and you restart the instance of SQL Server, the registry values are set as:

TCPDynamicPorts = Blank

TCPPort = New static port that you set by using the Server Network Utility.

- **Dynamic Port Allocation**

If you configure an instance of SQL Server to use dynamic port allocation, and you have not yet restarted the instance of SQL Server, the registry values are set as follows:

TCPDynamicPorts = Blank

TCPPort = 0

However, if you configure an instance of SQL Server to use dynamic port allocation, and you restart the instance of SQL Server, the registry values are set as follows:

TCPDynamicPorts = Current port used

TCPPort = Current port used

## Configuring a Static Port

To configure an instance of SQL Server to use a static port, do the following:

- 1 **Start the Server Network Utility by doing one of the following:**
  - Click **Start > Programs > Microsoft SQL Server > Server Network Utility**.
  - Click **Start > Run**. In the **Open** box, type `svrnetcn.exe`, and then click **OK**.

The **SQL Server Network Utility** dialog box appears.

- 2 **In the SQL Server Network Utility dialog box, click the General tab.**
- 3 **In the Instance(s) on this server list, select your instance of SQL Server.**
- 4 **In the Enabled Protocols list box, click TCP/IP, and then click Properties.**



If TCP/IP protocol is disabled, enable it by clicking **TCP/IP** in the **Disabled Protocols** list box, and then clicking **Enable**.

- 5 **In the Default port box, type a number for the static port and then click OK.**



The static port that you specify must not be the same as the dynamic port on which your instance of SQL Server is currently listening. For example, if your instance of SQL Server is currently listening on dynamic TCP/IP port 1400, type 1500 for the new static port.

- 6 **Click OK twice.**
- 7 **Restart the instance of SQL Server.**
- 8 **View the SQL Server error logs to verify whether or not the instance of SQL Server is currently using the static port.**

## Configuring a Dynamic Port

To configure your instance of SQL Server to use a dynamic port, do the following:

- 1 **Start the Server Network Utility by doing one of the following:**
  - Click **Start > Programs > Microsoft SQL Server > Server Network Utility**.
  - Click **Start > Run**. In the **Open** box, type `svrnetcn.exe`, and then click **OK**.
- 2 In the **SQL Server Network Utility** dialog box, click the **General** tab.
- 3 In the **Instance(s) on this server list**, select your instance of SQL Server.
- 4 In the **Enabled Protocols** list box, click **TCP/IP**, then click **Properties**.



If TCP/IP protocol is disabled, enable it by clicking TCP/IP in the Disabled Protocols list box, and then clicking **Enable**.

- 5 Type **0** in the **Default port** box, and then click **OK**.
- 6 Click **OK** twice.
- 7 Restart the instance of SQL Server.
- 8 View the SQL Server error logs to verify whether the instance of SQL Server is currently using the dynamic port.



If you have a clustered instance of SQL Server, and you follow the specified steps on a cluster node, you may notice that the `TCPDynamicPorts` and `TCPPEndpoint` registry values on other cluster nodes retain the old values. When you move the SQL Server group to the corresponding cluster node, and then bring SQL Server online on the cluster node, the registry values on the cluster nodes will reflect the updated values.

## SQL Server Scripts

You can do the following with SQL Server scripts:

- Determine whether or not values are in the database
- Check the buffer cache
- Run a script using [up.time](#)

You can use the following sample SQL Server script to check whether or not SQL Server is running on your system. The script determines if a customer ID was found. If the customer ID was found, the script returns the following:

OK – Customer id found

If the customer ID was not found, the script returns the following:

CRIT – Customer id not found

### Sample SQL Server Script

```
SET NOCOUNT ON
DECLARE
 @CONTENT varchar(255),
 @RC int
USE Northwind
SELECT @CONTENT = CustomerID from Customers where
CustomerID='ALFKIZ'
IF @@RowCount > 0
BEGIN
 SELECT @RC = 0
 SELECT "OK - Customer id found"
END
ELSE
BEGIN
 SELECT @RC = 2
 SELECT "CRIT - Customer id not found"
END
EXIT(select @RC)
go
```



## SQL Server Locks

Microsoft SQL Server uses locking to ensure the integrity of transactions and consistency in the database. Locking prevents data that users are reading from being changed by other users, and prevents multiple users from simultaneously changing the same data. If locking is not used, data within a database may become incorrect, and queries executed against that data may produce unexpected or invalid results. SQL Server automatically applies locks. However, you can make your applications more efficient by customizing or minimizing database locking.

The Locks object in SQL Server provides information about locks that are applied to individual resource types. Locks are held on SQL Server resources such as rows read or rows that were modified during a transaction, and prevent the concurrent use of a resource by multiple transactions.

SQL Server can lock the following resources:

| Resource | Description                                                      |
|----------|------------------------------------------------------------------|
| Database | A database.                                                      |
| Extent   | A contiguous group of eight data pages or index pages.           |
| Key      | A row lock within an index.                                      |
| Page     | An eight kilobyte data page or index page.                       |
| RID      | The identifier for a row in the database.                        |
| Table    | An entire table in the database, including all data and indexes. |

SQL Server locks resources using the following resource lock modes:

- Shared (S) Locks  
Allow concurrent transactions to read a resource. No other transactions can modify the data while shared (S) locks exist on the resource. Shared (S) locks are released when the data has been read, unless:
  - The transaction isolation level is set to repeatable read or higher.

- A locking hint is used to retain the shared (S) locks for the duration of the transaction.
- Update (U) Locks

Prevent a common form of deadlock. A typical update pattern consists of a transaction:

- Reading a record.
- Acquiring a shared (S) lock on the resource.
- Modifying the resource, which requires conversion to an exclusive (X) lock.

If two transactions acquire shared-mode locks on a resource and then attempt to update data concurrently, one transaction attempts to convert the lock to an exclusive (X) lock.

The shared-mode-to-exclusive lock conversion must wait because the exclusive lock for one transaction is not compatible with the shared-mode lock of the other transaction. In this case, a lock wait occurs.

The second transaction attempts to acquire an exclusive (X) lock for its update. Because both transactions are converting to exclusive (X) locks, and they are each waiting for the other transaction to release its shared-mode lock, a deadlock occurs.

Update (U) locks are used to avoid this problem. Only one transaction can obtain an update (U) lock to a resource at a time. If a transaction modifies a resource, the update (U) lock is converted to an exclusive (X) lock. Otherwise, the lock is converted to a shared-mode lock.

- Exclusive (X) Locks

Prevent access to a resource by concurrent transactions. No other transactions can read or modify data locked with an exclusive (X) lock.

- Intent Locks

Indicates that SQL Server wants to acquire a shared (S) lock or an exclusive (X) lock on some of the resources lower in the hierarchy. Intent locks include intent shared (IS), intent exclusive (IX), and shared with intent exclusive (SIX).

For example, a shared intent lock placed at the table level means that a transaction intends to place shared (S) locks on pages or rows within that

table. Setting an intent lock at the table level prevents another transaction from acquiring an exclusive (X) lock on the table containing that page.

- Schema (Sch) Locks

Used when a table data definition language (DDL) operation – such as adding a column or dropping a table – is being performed.

Schema stability (Sch-S) locks are used when compiling queries and do not block any transactional locks, including exclusive (X) locks. Other transactions can continue to run while a query is being compiled, including transactions with exclusive (X) locks on a table. However, DDL operations cannot be performed on the table.

- Bulk Update (BU) Locks

Used when bulk copying data into a table and either the TABLELOCK hint is specified, or the Table Lock on Build Load table option is set using `sp_Tableoption`. BU locks allow processes to bulk copy data concurrently into the same table while preventing other processes that are not bulk copying data from accessing the table.



# APPENDIX D

## Securing up.time

.....

This appendix explains how to secure the up.time monitoring station and agents:

*Implementing SSL for the Monitoring Station.....480*

*Securing the DataStore.....485*

*Securing Agents using SSL.....487*

## Implementing SSL for the Monitoring Station

SSL (Secure Sockets Layer) is a protocol that encrypts information travelling over a network. You can implement SSL to enable the [up.time](#) monitoring station to securely communicate with various agents and remote hosts.

Implementing SSL for the [up.time](#) monitoring station involves the following steps:

- Stopping the up.time Services
- Modifying the uptime.conf file
- Configuring SSL
- Starting the up.time Services

### Stopping the up.time Services

For information on stopping the up.time DataStore, core, and Web server services see “Stopping the up.time Services” on page 53.

### Modifying the uptime.conf file

The uptime.conf file contains configuration and operating parameters for [up.time](#). You must first add SSL-related entries to the uptime.conf file.

To modify the uptime.conf file, do the following:

**1 Open the uptime.conf file in a text editor.**

In Windows, the uptime.conf file is in the folder C:\Program Files\uptime software\uptime4\.

On Solaris and Linux systems, the uptime.conf file is in the directory /usr/local/uptime4.

**2 Add the following line to the file:**

httpContext=https://<hostname>

Where `<hostname>` is the name of the server on which the monitoring station is running, for example `localhost`.

Adding this line enables the monitoring station to communicate with the following using SSL:

- Agents installed on the hosts being monitored.
- The user's Web browser.

### **3 Save the file.**

For more information on the `uptime.conf` file, see

**3 Create a file named `passphrase.cmd` using a text editor.**

When run, this file prompts you for the password for your certificate key.

**4 Do the following:**

- Type the following in the file:

```
@echo your_ssl_server_password
```

Where `your_ssl_server_password` is the password for your generated certificate keys,

- Save the file.

**5 Copy the following files to the directory in which `up.time` is installed (The default installation directory is `C:\Program Files\uptime software\uptime4\`).**

- `uptime_ssl_server.key`
- `uptime_ssl_server.crt`
- `passphrase.cmd`

**6 Open the file `httpd.conf` (found in the directory `C:\Program Files\uptime software\uptime4\apache\conf`) in a text editor.**

**7 Add the following lines to the beginning of the file `httpd.conf`:**

- `SSLEngine on`
- `SSLMutex default`
- `SSLPassPhraseDialog "exec:C:/Program Files/uptime software/uptime4/passphrase.cmd"`
- `SSLCertificateFile "C:/Program Files/uptime software/uptime4/uptime_ssl_server.crt"`
- `SSLCertificateKeyFile "C:/Program Files/uptime software/uptime4/uptime_ssl_server.key"`
- `SSLSessionCache none`

**8 In the file `httpd.conf`, add the following after the entry `listen 9999`:**

```
listen 9997
```

You can set the listen port to any valid port number.



**9 Add the following lines to the end of the file `httpd.conf`:**

```
<VirtualHost *:9997>
SSLEngine off
DocumentRoot "C:/Program Files/uptime software/uptime4/
ui"
ServerName YourServerName
</VirtualHost>
```

Where `YourServerName` is the name of your server.

**10 Uncomment the following line in the file `httpd.conf`:**

```
LoadModule ssl_module/mod_ssl.so
```

**11 Save the file and exit the text editor.**

## Configuring SSL on Solaris or Linux

To configure SSL on Solaris or Linux, do the following:

**1 Generate or obtain a server certificate.**

You can generate your own non-recognized certificate using a tool like OpenSSL. A non-recognized certificate is one that does not come from a certificate-issuing authority. See <http://www.openssl.org> for more information.

You can also purchase a recognized certificate from a vendor such as Verisign (<http://www.verisign.com>) or Thawte (<http://www.thawte.com>).

**2 Create the following certificate keys:**

- `uptime_ssl_server.key`
- `uptime_ssl_server.crt`

**3 Do the following to create a file named `passphrase` that will prompt you for the password for your certificate key:**

- Open a text editor and create a new file.
- Type the following:  

```
echo your_SSL_server_password
```

## *Implementing SSL for the Monitoring Station*

Where your `_SSL_server_password` is the password for your generated certificate keys.

- Save the file and make it executable.

**4 Copy the files** `uptime_ssl_server.key` and `uptime_ssl_server.crt` **to the directory** `/usr/local/uptime4`.

**5 Add the following lines to the file** `/usr/local/uptime4/apache/conf/httpd.conf`:

- `SSLEngine on`
- `SSLPassPhrasedialog exec:/<path>/conf/passphrase`
- `SSLMutex default`
- `SSLCertificateFile /<path>/conf/uptime_ssl_server.crt`
- `SSLCertificateKeyFile /<path>/conf/uptime_ssl_server.key`
- `SSLSessionCache none`

In the above lines, `<path>` is the path in which the Web server is installed – `/usr/local/uptime4/apache/`.

## **Starting the up.time Services**

For information on stopping the up.time DataStore, core, and Web server services see “Starting the up.time Services” on page 52.

## Securing the DataStore

The **up.time** DataStore is a MySQL database that holds system and performance data, configuration information, user information, and reports. To further secure the DataStore, and ensure the safety and integrity of the information that it contains, you should change the passwords of the default database users.

### Database Users

**up.time** is installed with the following MySQL users:

- root (password: uptimerocks)  
The user who has complete control over the database.
- uptime (password: uptime)  
A privileged user who can create, delete, flush, and drop tables; create and delete indices on tables; select, insert, and delete data in tables. This user cannot perform administrative tasks, such as creating other MySQL users, killing threads, and the like.



If you change the uptime user's password in the **up.time** Web interface, you must update the user name in `uptime.conf` and then restart the core.

- reports (password: reports)  
An ODBC user who can select data in tables so they can generate reports using other applications. This user not used in **up.time**.  
For information on setting up an ODBC driver, see the uptime software Knowledge Base article “Connecting to the up.time DataStore via ODBC”.

### Change the Passwords for the Default Users

**up.time** is installed with three user accounts (see “Database Users” on page 485 for more information). While each user account has a password, you should change the passwords for these accounts. You can change passwords for MySQL users from the command line.

To change the passwords, do the following:

- 1 At the command prompt on the monitoring station, log into MySQL as user root by typing the following:**

```
mysql -u root
SET PASSWORD FOR 'user'@'localhost' = PASSWORD (new_password);
where:
```

- SET PASSWORD is the MySQL command that assigns or changes a password to a user account.
- user is user name for which you want to change the password.
- PASSWORD (); is the MySQL function that encrypts the new password.
- new\_password is the updated password.

- 2 Repeat step 1 for each MySQL user listed in the section “Database Users” on page 485.**

- 3 If you want to limit the domains from which users can access the database, type the following:**

```
SET PASSWORD FOR 'user'@'domain_name' = PASSWORD (new_password);
```

where domain\_name is the domain or IP address from which the user will be allowed to access the database. If the user attempts to access the database from a different domain or IP address, that user will not be able to log in.

## Securing Agents using SSL

You can encrypt the port on which the monitoring station communicates with an agent to ensure that communication between the monitoring station and the agents are secure. You can encrypt the communication port using the Stunnel utility, which is available from <http://www.stunnel.org>.

By default, agents communicate with the monitoring station using port 9998. You should encode a different port for secure communication – for example, port 9997.

To secure agents using SSL, do the following:

- 1 **Download, install, and configure stunnel for your platform.**
- 2 **Generate an stunnel encryption key and certificate by running the following command:**

```
make cert
```

This creates a file named `stunnel.pem`. For more information on generating certificates and keys, see the documentation for Stunnel.



To generate a key, you must have the OpenSSL software installed on the host. Contact your system administrator for more information.

- 3 **As user root, add the following command (which starts the agent in SSL mode) to the system startup script.**

The location of the startup scripts varies depending on your platform. For example, on Linux the script is `/etc/rc.local`.

```
/usr/sbin/stunnel -P /var/run -p /usr/local/etc/stunnel/
stunnel.pem -f -d 9997 -l /opt/uptime-agent/bin/
uptmagnt &
```



# APPENDIX E

## 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:

|                                                       |            |
|-------------------------------------------------------|------------|
| <i>NOTICE TO USER..</i>                               | <i>490</i> |
| <i>License.....</i>                                   | <i>491</i> |
| <i>Intellectual Property and Confidentiality.....</i> | <i>492</i> |
| <i>License Fees .....</i>                             | <i>494</i> |
| <i>Term and Termination.....</i>                      | <i>494</i> |
| <i>Remedies and Indemnification .....</i>             | <i>495</i> |
| <i>Disclaimer.....</i>                                | <i>496</i> |
| <i>Limitation of Liability .....</i>                  | <i>496</i> |
| <i>General Terms .....</i>                            | <i>497</i> |

## END USER LICENSE AGREEMENT

### NOTICE TO USER

#### Legal Contract

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”).

#### Binding Representation

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.

## 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

### 3.1 License Keys

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.

### 3.2 Refunds

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

### 4.1 Term

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.

### 4.2 Termination

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 Infringement or Piracy**

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, wilful 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 Claims, Actions, Liabilities, etc.**

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 lawyers 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

### 6.1 Licensed “As Is”

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.

### 6.2 Software Selection Responsibility

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

### 7.1 Cumulative 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.

## 7.2 Limitation Application

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](http://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

## A

acknowledging alerts 96  
action profiles 134  
    applying to applications 138  
    creating 134  
    editing 138  
    overview 16  
    viewing 137  
actions 16  
Active Directory monitor 170  
adding  
    alert settings 125  
    applications 90  
    groups 73  
    monitor information 118  
    multiple systems 434  
    nested groups 74  
    notification groups 300  
    Novell NRM system 440  
    OIDs 272  
    service groups 139  
    systems 76  
    user groups 298  
    views 93  
    VMware instance 86  
advanced monitors 114, 279  
    choosing 116  
    custom 115, 282  
    external check 115, 286  
    guidelines 281  
    overview 280  
    types 115  
    with retained data 115, 284  
agent monitors 113  
    File System Capacity 153  
    overview 152  
    Performance Check 155  
    Process Count Check 159  
agentless monitors 114  
agents  
    installing 43  
    Linux 45

pSeries 46  
    with HMC 46  
    without HMC 48  
Solaris 44  
UNIX 44  
Windows 43  
    overview 17  
alert profiles 129  
    custom formats 130  
    editing 132  
    understanding 15  
alerts  
    acknowledging 96  
    alert profiles 15  
    applying to applications 133  
    creating profiles 129  
    custom formats 130  
    editing profiles 132  
    monitor settings 125  
    overview 13  
    profiles 129  
Application Availability report 357  
application monitors  
    ESX v3 Workload 193  
    Exchange 176  
    IIS 179  
    Live Splunk Listener 198  
    Splunk Query 196  
    up.time Agent 174  
    WebLogic 182  
    WebSphere 187  
applications  
    adding 90  
    applying action profiles 138  
    applying alert profiles 133  
    deleting 95  
    editing 92  
    maintenance 103  
    offline 103  
    status in Global Scan 103  
    viewing details 92  
    viewing in Global Scan 103  
Archive Policy 424  
auto discovery 82  
    ESX v3 83  
    pSeries with HMC 84  
    using 83

## Index

### C

cloning service monitors [127](#)

Config Panel [9](#), [423](#)

    Archive Policy [424](#)

- Run Queue Length 390
- Run Queue Occupancy 390
- Usage 388
- disk performance statistics 409
- displaying process information 419
- exporting 384
- file system capacity 413
- formatting elements 383
- Graph Editor 380
- Graphing Tool 379
- instance motion 418
- Java 378
- LPAR entitlement 406
- memory usage 395
  - Cache Hit Rate 395
  - Free 395
  - Free Swap 396
  - generating 397
  - Paging Statistics 396
- Multi-CPU Usage 392
- network 407
  - errors 407
  - generating 408
  - I/O 407
- Novell NRM 416
- overview 12, 378, 386
- process 398
  - creation rate 399
  - generating 399
  - number of processes 398
  - running, blocked, waiting 398
- Quick Snapshot 387
- setting date and time ranges 27
- TCP retransmits 400
- top 10 disks 411
- trend lines 382
- user activity 401
- viewing quick snapshot 387
- viewing system status 387
- VXVM stats 414
- workload 402
- workload top 10 404
- groups
  - adding 73
  - adding nested 74
  - Global Scan 102

## H

- host check 4, 62
  - changing 142
  - ping 142
  - up.time agent 142
- HTTP (Web services) monitor 243

## I

- icons 10
  - Clone 10, 127
  - critical 103
  - Delete 11, 299
  - Edit 10, 132, 138, 142
  - View 10
- IIS monitor 179
- IMAP (Email Retrieval) monitor 247
- installation
  - agents 43
    - Linux 45
    - pSeries 46
      - with HMC 46
      - without HMC 48
    - Solaris 44
    - UNIX 44
    - Windows 43
  - guidelines 32
  - monitoring station 35
    - UNIX/Linux/AIX 38
    - Windows 36
  - post-installation tasks 42
  - requirements 33
    - browsers 34
    - hardware 34
    - monitoring station 33
    - up.time agents 34
  - upgrading 42
- Instance Motion graphs 418
- interface
  - overview 6
  - up.time tool bar 6
    - Config 9
    - Global Scan 7
    - My Enterprise 7
    - My Portal 7
    - Reports 9

Services 8  
Users 8

## L

LDAP monitor 249  
license information 428  
Live Splunk Listener monitor 198  
LPAR  
    adding 87  
    entitlement graphs 406  
    workload graphs 405  
LPAR workload graphs 405

## M

mail servers 427  
monitor template 117  
    Advanced View 119  
    alert settings 125  
    configuring response time 122  
    configuring settings 119  
    identification 117  
    monitoring period 126  
    response time 121  
    Standard View 119  
    thresholds 121  
    timing setting options 123  
    timing settings 122  
monitoring periods 429  
monitors 111  
    adding alert settings 125  
    adding information 118  
    adding timing settings 124  
    advanced 114, 279  
        external check 286  
        with retained data 284  
    agent  
        File System Capacity 153  
        overview 152  
        Performance Check 155  
        Process Count Check 159  
    agentless  
        NIS/YP 255  
    alert settings 125  
    application  
        Exchange 176  
        IIS 179

Live Splunk Listener 198  
Splunk query 196  
up.time Agent 174  
WebLogic 182  
WebSphere 187  
cloning 127  
comparisons 120  
database  
    MySQL (Advanced Metrics) 202  
    MySQL (Basic Checks) 209  
    Oracle (Advanced) 211  
    Oracle (Basic Checks) 214  
    Oracle Tablespace 217  
    SQL Server (Advanced Metrics) 224  
    SQL server (Basic Checks) 220  
    SQL Server Tablespace 228  
    Sybase 233  
editing performance 143  
getting help 126  
identification 117  
monitor settings 119  
monitoring period 126  
monitoring periods 429  
network  
    DNS 238  
    FTP 241  
    HTTP 243  
    IMAP 247  
    LDAP 249  
    NFS 253  
    NIS/YP 255  
    NNTP 257  
    Ping 261  
    POP 263  
    SMTP 267  
    SNMP 269  
    SSH 265  
    TCP 277  
overview 21, 112  
template 117  
testing 128  
timing setting options 123  
timing settings 122  
types 113

- Windows
  - Active Directory 170
  - Event Log Scanner 162
  - Service Check 165
  - SMB 168
- Multi-System CPU report 322
- My Enterprise 71
  - acknowledge alerts 96
  - adding applications 90
  - adding groups 73
  - adding nested groups 74
  - adding systems 76
  - application details 92
  - editing applications 92
  - editing system profile 88
  - overview 72
  - views 93
- My Enterprise panel 7
- My Portal panel 7, 67, 68
- MySQL (Advanced Metrics) monitor 202
- MySQL (Basic Checks) monitor 209

## N

- Net-SNMP 269
- Network Bandwidth Report 342
- Network graphs 407
- network monitors
  - DNS 238
  - FTP 241
  - HTTP 243
  - IMAP 247
  - LDAP 249
  - NFS 253
  - NIS/YP 255
  - NNTP 257
  - Ping 261
  - POP 263
  - SMTP 267
  - SNMP 269
  - SSH 265
  - TCP 277
- NFS monitor 253
- NIS/YP monitor 255
- NNTP (Network News) monitor 257
  - command implementation 257
  - response category 258
  - response codes 258

- notification groups 300
  - adding 300
  - editing 301
  - overview 300
  - viewing 301
- Novell NRM 440
  - Abended Thread Count 443
  - abended thread count 443
  - adding 440
  - Allocated Server Processes 442
  - allocated server processes 442
  - Available Disk Space 445
  - available disk space 445
  - Available Event Control Blocks (ECBs) 445
  - Available Memory 444
  - available memory 444
  - Available Server Processes 442
  - available server processes 442
  - Connection Usage 444
  - connection usage 444
  - CPU Utilization 443
  - CPU utilization 443
  - Disk Throughput 445
  - disk throughput 445
  - DS Thread Usage 444
  - DS thread usage 444
  - ECBs 445
  - LAN Traffic 445
  - LAN traffic 445
  - Packet Receive Buffers 444
  - packet received buffers 444
  - Statistics Available 441
  - statistics captured 441
  - view/graph information 446
  - Work To Do Response Time 442
  - work to do response time 442

## O

- OIDs
  - deleting 272
  - manually adding 272
- Oracle (Advanced Metrics) monitor 211
- Oracle (Basic Checks) monitor 214
- Oracle Tablespace Check monitor 217

### P

- Performance Check monitor 155
- Ping monitor 261
- POP (Email Retrieval) monitor 263
- Problem Reporting 431
- Process Count Check 159
  - Configuring 159
  - Settings 159
- Process Count Check monitor 159
- pSeries
  - adding LPARs 87
  - enabling SSH on HMC 447

### R

- Report Log 313
  - completed reports 314
  - deleting 315
  - pending reports 313
  - running reports 313
  - viewing 314
- reports 317
  - Application Availability 357
  - background 305
  - CPU Run Queue Threshold 348
  - CPU Utilization Ratio 326
  - CPU Utilization Summary 323
  - Disk I/O Bandwidth 345
  - dynamic 305
  - Enterprise CPU Utilization 332
  - File System Capacity Growth 335
  - File System Service Time Summary 353
  - generating 305
  - generation options 305
    - email 306
    - to screen 305
    - XML 306
  - Multi-System CPU 322
  - Network Bandwidth 342
  - overview 12
  - Report Log 313
  - Resource Usage 318
  - saving 307
  - saving to file system 307
  - scheduling 310
  - searching saved 309
  - Server Virtualization 336
  - Service Monitor Availability 358

- Service Monitor Metrics 329
- Service Monitor Outages 359
- setting date and time ranges 27
- Solaris Mutex Exception 340
- VI3 Workload Profile 371
- viewing saved 308
- VMware Workload 368
- Wait I/O 327
- WebSphere 361
- Reports panel 9
- Resource Scan 109
  - chart 109
  - gauges 109
- Resource Usage report 318
- Response Time 121

### S

- scheduled maintenance 147
  - assigning to a host 149
  - assigning to a service 150
  - profiles 147
  - viewing profiles 149
- search box 63
- searching 63
- Server Virtualization report 336
  - using 339
- service groups 139
  - creating 139
  - editing 140
  - overview 24
- service monitor
  - File System Capacity 153
  - Process Count Check 159
  - SQL Server (Advanced) 224
  - up.time Agent 174
- Service Monitor Availability report 358
- Service Monitor Metrics report 329
- Service Monitor Outages report 359
- service monitors 111
  - adding alert settings 125
  - adding information 118
  - adding timing settings 124
  - advanced 114, 279
    - external check 286
    - with retained data 284
- agent
  - File System Capacity 153

- overview 152
- Performance Check 155
- Process Count Check 159
- alert settings 125
- application
  - ESX v3 Workload 193
  - Exchange 176
  - IIS 179
  - Live Splunk Listener 198
  - Splunk query 196
  - up.time Agent 174
  - WebLogic 182
  - WebSphere 187
- cloning 127
- comparisons 120
- database
  - MySQL (Advanced Metrics) 202
  - MySQL (Basic Checks) 209
  - Oracle 214
  - Oracle (Advanced) 211
  - Oracle tablespace 217
  - SQL Server (Advanced Metrics) 224
  - SQL server (Basic Checks) 220
  - SQL Server Tablespace 228
  - Sybase 233
- getting help 126
- identification 117
- monitor settings 119
- monitoring period 126
- network
  - DNS 238
  - FTP 241
  - HTTP 243
  - IMAP 247
  - LDAP 249
  - NFS 253
  - NIS/YP 255
  - NNTP 257
  - Ping 261
  - POP 263
  - SMTP 267
  - SNMP 269
  - SSH 265
  - TCP 277
- overview 21, 112
- template 117
- testing 128
- timing settings 122
- types 113
- Windows
  - Active Directory 170
  - Event Log Scanner 162
  - Service Check 165
  - SMB 168
- services
  - filtering 64
  - service groups 24
  - viewing 108
- Services panel 8
- SMB Check monitor 168
- SMTP monitor 267
- SNMP 269
  - MIB Browser 270
  - supported versions 270
- SNMP monitor 269
  - adding OIDs 272
  - deleting OIDs 272
  - loading MIBs 271
  - manually adding OIDs 272
  - MIB browser 270
- Solaris Mutex Exception report 340
- Splunk
  - action profile 137
  - Live Splunk Listener monitor 198
  - query monitor 196
- SQL Server (Advanced Metrics) monitor 224
- SQL Server (Advanced) 224
- SQL Server (Advanced) monitor
  - multiple 224
- SQL Server (Basic Checks) monitor 220
- SQL Server locks 475
- SQL Server ports 470
  - configuring dynamic 473
  - configuring static 472
  - finding TCP/IP port 470
  - verifying configuration 470, 471
- SQL Server scripts 474
- SQL Server Tablespace Check monitor 228
- SSH
  - enabling on HMC 447
- SSH monitor 265
- starting up.time 56
- supported Web browsers 34

## Index

Sybase monitor 233  
System List (Syslist) 9  
systems

adding 78  
deleting 95

## T

testing monitors 128  
Top 10 Disks graph 411  
topological dependencies 145  
adding 146  
viewing 146

## U

UNIX vs. Windows 386  
up.time  
administrator account 55  
architecture 3  
exiting 56  
installing 31  
interface 6  
monitoring concepts 4  
overview 2  
service information 59  
SSL 480  
starting 56  
starting and exiting 55  
system information 57  
tool bar 6  
Config 9  
Global Scan 7  
My Enterprise 7  
My Portal 7  
Reports 9  
Services 8  
Users 8  
uptime.conf file 456  
viewing information 57  
up.time Agent monitor 174  
upgrading to up.time 4 42  
uptime.conf 432, 456  
database 457  
Global Scan thresholds 465  
mail server 459  
remote reporting 466  
RSS feed 467

SNMP trap 460  
Splunk 468  
system 462

user groups  
adding 298  
deleting 299  
editing 298  
overview 297  
viewing 298

user roles  
adding 290  
editing 292  
overview 290  
viewing 291

users  
adding 293  
configuring 289  
editing 296  
notification groups 300  
overview 293  
roles 290  
viewing 296

Users panel 8

## V

VI3 Workload Profile report 371  
viewing  
action profiles 137  
all elements 106  
detailed process information 419  
notification groups 301  
Quick Snapshot 387  
report logs 314  
Resource Scan 109  
scheduled maintenance profiles 149  
service information 59  
services 108  
system information 57  
system status 387  
user groups 298  
user roles 291  
users 296  
views  
adding 93  
deleting 95  
VMware 76, 86, 336, 368, 371  
Instance Motion graph 418



VMware instance [86](#)  
VMware Workload report [368](#)  
    using [370](#)  
VXVM Stats graph [414](#)

## W

Wait I/O report [327](#)  
warning threshold [121](#)  
WebLogic monitor  
    configuring [185](#)  
    metrics [182](#)  
Weblogic monitor [182](#)  
WebLogic report  
    using [366](#)  
WebSphere monitor [187](#)  
    configuring [191](#)  
    counters [187](#)  
WebSphere report [361](#)  
    using [363](#)  
Windows Event Log Scanner monitor [162](#)  
Windows File Shares monitor [168](#)  
Windows monitors  
    Active Directory [170](#)  
    Event Log Scanner [162](#)  
    Service Check [165](#)  
    SMB [168](#)  
Windows Service Check monitor [165](#)  
Workload graphs [402](#)  
Workload Top 10 graphs [404](#)