



**up.time User Guide  
version 6**

## **Copyright © 2011 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 685) 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 mechantability 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.

iText is used under the Lesser General Public License (LGPL).

Oracle and Solaris are registered trademarks, 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.

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  
Canada  
M5V 3B1

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

## **Contacting Sales**

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

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

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

## **Contacting Support**

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

uptime software offers user support through the following:

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

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

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

# TABLE OF CONTENTS

## Welcome to up.time

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

## Understanding up.time

<i>Understanding the up.time Interface</i> .....	<b>6</b>
<i>up.time Tool Bar</i> .....	6
<i>System List</i> .....	10
<i>Icons</i> .....	11
<i>System Icons</i> .....	11
<i>Understanding Reports and Graphs</i> .....	<b>13</b>
<i>Understanding Reports</i> .....	13
<i>Understanding Graphs</i> .....	13
<i>Understanding Agents</i> .....	<b>14</b>
<i>Understanding Major and Minor Versions</i> .....	14
<i>Understanding the up.time DataStore</i> .....	<b>16</b>
<i>Connecting to the DataStore Using ODBC</i> .....	16
<i>Understanding Service Monitors</i> .....	<b>18</b>
<i>Understanding Database Monitors</i> .....	18

<i>Understanding Agentless Monitors Using Net-SNMP</i> .....	18
<b><i>Understanding Services</i></b> .....	<b>21</b>
<i>Understanding Service Groups</i> .....	21
<b><i>Understanding the Status of Services</i></b> .....	<b>22</b>
<b><i>Understanding Dates and Times</i></b> .....	<b>23</b>
<b><i>Understanding Retained Data</i></b> .....	<b>25</b>

## **Installing up.time**

<b><i>Installation Plan</i></b> .....	<b>28</b>
<b><i>Installation Requirements</i></b> .....	<b>29</b>
<i>up.time Monitoring Station</i> .....	29
<i>up.time Agents</i> .....	30
<b><i>Installing the up.time Monitoring Station</i></b> .....	<b>31</b>
<i>Before You Begin</i> .....	31
<i>Installing the Monitoring Station on Windows</i> .....	32
<i>Installing the Monitoring Station on Solaris or Linux</i> .....	34
<i>Installing the Monitoring Station as a Virtual Appliance</i> .....	37
<b><i>Post-Installation Tasks</i></b> .....	<b>39</b>
<i>Configuring the Monitoring Station to Use Oracle</i> .....	39
<b><i>Upgrading to up.time 5</i></b> .....	<b>41</b>
<b><i>Installing Agents</i></b> .....	<b>42</b>
<i>Installing Agents on Windows</i> .....	42
<i>Installing Agents on Solaris</i> .....	43
<i>Installing Agents on UNIX</i> .....	44
<i>Installing Agents on Linux</i> .....	44
<i>Installing Agents on IBM pSeries Servers</i> .....	45

## Getting Started

<b>Accessing and Exiting up.time</b> .....	<b>50</b>
Setting Up the Administrator Account .....	50
Accessing up.time .....	51
Exiting up.time.....	51
<b>Viewing System and Service Information</b> .....	<b>52</b>
Viewing System Information.....	52
Viewing Service Information.....	54
<b>Searching and Filtering</b> .....	<b>59</b>
Using the Search Box .....	59
Filtering Service Instances .....	60
<b>Audit Logging</b> .....	<b>62</b>
Enabling the Audit Log.....	62

## Using My Portal

<b>Overview</b> .....	<b>64</b>
Assistance.....	64
My Preferences.....	65
Latest up.time Articles.....	65
up.time Information .....	65
My Alerts .....	65
Saved Reports.....	66
Custom Dashboards .....	66

## Defining and Managing Your Infrastructure

<b>Overview</b> .....	<b>68</b>
<b>Working with Systems</b> .....	<b>69</b>
Adding Systems or Network Devices .....	71

<i>Auto Discovery</i> .....	76
<i>Adding VMware Instances to up.time</i> .....	81
<i>Adding Individual LPARs to up.time</i> .....	83
<i>Agentless WMI Systems</i> .....	83
<i>Novell NRM Systems</i> .....	88
<i>Adding Multiple Systems</i> .....	94
<i>Editing a System Profile</i> .....	100
<b><i>Working with Applications</i></b> .....	<b>103</b>
<i>Adding Applications</i> .....	103
<i>Viewing Details About Applications</i> .....	105
<i>Editing Applications</i> .....	106
<b><i>Working with SLAs</i></b> .....	<b>107</b>
<b><i>Working with Groups</i></b> .....	<b>108</b>
<i>Adding Groups</i> .....	108
<i>Adding Nested Groups</i> .....	109
<i>Editing Groups</i> .....	110
<b><i>Working with Views</i></b> .....	<b>111</b>
<i>Adding Views</i> .....	111
<i>Adding Nested Views</i> .....	112
<i>Editing Views</i> .....	113
<b><i>Deleting Elements, Applications, and Views</i></b> .....	<b>114</b>
<b><i>Acknowledging Alerts</i></b> .....	<b>115</b>
<b>Overseeing Your Infrastructure</b>	
<b><i>Overview</i></b> .....	<b>118</b>
<i>Viewing More Information</i> .....	119
<i>Groups and Views in the Global Scan Panel</i> .....	120
<b><i>Viewing All SLAs</i></b> .....	<b>121</b>



- SLA Status Indicators..... 121
- Generating an SLA Detailed Report..... 123
- SLA View Types ..... 123
  
- Viewing All Applications..... 126**
  - Condensed View ..... 127
  - Detailed View ..... 128
  
- Viewing All Elements..... 129**
  
- Viewing All Services..... 131**
  
- Viewing the Resource Scan Report..... 132**
  - Performance Gauges ..... 132
  - 24-Hour Performance Graphs..... 133
  - Elements Chart..... 133
  
- Viewing Scrutinizer Status ..... 135**
  
- Changing Reporting Thresholds ..... 136**
  
- Using Service Monitors**
  
- Overview..... 138**
  
- Using Service Monitors..... 139**
  - Using Agent Monitors..... 139
  - Using Agentless Monitors ..... 140
  - Using Advanced Monitors ..... 141
  - Selecting a Monitor ..... 143
  
- The Monitor Template ..... 144**
  - Monitor Identification ..... 144
  - Adding Monitor Identification Information..... 145
  - Monitor Settings Configuration..... 145
  - Configuring Warning and Critical Thresholds..... 147

Monitor Timing Settings .....	149
Monitor Alert Settings.....	151
Monitoring Period Settings .....	153
Getting Additional Help .....	153
<b>Cloning Service Monitors.....</b>	<b>154</b>
<b>Testing Service Monitors.....</b>	<b>155</b>
<b>Service Groups.....</b>	<b>156</b>
Creating Service Groups.....	156
Creating VMware vSphere Service Groups .....	157
Editing Service Groups.....	158
Editing VMware vSphere Service Groups.....	160
<b>Changing Host Checks.....</b>	<b>161</b>
Change a Host Check.....	161
<b>The Platform Performance Gatherer.....</b>	<b>162</b>
Editing the Platform Performance Gatherer .....	162
<b>Topological Dependencies.....</b>	<b>164</b>
Adding Topological Dependencies .....	165
Viewing Topological Dependencies.....	166
<b>Scheduling Maintenance .....</b>	<b>167</b>
Creating Scheduled Maintenance Profiles .....	167
Viewing Scheduled Maintenance Profiles .....	168
Scheduling Maintenance for a Host .....	168
Scheduling Maintenance for a Service.....	169
Putting an Element or Group into Temporary Maintenance Mode .....	169
	169
<b>Monitoring VMware vSphere</b>	
<b>Overview.....</b>	<b>172</b>

<b>VMware vSphere Monitoring Concepts .....</b>	<b>173</b>
About up.time vSync .....	173
vSphere Components as up.time Elements .....	174
vSphere Components and Topological Dependencies .....	174
vSphere Components and Service Groups .....	175
Reports and Graphs for vCenter Server Metrics .....	176
 <b>Managing vSync .....</b>	 <b>178</b>
Managing vCenter Inventories in up.time .....	178
vCenter vSync Settings .....	181
 <b>Standalone Monitoring for vCenter VMs .....</b>	 <b>184</b>
Detaching Ignored VMs .....	186
 <b>vCenter Configuration Mirroring .....</b>	 <b>188</b>
 <b>vSphere Element Configurations .....</b>	 <b>192</b>
vCenter Element Profiles .....	192
ESX Server Element Profiles .....	193
VM Element Profiles .....	195
 <b>Agent Monitors</b>	
 <b>Overview .....</b>	 <b>200</b>
 <b>File System Capacity .....</b>	 <b>201</b>
Configuring File System Capacity Monitors .....	201
 <b>Performance Check .....</b>	 <b>204</b>
Configuring Performance Check Monitors .....	204
 <b>Process Count Check .....</b>	 <b>208</b>
Configuring Process Count Check Monitors .....	208

## VMware Monitors

<b>Overview</b> .....	<b>212</b>
<b>vSphere Performance Monitors</b> .....	<b>213</b>
Datacenter and Cluster Performance.....	213
Resource Pool and vApp Performance.....	215
<b>ESX Server Monitors</b> .....	<b>218</b>
vSphere ESX Server Performance .....	218
ESX (Advanced Metrics) .....	222
ESX Workload.....	225
<b>Power State Monitors</b> .....	<b>229</b>
ESX Server Power State .....	229
VM Instance Power State.....	231

## Microsoft Windows Monitors

<b>Windows Event Log Scanner</b> .....	<b>236</b>
Configuring Windows Event Log Scanner Monitors.....	236
<b>Windows Service Check</b> .....	<b>240</b>
Configuring Windows Service Check Monitors .....	240
<b>Windows File Shares (SMB)</b> .....	<b>243</b>
Configuring Windows File Shares (SMB) Monitors .....	243
<b>Active Directory</b> .....	<b>245</b>
Configuring Active Directory Monitors.....	245

## Application Monitors

<b>Uptime Agent</b> .....	<b>250</b>
Configuring Uptime Agent Monitors .....	250

<b>Exchange</b> .....	<b>252</b>
<i>Configuring Exchange 2003 Monitors</i> .....	252
<i>Configuring Exchange Monitors</i> .....	254
<b>IIS</b> .....	<b>258</b>
<i>Configuring IIS Monitors</i> .....	258
<b>WebLogic</b> .....	<b>261</b>
<i>Monitoring WebLogic 8</i> .....	264
<i>Configuring WebLogic 8 Monitors</i> .....	265
<i>Monitoring WebLogic 9–11</i> .....	266
<i>Configuring WebLogic Monitors</i> .....	267
<b>WebSphere</b> .....	<b>269</b>
<i>Deploying the WebSphere Performance Servlet</i> .....	272
<i>Configuring WebSphere Monitors</i> .....	273
<b>Web Application Transactions</b> .....	<b>275</b>
<i>Using the Web Application Transaction Monitor</i> .....	275
<i>Configuring Web Application Transaction Monitors</i> .....	276
<i>Viewing and Diagnosing Web Transaction Performance</i> .....	279
<i>Using Web Transaction Performance in SLA Reports</i> .....	280
<b>Email Delivery Monitor</b> .....	<b>282</b>
<i>Configuring Email Delivery Monitors</i> .....	282
<i>Diagnosing and Reporting Email Delivery Problems</i> .....	285
<b>Splunk Query</b> .....	<b>288</b>
<i>Configuring Splunk Query Monitors</i> .....	288
<b>Live Splunk Listener</b> .....	<b>290</b>
<i>Before You Begin</i> .....	291
<i>Configuring the Live Splunk Listener Monitor</i> .....	292

## Database Monitors

<b>MySQL (Advanced Metrics)</b> .....	<b>296</b>
<i>Configuring MySQL (Advanced Metrics) Monitors</i> .....	296
<b>MySQL (Basic Checks)</b> .....	<b>303</b>
<i>Configuring MySQL (Basic Checks) Monitors</i> .....	303
<b>Oracle (Advanced Metrics)</b> .....	<b>305</b>
<i>Configuring Oracle (Advanced Metrics) Monitors</i> .....	305
<b>Oracle (Basic Checks)</b> .....	<b>308</b>
<i>Configuring Oracle (Basic Checks) Monitors</i> .....	308
<b>Oracle Tablespace Check</b> .....	<b>311</b>
<i>Configuring Oracle Tablespace Check Monitors</i> .....	311
<b>SQL Server (Basic Checks)</b> .....	<b>314</b>
<i>Configuring SQL Server (Basic Checks) Monitors</i> .....	314
<b>SQL Server (Advanced Metrics)</b> .....	<b>318</b>
<i>Using Multiple SQL Server (Advanced Metrics) Monitors</i> .....	318
<i>Configuring SQL Server (Advanced Metrics) Monitors</i> .....	319
<b>SQL Server Tablespace Check</b> .....	<b>322</b>
<i>Structure of a SQL Server Database</i> .....	322
<i>Configuring SQL Server Tablespace Check Monitors</i> .....	323
<b>Sybase</b> .....	<b>327</b>
<i>Configuring Sybase Monitors</i> .....	327

## Network Service Monitors

<b>DNS</b> .....	<b>332</b>
<i>Before You Begin</i> .....	332
<i>Configuring DNS Monitors</i> .....	333

<b>FTP</b> .....	<b>335</b>
Configuring FTP Monitors .....	335
<b>HTTP (Web Services)</b> .....	<b>337</b>
Configuring HTTP (Web Services) Monitors .....	337
<b>IMAP (Email Retrieval)</b> .....	<b>341</b>
Configuring IMAP (Email Retrieval) Monitors.....	341
<b>LDAP</b> .....	<b>343</b>
Before You Begin .....	343
Configuring LDAP Monitors.....	344
<b>NFS</b> .....	<b>347</b>
Configuring NFS Monitors.....	347
<b>NIS/YP</b> .....	<b>349</b>
Configuring NIS/YP Monitors .....	349
<b>NNTP (Network News)</b> .....	<b>351</b>
Command Implementation .....	351
Response Category .....	352
Response Codes.....	352
Configuring NNTP (Network News) Monitors.....	353
<b>Ping</b> .....	<b>355</b>
Configuring Ping Monitors.....	355
<b>POP (Email Retrieval)</b> .....	<b>357</b>
Configuring POP (Email Retrieval) Monitors.....	357
<b>SSH (Secure Shell)</b> .....	<b>359</b>
Configuring SSH (Secure Shell) Monitors.....	359
<b>SMTP (Email Delivery)</b> .....	<b>361</b>
Configuring SMTP (Email Delivery) Monitors.....	361

<b>SNMP</b> .....	<b>363</b>
<i>Net-SNMP</i> .....	363
<i>SNMP MIB Browser</i> .....	364
<i>Supported Versions of SNMP</i> .....	364
<i>Using the SNMP MIB Browser</i> .....	364
<i>Configuring SNMP Monitors</i> .....	367

<b>TCP</b> .....	<b>370</b>
<i>Configuring TCP Monitors</i> .....	370

## Advanced Monitors

<b>Overview</b> .....	<b>374</b>
<i>Before You Begin</i> .....	375

<b>Custom Monitors</b> .....	<b>376</b>
<i>Configuring Custom Monitors</i> .....	376

<b>Custom with Retained Data</b> .....	<b>378</b>
<i>Configuring Custom Monitors with Retained Data</i> .....	378

<b>External Check</b> .....	<b>380</b>
<i>Configuring External Check Monitors</i> .....	381

<b>Plug-In Monitors</b> .....	<b>382</b>
<i>Installing Plug-In Monitors</i> .....	382

## Configuring Users

<b>Working with User Roles</b> .....	<b>386</b>
<i>Adding User Roles</i> .....	386
<i>Viewing User Roles</i> .....	387
<i>Editing User Roles</i> .....	388

<b>Working with Users</b> .....	<b>389</b>
---------------------------------	------------



Adding Users.....	389
Viewing Users .....	392
Editing User Information.....	392
<b>Working with User Groups .....</b>	<b>393</b>
Adding User Groups.....	394
Viewing User Groups .....	394
Editing User Groups.....	394
Deleting User Groups.....	395
<b>Managing Distribution Lists .....</b>	<b>396</b>
Adding Distribution Lists .....	396
Viewing Distribution Lists .....	397
Editing Distribution Lists .....	397
<b>Working with Notification Groups .....</b>	<b>399</b>
Adding Notification Groups.....	399
Viewing Notification Groups .....	400
Editing Notification Groups.....	400
<b>Changing How Users Are Authenticated.....</b>	<b>401</b>
Active Directory Authentication .....	401
LDAP Authentication.....	404
up.time DataStore Authentication.....	406
<b>Working with Service Level Agreements</b>	
<b>Overview.....</b>	<b>410</b>
<b>SLAs, Service Monitors, and SLOs .....</b>	<b>411</b>
<b>Viewing Service Level Agreements.....</b>	<b>412</b>
Viewing SLA Status.....	412
Viewing SLA Details.....	412
<b>SLA Compliance Calculation .....</b>	<b>415</b>

<i>Reporting SLA Status</i> .....	415
<i>Handling Simultaneous Service Downtime</i> .....	416
<i>A Note About SLOs and Compliance</i> .....	417
<b><i>SLA-Creation Strategies</i> .....</b>	<b>418</b>
<i>Setting Up and Gathering Data for Monitors</i> .....	418
<i>Identifying Outages and Improvable Performance</i> .....	418
<i>Developing Baselines</i> .....	420
<b><i>Working with SLA Reports</i> .....</b>	<b>422</b>
<b><i>Adding and Editing SLA Definitions</i> .....</b>	<b>423</b>
<i>Adding a Service Level Agreement</i> .....	423
<i>Adding Service Level Objectives to an SLA</i> .....	425
<i>Associating Alert and Action Profiles to an SLA</i> .....	426

## Alerts and Actions

<b><i>Understanding Alerts</i> .....</b>	<b>430</b>
<i>Understanding the Alert Flow</i> .....	431
<b><i>Alert Profiles</i> .....</b>	<b>433</b>
<i>Enabling the Windows Messaging Service</i> .....	433
<i>Creating Alert Profiles</i> .....	434
<i>Viewing Alert Profiles</i> .....	435
<i>Editing Alert Profiles</i> .....	436
<i>Associating Alert Profiles to Elements</i> .....	436
<b><i>Working with Custom Alert Formats</i> .....</b>	<b>437</b>
<i>Custom Alert Format Variables</i> .....	438
<b><i>Action Profiles</i>.....</b>	<b>441</b>
<i>VMware vCenter Orchestrator Workflow Actions</i> .....	441
<i>SNMP Trap Actions</i> .....	442
<i>Creating Action Profiles</i> .....	443
<i>Viewing Action Profiles</i> .....	447

<i>Editing Action Profiles</i> .....	447
<b><i>Monitoring Periods</i></b> .....	<b>449</b>
<i>Adding Monitoring Periods</i> .....	449

## Understanding Report Options

<b><i>Overview</i></b> .....	<b>452</b>
<b><i>Generating Reports</i></b> .....	<b>453</b>
<i>Report Generation Options</i> .....	454
<b><i>Saving Reports</i></b> .....	<b>456</b>
<i>Saving Reports to the File System</i> .....	456
<i>Viewing Saved Reports</i> .....	457
<b><i>Scheduling Reports</i></b> .....	<b>459</b>
<b><i>The Report Log</i></b> .....	<b>462</b>
<i>Viewing Report Logs</i> .....	463
<i>Deleting Report Log Entries</i> .....	464

## Using Reports

<b><i>Reports for Performance and Analysis</i></b> .....	<b>466</b>
<i>Resource Usage Report</i> .....	466
<i>Multi-System CPU Report</i> .....	470
<i>CPU Utilization Summary Report</i> .....	471
<i>CPU Utilization Ratio Report</i> .....	474
<i>Wait I/O Report</i> .....	475
<i>Service Monitor Metrics Report</i> .....	477
<b><i>Reports for Capacity Planning</i></b> .....	<b>480</b>
<i>Enterprise CPU Utilization Report</i> .....	480
<i>File System Capacity Growth Report</i> .....	483

<i>Server Virtualization Report</i> .....	484
<i>Solaris Mutex Exception Report</i> .....	488
<i>Network Bandwidth Report</i> .....	490
<i>Disk I/O Bandwidth Report</i> .....	493
<i>CPU Run Queue Threshold Report</i> .....	497
<i>File System Service Time Summary Report</i> .....	501
<b><i>Reports for Service Level Agreements</i></b> .....	<b>505</b>
<i>SLA Summary Report</i> .....	505
<i>SLA Detailed Report</i> .....	506
<b><i>Reports for Availability</i></b> .....	<b>508</b>
<i>Application Availability Report</i> .....	508
<i>Incident Priority Report</i> .....	509
<i>Service Monitor Availability Report</i> .....	512
<i>Service Monitor Outages Report</i> .....	513
<b><i>Reports for J2EE Applications</i></b> .....	<b>515</b>
<i>WebSphere Report</i> .....	515
<i>WebLogic Report</i> .....	518
<b><i>Reports for Virtual Environments</i></b> .....	<b>522</b>
<i>VM Sprawl Report</i> .....	522
<i>vSphere Workload Report</i> .....	524
<i>VMware Workload Report</i> .....	528
<i>VMware Infrastructure Density Report</i> .....	531
<i>LPAR Workload Report</i> .....	532

## Understanding Graphing

<b><i>Graphing in up.time</i></b> .....	<b>536</b>
<i>Graphing Tool</i> .....	537
<b><i>Using the Graph Editor</i></b> .....	<b>538</b>
<i>Working with Trend Lines</i> .....	540
<i>Formatting Individual Graph Elements</i> .....	541

Exporting Graphs .....	542
Changing the Look and Feel of a Graph .....	542

## Using Graphs

<b>Overview</b> .....	<b>544</b>
UNIX vs. Windows Performance Monitoring .....	544
<b>Viewing the Status of a System</b> .....	<b>545</b>
Viewing a Quick Snapshot .....	546
<b>Monitoring CPU Performance</b> .....	<b>547</b>
Usage (% busy) .....	547
Run Queue Length .....	549
Run Queue Occupancy .....	549
Generating a CPU Performance Graph .....	550
<b>Multi-CPU Usage</b> .....	<b>551</b>
Generating a Multi-CPU Usage Graph .....	551
<b>Graphing Memory Usage</b> .....	<b>554</b>
Used .....	554
Cache Hit Rate .....	554
Paging Statistics .....	555
Free Swap .....	555
Generating a Memory Usage Graph .....	556
<b>Graphing Processes</b> .....	<b>557</b>
Number of Processes .....	557
Process Running, Blocked, Waiting .....	557
Process Creation Rate .....	558
Generating a Process Graph .....	558
<b>Graphing TCP Retransmits</b> .....	<b>559</b>
Generating a TCP Retransmits Graph .....	559

<b>Graphing User Activity .....</b>	<b>560</b>
<i>Generating a User Activity Graph .....</i>	<i>560</i>
<b>Workload Graphs .....</b>	<b>561</b>
<i>Generating a Workload Graph .....</i>	<i>562</i>
<b>Workload Top 10 Graphs .....</b>	<b>564</b>
<i>Generating a Workload Top 10 Graph .....</i>	<i>564</i>
<b>LPAR Workload Graphs .....</b>	<b>565</b>
<i>Generating an LPAR Workload Graph .....</i>	<i>565</i>
<i>LPAR CPU Utilization Graphs .....</i>	<i>566</i>
<b>Network Graphs .....</b>	<b>567</b>
<i>I/O .....</i>	<i>567</i>
<i>Errors .....</i>	<i>567</i>
<i>NetFlow .....</i>	<i>568</i>
<i>Generating a Network Graph .....</i>	<i>568</i>
<b>Disk Performance Statistics Graph .....</b>	<b>570</b>
<i>Generating a Disk Performance Statistics Graph .....</i>	<i>570</i>
<b>Top 10 Disks Graph .....</b>	<b>572</b>
<i>Generating a Top 10 Disks Graph .....</i>	<i>572</i>
<b>File System Capacity Graph .....</b>	<b>574</b>
<i>Generating a File System Capacity Graph .....</i>	<i>574</i>
<b>VXVM Stats Graph .....</b>	<b>575</b>
<i>Generating a VXVM Stats Graph .....</i>	<i>575</i>
<b>Novell NRM Graphs .....</b>	<b>577</b>
<i>Generating a Novell NRM Graph .....</i>	<i>578</i>
<b>Instance Motion Graphs .....</b>	<b>579</b>
<i>Generating an Instance Motion Graph .....</i>	<i>579</i>

**Displaying Detailed Process Information ..... 580**  
    Generating Detailed Process Information ..... 581

## **Graphing VMware vSphere Components**

**Overview..... 584**

**Viewing the Status of a vSphere Element..... 585**

### **CPU Performance Graphs ..... 593**

    CPU Workload ..... 593  
    CPU Usage ..... 594  
    Multi-CPU Usage ..... 596  
    Wait and Ready Time..... 597

### **Memory Usage Graphs ..... 600**

    Memory Workload..... 600  
    Memory Usage..... 602  
    Memory Profile..... 604  
    Memory: Guest Paging ..... 605  
    Memory Swap..... 606

### **Process Graphs ..... 609**

    Process Workload..... 609  
    Detailed Process Information..... 610  
    Number of Processes..... 613

### **Network Graphs ..... 615**

    Network Workload..... 615  
    Network I/O..... 616  
    Network Errors ..... 618

### **Disk Graphs..... 621**

    Disk Workload..... 621  
    Disk I/O..... 622  
    Disk Errors ..... 625

<i>Disk Latency</i> .....	626
<i>Disk Storage Capacity</i> .....	628
<i>Disk Storage Profile</i> .....	630

<b><i>VMware Management Graphs</i></b> .....	<b>632</b>
<i>VM Instance Motion</i> .....	632
<i>Generating an Instance Motion Graph</i> .....	632
<i>Power Consumption</i> .....	633
<i>Power States</i> .....	634

## Configuring and Managing up.time

<b><i>Overview</i></b> .....	<b>638</b>
<i>Modifying up.time Config Panel Settings</i> .....	639
<i>Modifying uptime.conf File Settings</i> .....	639
<i>Stopping and Restarting up.time Services</i> .....	640

<b><i>Interfacing with up.time</i></b> .....	<b>642</b>
<i>Database Settings</i> .....	642
<i>Monitoring Station Web Server</i> .....	644
<i>SMTP Server</i> .....	644
<i>Configuring Global Data Collection Methods</i> .....	646
<i>RSS Feed Settings</i> .....	647
<i>VMware vCenter Orchestrator Integration</i> .....	649
<i>Web Application Monitor Proxy Settings</i> .....	650
<i>Remote Reporting Settings</i> .....	651
<i>User Interface Instance Settings</i> .....	652
<i>Scrutinizer Settings</i> .....	652
<i>Splunk Settings</i> .....	653

<b><i>Archiving the DataStore</i></b> .....	<b>655</b>
<i>Archive Categories</i> .....	656
<i>Configuring an Archive Policy</i> .....	656
<i>Restoring Archived Data</i> .....	657
<i>Exporting and Importing the DataStore</i> .....	658



<b><i>up.time Diagnosis</i></b> .....	<b>661</b>
System Event Logging .....	661
Audit Logs .....	662
Problem Reporting .....	662
<b><i>up.time Measurement Tuning</i></b> .....	<b>664</b>
Service Monitor Thread Counts.....	664
Status Thresholds .....	664
Platform Performance Gatherer Check Intervals .....	667
<b><i>Report Storage Options</i></b> .....	<b>668</b>
Changing the Number of Days Reports Are Cached .....	668
Changing the Published Report Location.....	669
<b><i>Resource Usage Report Generation</i></b> .....	<b>670</b>
<b><i>Monitoring Station Interface Changes</i></b> .....	<b>671</b>
Status Alert Acknowledgement .....	671
3D Graphs.....	671
Custom Dashboard Tabs.....	672
<b><i>License Information</i></b> .....	<b>673</b>

## Reference

<b><i>Frequency Definitions</i></b> .....	<b>676</b>
<b><i>Time Period Definitions</i></b> .....	<b>677</b>
Building Blocks.....	677
Basic Expressions.....	679
Combining Expressions and Excluding Time Periods.....	682

## End User License Agreement

<b><i>NOTICE TO USER</i></b> .....	<b>686</b>
------------------------------------	------------

1. License.....	686
2. Intellectual Property and Confidentiality.....	688
3. License Fees.....	689
4. Term and Termination.....	690
5. Remedies and Indemnification.....	690
6. Disclaimer .....	691
7. Limitation of Liability.....	691
8. General Terms.....	692

## Index

# CHAPTER 1

## Welcome to up.time

---

This chapter introduces [up.time](#) in the following sections:

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

## Introducing up.time

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

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

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

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

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

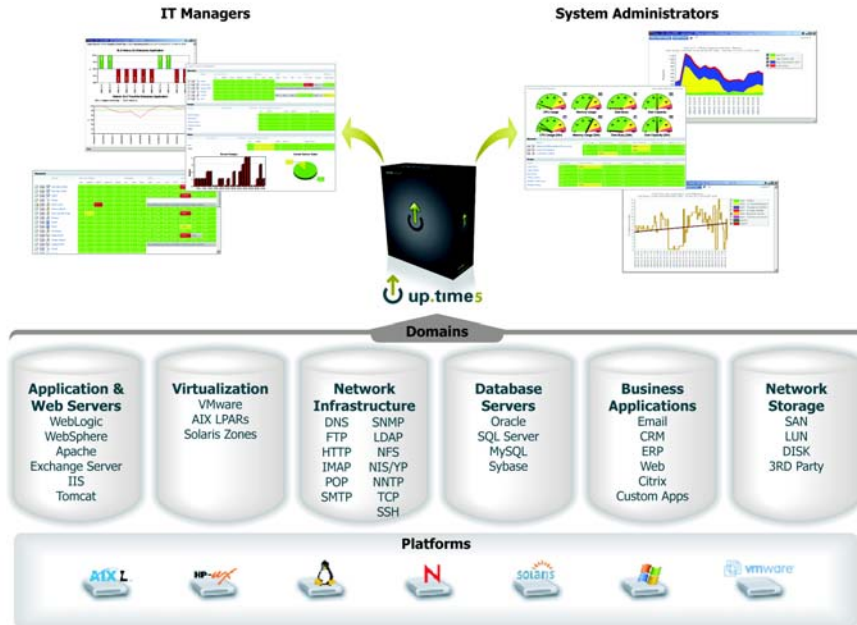
## Who Should Read This Guide

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

- system administrators who want to use up.time to monitor a single system or multiple systems in a distributed environment at a single datacenter
- 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 (i.e., an *agent*) 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

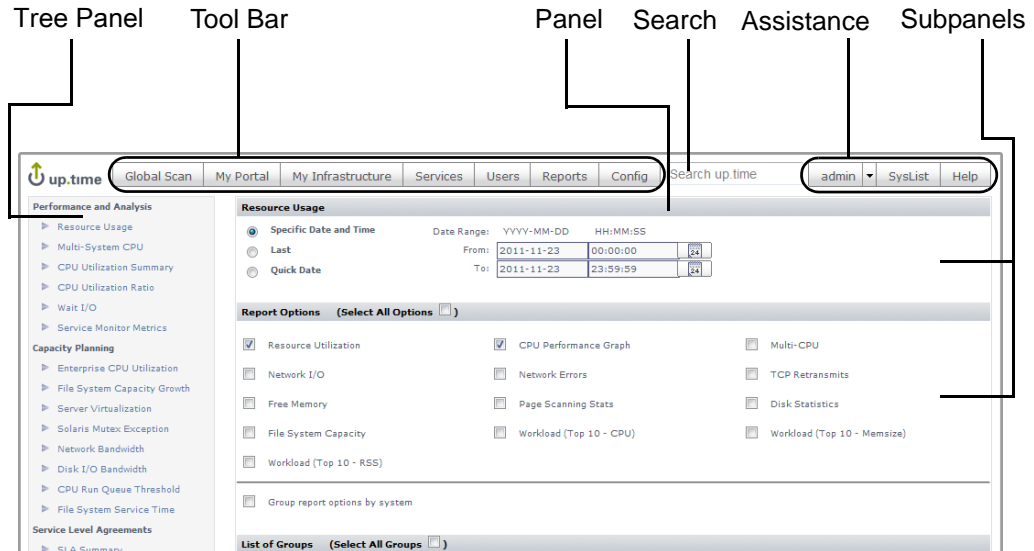
---

This chapter explains underlying concepts in the following sections:

<i>Understanding the up.time Interface</i> .....	6
<i>Understanding Reports and Graphs</i> .....	13
<i>Understanding Agents</i> .....	14
<i>Understanding the up.time DataStore</i> .....	16
<i>Understanding Service Monitors</i> .....	18
<i>Understanding Services</i> .....	21
<i>Understanding the Status of Services</i> .....	22
<i>Understanding Dates and Times</i> .....	23
<i>Understanding Retained Data</i> .....	25

## Understanding the up.time Interface

The up.time Web interface consists of seven 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 Infrastructure
- 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 manage the resources in your infrastructure.

For more information about using the **Global Scan** panel, see “Overseeing Your Infrastructure” on page 117.

## 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 63.

## My Infrastructure

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

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

For more information about using the **My Infrastructure** panel, see “Defining and Managing Your Infrastructure” on page 67.

### Services

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

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

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

For more information about using the **Services** panel, see “Using Service Monitors” on page 137 .

### Users

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

- users
- user groups
- Notification Groups
- user roles

For more information about using the **Users** panel, see “Configuring Users” on page 385.

## 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 465.

## Config

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

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

You can also generate problem reports and edit the `uptime.conf` file from the **Config** panel. For more information about using the **Config** panel, see “Configuring and Managing up.time” on page 637.

## Search

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

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

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

## System List

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

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

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








Display Name	Host Name	Entity Group Name
 <b>AIX DEV LPAR</b>	10.1.1.57	Unix Boxes
 <b>AIX QA LPAR</b>	10.1.1.56	Unix Boxes
 <b>AIX5</b>	aix51	Unix Boxes
 <b>Development Group</b>	Development Group	My Enterprise
 <b>Email Delivery</b>	Email Delivery	My Enterprise
 <b>Enterprise Application</b>	Enterprise Application	My Enterprise
 <b>ESX4</b>	vmh-esx4	VMware Boxes
 <b>ESX7</b>	vmh-esx7	VMware Boxes
 <b>Exchange</b>	uptime-exchange	Email Systems
 <b>lab-t1-2</b>	10.1.1.232	Solaris Boxes
 <b>lab-t1-4</b>	10.1.1.234	Solaris Boxes
 <b>vmh-prod</b>	vmh-prod	VMware Boxes
 <b>WebSphere</b>	lab-websphere51	Windows 2003 Boxes

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 Infrastructure** panels, and identify the type of system that [up.time](#) is monitoring:



Linux



AIX



Solaris



Novell NRM



Windows



HP-UX

## Understanding up.time *Understanding the up.time Interface*



VMware ESX



Net-SNMP



HMC/VIO

## Understanding Reports and Graphs

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

### Understanding Reports

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

For detailed information about reports, see “Using Reports” on page 465.

If you need to regularly run certain reports, you can save them to the **My Portal** panel. See “Scheduling Reports” on page 459 for more information.

### Understanding Graphs

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

You can view graphs in two ways:

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

For more information on graphing, see “Understanding Graphing” on page 535 and “Using Graphs” on page 543

## 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 . 0.

- Minor version

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

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

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

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

## Understanding the up.time DataStore

The DataStore is a database in which [up.time](#) stores different types of information:

- configuration information for [up.time](#)
- configuration and system information for the hosts that you are monitoring
- the performance data gathered by monitors, which is used for generating graphs and reports
- user information, including user names and passwords (encrypted if it is sensitive information)
- the settings for service monitors, Alert and Action Profiles, scheduled maintenance, and host checks
- reports that Monitoring Station users have saved, and are scheduled to run at specific intervals.

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

Access to the DataStore is determined by one of the three installed user accounts: root, uptime, and reports. Each account gives users varying levels of access to the contents of the DataStore. For more information about these accounts, see the uptime software Knowledge Base article “Securing MySQL Database and Adding Users”.

[up.time](#) can also use either an Oracle or MS SQL Server database as its DataStore. If you plan to use either of these databases, refer to our Knowledge Base 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 137.

## 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 363.

## 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 137 .

## Understanding Service Groups

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

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

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

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

See “VMware vSphere Monitoring Concepts” on page 173 VMware vSphere Monitoring Concepts for information on vSync, and “Service Groups” on page 156 for information on creating service groups.

## Understanding the Status of Services

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

- 0 – OK  
The services are functioning properly.
- 1 – Warning  
There is a potential problem with one of more of the services.
- 2 – Critical  
There is a critical problem with one or more services.
- 3 – Unknown  
This status is returned when:
  - The host on which the service sits is offline.
  - The host on which the service sits is in a scheduled maintenance or downtime period.
  - The Monitoring Station could not execute the service monitor.

Each status reflects the state of the service that has been assigned to the system that you are currently viewing. up.time picks up these error codes and triggers an alert or an action. If a service is in a warning or critical state, you can acknowledge an alert so that up.time does not generate subsequent notifications.

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

	Name	Service Status						
		OK	WARN	CRIT	MAINT	UNKN	ACK	1
	Dev1-w2k3se	0	0	1	0	1	0	0
	The Vault	2	0	0	0	0	0	0
	WebSphere	4	0	1	0	0	0	0

The figures in each column in the Global Scan panel indicate the number of services for that particular machine that are in each state. Click a number to view the **System Status** screen for a particular system. See “Viewing the Status of a System” on page 545 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:

<input checked="" type="radio"/> <b>Specific Date and Time</b>	Date Range: YYYY-MM-DD	HH:MM:SS
<input type="radio"/> <b>Last</b>	From: 2008-04-21	00:00:00 
<input type="radio"/> <b>Quick Date</b>	To: 2008-04-21	23:59:59 

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

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

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

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

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

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



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

## Understanding Retained Data

up.time enables you to save some or all of the metrics that its monitors collect to the DataStore. You can use the retained data to generate a Service Metrics report (see “Service Monitor Metrics Report” on page 477) or a Service Metrics graph (see “Viewing System and Service Information” on page 52).

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			
Port	▼	9998	
Use SSL	▼	<input type="checkbox"/> Use SSL	
<b>Web Mail Sends Per Second</b> ▼			Save for Graphing <input checked="" type="checkbox"/>
Warning	▼	is greater than or equal to	70
Critical	▼	is greater than or equal to	100
<b>Web Mail Auths Per Second</b> ▼			Save for Graphing <input checked="" type="checkbox"/>
Warning	▼	is greater than or equal to	40
Critical	▼	is greater than or equal to	75
<b>SMTP Bytes Sent Per Second</b> ▼			Save for Graphing <input type="checkbox"/>



# CHAPTER 3

## Installing up.time

---

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

<i>Installation Plan</i> .....	28
<i>Installation Requirements</i> .....	29
<i>Installing the up.time Monitoring Station</i> .....	31
<i>Installing Agents</i> .....	42

## Installation Plan

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

- identify the system 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 Monitoring Station system must have a CD-ROM drive from which to load the server software. A CD-ROM drive is not required if you have downloaded the **up.time** software from the Internet.

The installation procedure creates the user ID `uptime` on the 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 **up.time** Monitoring Station and **up.time** Agents. Before installation, it is recommended that you check the uptime software Web site (<http://www.uptimesoftware.com>) for the most up-to-date list of hardware and software requirements.

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

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

**Note** – Suse Linux systems may require additional SSL libraries.

### Supported Web Browsers

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

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

### Minimum Hardware Configuration

The hardware configurations for 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 following is the recommended minimum hardware:

- 2.4 GHz dual-core processor
- 2 GB of memory
- 80 GB of disk storage
- 100 Mbps 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 140.



## Installing the up.time Monitoring Station

The Monitoring Station is installed a single directory:

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

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



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

In addition to the (included) MySQL database, **up.time** 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 our Knowledge Base for the additional steps required to enable **up.time** to work with these databases.

### Before You Begin

There are three ways in which to install the **up.time** Monitoring Station:

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

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

**2 From the distribution CD.**

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

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

`up.time_MonitoringStation`

### 3 Imported as a VMware Virtual Appliance.

If you are installing [up.time](#) as an appliance on an ESX server, you can download the package from the [uptime software web site](#), either directly or through the VMware Virtual Appliance Marketplace. Unarchive the Virtual Appliance package and note its location; you will need to locate the `.ovf` file during the import procedure.

Once preparations have been made, refer to the procedures in the “Installing the Monitoring Station on Windows” on page 32, “Installing the Monitoring Station on Solaris or Linux” on page 34, or “Installing the Monitoring Station as a Virtual Appliance” on page 37 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 **If you are upgrading, ensure you have logged out of the [up.time](#) Web application by clicking the Logout button.**
- 2 **Ensure you are logged in to the Monitoring Station system as the local administrator.**

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

### 3 Double click the following file:

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

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

`up.time-5.0.455-win32-x86.exe`

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

- Click **Next** to accept the default location (C:\Program Files\uptime software\uptime).
- In the **Please Choose a Folder** field, type the name of the directory where you want to install the application and then click **Next**.
- Click **Choose** and select a directory from the **Browse for Folder** window.
- To recover the default directory, click **Restore Default Folder**.

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

- Click **Next** to accept the default location (C:\Program Files\uptime software\uptime\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.

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

- Click **Next** to accept the defaults.
- Enter information in the following fields:
  - Email address  
The email address from which the 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`.

## Installing up.time *Installing the up.time Monitoring Station*

- 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`.

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

**11 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.

**12 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.

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

**14 Click Finish.**

## Installing the Monitoring Station on Solaris or Linux

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

## Installing the Monitoring Station

To install the [up.time](#) Monitoring Station on Solaris or Linux, do the following:

- 1 If you are upgrading, ensure you have logged out of the [up.time](#) Web application by clicking the Logout button.**
- 2 Ensure you have logged in to the Monitoring Station system as root.**

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

- 3 Type the following command:**

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

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

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

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

- 4 On the Introduction page, press Enter to continue.**
- 5 On the License Agreement page, carefully read the [up.time](#) end user license agreement. Press Enter to scroll through the agreement.**
- 6 At the DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N) prompt, type `y` and press Enter.**
- 7 Do one of the following to set the directory in which [up.time](#) will be installed:**
  - Press Enter to accept the default location (`/opt/uptime` on Solaris, and `/usr/local/uptime/` on Red Hat and SLES)

- Type a new location at the command prompt (for example, `/opt/uptime` on Solaris), then press Enter.



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

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

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

### 9 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`.
  - **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`.

**10 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.

**11 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

**12 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.

## Installing the Monitoring Station as a Virtual Appliance

To install the [up.time](#) Monitoring Station as a Virtual Appliance, do the following:

- 1 In the Virtual Infrastructure Client, start the procedure to import a virtual appliance.**
- 2 Select the Import from file option, and locate the `up.time.ovf` file you downloaded from the [uptime software web site](#). Click Next.**

- 3 After viewing the Virtual Appliance Details, click Next.**
- 4 On the License Agreement screen, review the up.time end user license agreement, click the Accept all license option, then click Next.**
- 5 Provide configuration information for install:**
  - the name and location of the up.time Virtual Appliance
  - the host or cluster on which the Virtual Appliance will run
  - the resource pool within which it will be run
  - the datastore in which the appliance's data will be kept
  - the network the appliance will use
- 6 Review your selections, then click Finish.**

Wait for the import process to complete.
- 7 In the Virtual Infrastructure Client, navigate to, select the up.time appliance, and power it on.**
- 8 Click the Console tab for the appliance.**
- 9 After initialization, ensure the appliance time is correct.**

The default time zone is PST. The appliance time zone must match that of your monitored infrastructure in order to correctly collect and report performance data.
- 10 After the appliance configuration has been completed, you can log in to the Monitoring Station to begin setting up your monitored environment.**



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 50)
- provide the host name of the SMTP server when you first log in (see “SMTP Server” on page 644)
- install the license for [up.time](#) (see “License Information” on page 673)
- add users and systems (see “Configuring Users” on page 385 and “Defining and Managing Your Infrastructure” on page 67)

## Configuring the Monitoring Station to Use Oracle

If this Monitoring Station installation is for a standalone [up.time](#) instance that is *not* part of a multi-datacenter deployment, skip this section and use the default bundled MySQL implementation; otherwise, you must configure the Monitoring Station to write to an Oracle database instance instead of MySQL. To switch the database used by the Monitoring Station, edit the `uptime.conf` file.

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

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

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

- 2 **Show (i.e., uncomment) the Oracle database settings.**
- 3 **For the `dbHostname` and `dbPort` settings, enter the address and port for your Oracle database server.**

- 4 For the `dbName` setting, provide a name for the Enterprise Monitoring Station's Oracle database instance.
- 5 In the `dbUsername` and `dbPassword` fields, enter the authentication details to access and write to the database.
- 6 Save your changes.
- 7 Use the `resetdb` utility with the `really` option to delete, then recreate the database structure that is used by `up.time` by running the appropriate command:
  - Linux: `/usr/local/uptime/resetdb really`
  - Solaris: `/opt/uptime/resetdb really`
  - Windows: `C:\Program Files\uptime software\uptime\resetdb really`

## Upgrading to up.time 5

If you are using a previous version of **up.time** and intend to upgrade to version 5, 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](mailto:support@uptimesoftware.com) and request the extracted agent which can be installed without using the Windows installer.

**Note** – If the Windows installer requires unavailable service packs – for example, SiteServer or Terminal Server – send an email to [support@uptimesoftware.com](mailto: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 **Log in to the Monitoring Station as the local administrator.**

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

- 3 In Windows Explorer, double click the file `setup.exe`.**
- 4 On the installer Welcome screen, click Next.**
- 5 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.
- 6 Select the checkbox Make available for Everyone option.**
- 7 Click Next.**
- 8 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-<version>.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` can collect workload information from IBM pSeries servers that have logical partitions (LPARs). To have `up.time` collect this information, you must install the latest AIX or Linux agents on the 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 that uses the Integrated Virtual Manager (IVM)

In both cases, you will need to install the agent on each LPAR; whether you use an HMC determines how the agent is installed on 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 LPAR that is on IBM pSeries server with an HMC, do the following:

**1 Ensure you are logged in to the HMC as a super-administrator-level user.**

`up.time` communicates with the HMC to acquire LPAR information.

**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 (e.g., `UptimeAgent-Linux-<version>.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 (e.g., `uptmagnt-AIX-<version>.tar`).

- Run the following command to install the agent:



```
./INSTALL.sh
```



If you are using an HMC, do not install the agent as a Virtual I/O Server by using the “-vio” attribute with the install command. Doing so may lead to conflicts with HMC-managed systems, and can result in incorrect performance statistics.

#### 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 (e.g., `uptmagnt-AIX-<version>.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 LPAR that is not managed by an HMC, but whose partitions are managed by the Integrated Virtual Manager (IVM).

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 (e.g., UptimeAgent-Linux-<version>.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, `uptmagnt-AIX-<version>.tar`.

- Run the following command to install the agent as a Virtual I/O Server:

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

### 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, `uptmagnt-AIX-<version>.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>Accessing and Exiting up.time</i> .....	50
<i>Viewing System and Service Information</i> .....	52
<i>Searching and Filtering</i> .....	59
<i>Audit Logging</i> .....	62

## Accessing and Exiting up.time

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

## Setting Up the Administrator Account

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

To set up the administrator account, do the following:

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

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

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

```
http://localhost:9999
```

The [up.time](#) log in window opens in a Web browser.

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

## Accessing up.time

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

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

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


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

Where <uptime\_hostname> is the name or IP address of the server that is hosting the Enterprise Monitoring Station.

The [up.time](#) log in window opens in the Web browser.

- 3 **Enter your assigned user name in the User Name field.**
- 4 **Enter your assigned password in the Password field.**
- 5 **Click the Login button.**

## Exiting up.time

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

## Viewing System and Service Information

You can view information about the following:

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

### Viewing System Information

To view system information, do the following:

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

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

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

- Info & Rescan

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

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

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

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



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

- **CPU Information**  
Lists the speed (in MHz) of all of the CPUs on the system.
- **Network**  
Lists the network interfaces on the system, as well as the IP addresses of those interfaces.
- **Disks/File System**  
Lists the disks that are on Solaris and Linux systems and the names of the file systems that [up.time](#) is monitoring.
- **Poll Agent**  
Displays the output from an [up.time](#) agent that you suspect may have a problem. You can forward the output to uptime software Client Care when you encounter problems with [up.time](#).
- **Services**  
Lists the services assigned to the system, as well as the interval (in minutes) at which the services are checked.
- **User Groups**  
Lists the user groups that are associated with the system.

## Viewing Service Information

To view system information, do the following:

- 1 In the **Global Scan** or **My Infrastructure** 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 115.

- 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 22.

- **Outages**

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

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

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

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

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

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

- **Availability**

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

## Getting Started *Viewing System and Service Information*

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

The Availability table is shown below:


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

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

[Generate Graph](#)

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



- **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 137.

- **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 167.

**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 25.**

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 23.

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

**Service Monitor Metrics**

**Specific Date and Time**      Date Range:    YYYY-MM-DD    HH:MM:SS  
 **Last**      From:    2008-04-18    00:00:00    24  
 **Quick Date**      To:    2008-04-18    23:59:59    24

---

**Current Retained Service Metrics**

Instance Name	Instance Description	<input type="checkbox"/> Select	Variable	Units	Data Type
Exchange		<input type="checkbox"/>	Response time ▾	ms	integer
		<input type="checkbox"/>	SMTP Bytes Received Per Second ▾		integer
		<input type="checkbox"/>	SMTP Bytes Sent Per Second ▾		integer
		<input type="checkbox"/>	SMTP Bytes Total Per Second ▾		integer
		<input type="checkbox"/>	SMTP Connection Errors Per Second ▾		integer
		<input type="checkbox"/>	SMTP Inbound Connections ▾		integer
		<input type="checkbox"/>	SMTP Local Queue Length ▾		integer
		<input type="checkbox"/>	SMTP Messages Per Second ▾		integer
		<input type="checkbox"/>	SMTP Outbound Connections ▾		integer
		<input type="checkbox"/>	Web Mail Auths Per Second ▾		integer
		<input type="checkbox"/>	Web Mail Sends Per Second ▾		integer

[Generate Graph](#)

- Click **Generate Graph**.

## Searching and Filtering

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

### Using the Search Box

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

To use the search box, do the following:

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

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



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

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

**2 Click Go.**

The following information is displayed:

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

- any information in the four custom fields in the system profile (e.g., the job being done by the system, and its physical location)  
For more information, see “Editing a System Profile” on page 100.

## 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.



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

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

The following is an example of an audit log entry:

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

## Enabling the Audit Log

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

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

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



# CHAPTER 5

## Using My Portal

---

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

## Overview

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

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

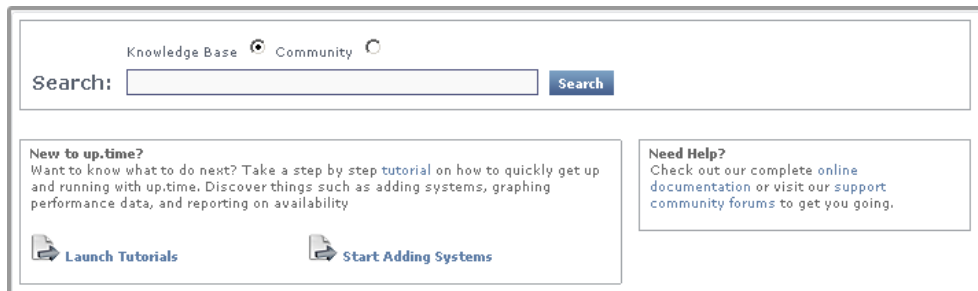
## Assistance

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

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



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

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



## My Preferences

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

- View your user account settings. Click the **View** icon (  ) or your user name to open your account settings in the subpanel. You can also edit your user information by clicking **Edit User**.
- Change your user account settings. Click the **Edit** icon (  ). The Edit User window appears. See “Editing User Information” on page 392 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 Alerts

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

## Saved Reports

The **Saved Reports** tab lists the reports that you have scheduled and saved. For more information on scheduling reports, see “Scheduling Reports” on page 459.

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.

## Custom Dashboards

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

Custom dashboards are visible to members of specific, dashboard-related User Groups. For information on configuring a custom dashboard, see “Custom Dashboard Tabs” on page 672.

# CHAPTER 6

## Defining and Managing Your Infrastructure

---

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

<i>Overview</i> .....	68
<i>Working with Systems</i> .....	69
<i>Working with Applications</i> .....	103
<i>Working with SLAs</i> .....	103
<i>Working with Groups</i> .....	108
<i>Working with Views</i> .....	111
<i>Deleting Elements, Applications, and Views</i> .....	114
<i>Acknowledging Alerts</i> .....	115

## Overview


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

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

## 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 Infrastructure** panels, agent systems are denoted by this icon: 

- Net-SNMP v2 or Net-SNMP v3

Systems that use version 2 of the Net-SNMP protocol, or 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 version 2.

All of the data gathered from Net-SNMP is based on the following MIB implementations:

- RFC 1213 (Management Information Base for Network Management of TCP/IP-based internets)

Presents network interface information.


- UCD-SNMP-MIB

Presents general system state information.

- Host Resources MIB (RFC 2790)

Presents system performance data.

- Node

A device without an agent, but with which [up.time](#) can communicate using an IP address. In the **Global Scan** and **My Infrastructure** panels, nodes are denoted by this icon: 

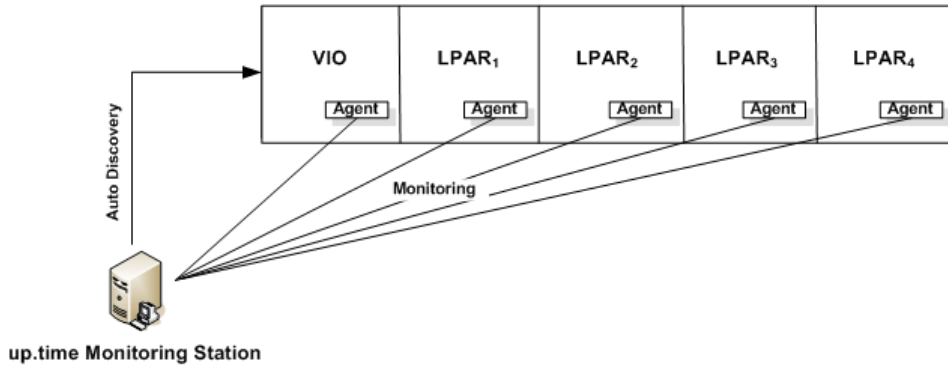
- Novell NRM

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.


- 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 45 for more information.

 You can also add pSeries servers that are managed by a Hardware Management Console (HMC) to [up.time](#) either manually, or 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” on page 94 for more information.

- **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 Infrastructure** panels, virtual nodes are denoted by this icon: 
- **VMware ESX**  
 A system that is running version 3 or 4 of the VMware ESX server software, which enables a single host to run multiple virtual servers and their applications. ESX 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 server.

- **VMware vCenter Server**

A central control point for a VMware vSphere datacenter that includes ESX hosts, VMs, as well as groupings such as clusters, datacenters, vApps, and resource pools. A VMware vCenter server's inventory, system configurations, storage profiles, and performance data can be represented in [up.time](#) alongside physical systems and network devices.

- **WMI Agentless**

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



WMI-based monitoring only works if the Monitoring Station is running on Windows.

## Adding Systems or Network Devices

To add systems or network devices, do the following:

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

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](#) 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
- Net-SNMP v2
- Net-SNMP v3
- Node
- Novell NRM
- pSeries LPAR Server (VIO)
- pSeries LPAR Server (HMC)
- Virtual Node
- VMware ESX
- VMware vCenter Server
- WMI Agentless (only present on Monitoring Stations running on Windows)

**5 Enter the host name of the system in the Host Name 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 If applicable, 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 Configure the system- or device-specific settings:**

- If you selected Agent in step 4 and want to securely access the system, click the **Use SSL** option.
- If you selected **Net-SNMP v2** 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).

- 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 Method (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.

- If you selected **Node** in step 4, optionally select the following check boxes:
  - Is Node Pingable?

This options specifies whether [up.time](#) can contact the node using the ping utility.

There are scenarios in which you might not want the node to be pingable (e.g., you have a firewall in place). Before selecting this check box, you should try to contact the node using the ping utility. If you cannot ping the node, ensure the check box is left cleared. Then, change the default host check for the node. See “Changing Host Checks” on page 161 for more information.
  - Exports NetFlow Data to Scrutinizer?

If Scrutinizer has been integrated with [up.time](#), and is also receiving NetFlow data from the node, select this check box. You will then be able to call a Scrutinizer instance directly from the node’s Graphing tab in [up.time](#).
- 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.
- If you selected **VMware ESX** in step 4, enter information in the following fields:
  - User Name

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

The password required to log into the VMware ESX server.

- If you selected **VMware vCenter Server** in step 4, enter information in the following fields:
  - User Name  
The name of the VMware vCenter administrator account.
  - Password  
The password for the VMware vCenter account.
  - Notification Settings  
When a vSync operation is performed to check for changes to the VMware vCenter inventory, these check boxes indicate whether [up.time](#) will send notifications about, or perform scripted actions in response to, newly discovered ESX servers or VMs . (For more information, see Managing vSync“Managing vSync” on page 178.)
  - Additional VM Guest Performance Data Collection  
These options enable additional monitoring for VMware vCenter VMs using the [up.time](#) agent or WMI. (See Standalone Monitoring for vCenter VMs“Standalone Monitoring for vCenter VMs” on page 184.)  
  
For the [up.time](#) agent, indicate the port on which it is listening, and whether it is securely communicating with [up.time](#) using SSL.  
  
For data collection via WMI, indicate the host and domain on which WMI has been implemented, and the username and password required for access.
- If you selected **WMI Agentless** in step 4, enter information in the following fields:
  - Windows Domain  
The Windows domain in which WMI has been implemented.
  - User Name  
The name of the account with access to WMI on the Windows domain.
  - Password

The password for the account with access to WMI on the windows domain.

- 8 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 68 for more information on defining groups.

- 9 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 156 for more information.

- 10 Click Save.**

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

- 11 If you want to add another system or network device, 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.

- 12 Click Save.**

## Auto Discovery

It can be time consuming to add a large number 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.

- Performs an agent check by scanning systems to determine whether or not agents are installed on them.
- Performs a WMI check by checking whether systems are using WMI to gather metrics (optional).
- Performs an SNMP probe to find any systems that use Net-SNMP (optional).

Systems that are repeatedly discovered through additional checks (e.g., both an agent and WMI implementation are detected on the same system) will, by default, be assigned a type based on the first check that resulted in its discovery. The auto-discovery order is as follows: agent check, WMI check, SNMP probe, node discovery.

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

See [for more information](#).

You can also use the Auto Discovery feature to add VMware ESX systems that are being monitored by Virtual Infrastructure 3 or vSphere 4, 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 Systems” on page 79.
- “Using Auto Discovery to Add pSeries Servers Managed by an HMC” on page 80.

## Using Auto Discovery

To use Auto Discovery, do the following:

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

The **Auto Discovery** window appears.

- 2 To scan for agents, in the Network Address field, indicate which IP address ranges you want [up.time](#) to scan.**

The following formats are acceptable:

- a single IP address (e.g., 10.1.52.1)
- a single subnet (e.g., 10.1.52)
- an IP address range (e.g., 10.1.52-100)

- multiple, comma-separated subnet entries (e.g., 10.1.1, 10.1.4, 10.1.52)
- multiple subnets and an IP address range (e.g., 10.1.1, 10.1.4, 10.1.52-100.)

**3 If you would like to scan for systems using WMI to collect metrics, enter the login information for an administrative Windows account in the following fields:**

- **Windows Domain (optional)**  
The Windows domain in which WMI has been implemented.
- **User Name**  
The name of the account with access to WMI on the Windows domain.
- **Password**  
The password for the account with access to WMI on the windows domain.

Note that this option is only available on Monitoring Stations running on the Windows platform.

**4 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 (e.g., `private`)

**5 Click Continue.**

`up.time` returns a list of the systems that have an IP address within a range that you specified.

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

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

**7 If necessary, edit the details of the system as described in the section “Adding Systems or Network Devices” on page 71.**

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

**8 Repeat steps 4 and 5 for any other systems that you want to add.**



## Using Auto Discovery to Add ESX Systems

Virtual Infrastructure 3 (VI3; also called Virtual Center) is a software suite that manages multiple, physical VMware ESX 3 servers. The latest version, which supports ESX 4, is known as vSphere 4 (vCenter). You cannot directly add VI3 or vCenter systems to [up.time](#); you can, however, use the Auto Discovery feature to point [up.time](#) to a VI3 or vSphere 4 system, then add any or all of the ESX servers it is managing.

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

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

The **Auto Discovery** window appears.

**2 Click the ESX Discovery option.**

**3 Complete the following fields:**

- **Virtual Center Host Name**  
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 servers that are being managed by the VI3 or vSphere 4 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 71.**

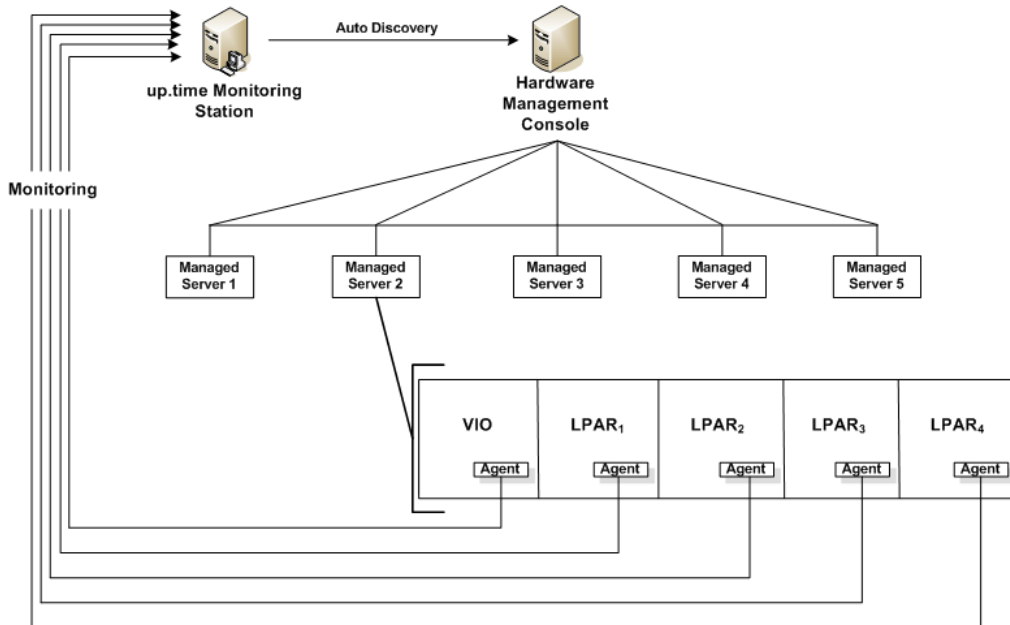
Otherwise, click **Save** in the **Add System/Network Device** window.

**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. 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 latter. See the [up.time software Knowledge Base article](#) entitled “Enabling SSH on the Hardware Management Console” for more information.

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

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

The **Auto Discovery** window appears.

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

**3 Complete the following fields:**

- HMC Host Name

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 71.**

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 Infrastructure 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			
VMware Display Name	IP	Guest OS	Is On?
Dev1-w2k3se		Microsoft Windows Server 2003, Enterprise Edition	N
dev1-rhes4		Red Hat Enterprise Linux 4	N
dev1-sles9	10.1.1.123	Suse Linux Enterprise Server	Y
Dev1-w2k3ee		Microsoft Windows Server 2003, Enterprise Edition	N
css11-w2k3ee-x86		Microsoft Windows Server 2003, Enterprise Edition	N
dev-sles10-x86		Suse Linux Enterprise Server	Y
			<a href="#">Add to up.time</a>
dev1-w2k3se-r2-x86	10.1.1.130	Microsoft Windows Server 2003, Standard Edition	Y
			<a href="#">Add to up.time</a>
dev1-rhes4u5-x86		Red Hat Enterprise Linux 4	Y
			<a href="#">Add to up.time</a>
dev1-rhes4u6-x86		Red Hat Enterprise Linux 4	N
dev2-sles10-x86		Suse Linux Enterprise Server	Y
			<a href="#">Add to up.time</a>
dev1-vista32		Microsoft Windows Vista	N
dev2-rhes4u5-x86		Red Hat Enterprise Linux 4	N
dev2-vista32		Microsoft Windows Vista	N
css4-w2k3ee	10.1.0.221	Microsoft Windows Server 2003, Enterprise Edition	Y
			<a href="#">Add to up.time</a>
lab-v5-sla		Red Hat Enterprise Linux 4	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 Infrastructure 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](#).**

## Agentless WMI Systems

If the Windows-based component of your infrastructure already makes use of WMI (Windows Management Instrumentation), Windows Elements can be configured to use it for data collection as an alternative to the [up.time](#) Agent. Using WMI allows you to avoid the overhead associated with

managing and updating all of the systems on which an **up.time** Agent has been installed.



WMI-based monitoring can only be performed if the Monitoring Station itself is running on Windows.

An Element can be set to use WMI through the following methods:

- its system type is set to “WMI Agentless” when it is first added to **up.time**
- its system type was set to “Agent” when originally added to **up.time**, but is being individually modified to use WMI
- it is part of a bulk agent-to-WMI conversion with other agent-based Elements

Globally defined WMI credentials can be used for the second and third method. In the latter’s case, configuring these is mandatory. Refer to “Configuring Global WMI Credentials” on page 646 for more information.

Regardless of which method is used, when changing a Windows Element’s data collection method, all historical data is retained.

### WMI Requirements

In order to monitor agentless systems through WMI in a secure environment (e.g., through a firewall), you need to create an exception for WMI on the host end. For example, to allow WMI access through Windows Firewall, refer to the following MSDN articles:

- for Windows XP or Windows Server 2003:  
<http://msdn.microsoft.com/en-us/library/aa389286%28v=VS.85%29.aspx>
- for Windows Vista or Windows Server 2008:  
<http://msdn.microsoft.com/en-us/library/aa822854%28v=VS.85%29.aspx>

### Adding a WMI System to up.time

To add an agentless WMI system to **up.time**, do the following:

- 1 On the **up.time** tool bar, click **My Infrastructure**, then click **Add System/Network Device**.**

**2 Complete the Display name in up.time and Description fields.**

See “Adding Systems or Network Devices” on page 71 for more information.

**3 Select WMI Agentless from the Type of System/Device dropdown list.**

**4 In the Host Name field, enter the actual name or IP address of the machine that up.time will be monitoring.**

**5 Select the Use WMI Global Credentials check box if they have been configured, and you would like to use them (see “Configuring Global WMI Credentials” on page 646 for more information); otherwise complete the following fields:**

- **Windows Domain**  
The Windows domain in which WMI has been implemented.
- **Username**  
The name of the account with access to WMI on the Windows domain.
- **Password**  
The password for the account with access to WMI on the windows domain.

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

**7 If you want to associate this system with a Service Group, select its name Service Group dropdown list.**

**8 Click Save.**

## Switching an Element to WMI Data Collection

To change the data collection source for an individual Windows Element from the [up.time](#) Agent to WMI, do the following:

**1 In the Global Scan or My Infrastructure panels, click the name of the Windows server.**

**2 Click the Info tab, then click Info & Rescan.**

- 3 Click the **Edit Collection Method** link found beside the **Collection Method** setting, as shown below:



The **Edit Data Collection Method** window appears.

- 4 Select the **WMI Agentless** data collection option.
- 5 Select the **Use WMI Global Credentials** check box if they have been configured, and you would like to use them (see “**Configuring Global WMI Credentials**” on page 646 for more information); otherwise complete the following fields:

- **Windows Domain**  
The Windows domain in which WMI has been implemented.
- **Username**  
The name of the account with access to WMI on the Windows domain.
- **Password**  
The password for the account with access to WMI on the windows domain.

- 6 Click **Save** to retain your changes and close the pop-up window.

## Switching an Element to Agent-Based Data Collection

To change the data collection source for an individual Windows Element from WMI to the [up.time](#) Agent, do the following:

- 1 In the **Global Scan** or **My Infrastructure** panels, click the name of the **Windows server**.
- 2 Click the **Info** tab, then click **Info & Rescan**.
- 3 Click the **Edit Collection Method** link found beside the **Collection Method** setting, as shown below:





The **Edit Data Collection Method** window appears.

- 4 **Select the up.time Agent data collection option.**
- 5 **Select the Use up.time Agent Global Configuration check box if it has been configured, and you would like to use it (see “Configuring a Global up.time Agent Configuration” on page 647 for more information); otherwise complete the following options:**
  - **Port**

The port through which the up.time Agents communicate with the up.time Monitoring Station.
  - **Use SSL**

Select this check box if the agent securely communicates with the Monitoring Station using SSL.
- 6 **Click Save to retain your changes and close the pop-up window.**

## Converting Multiple Elements to WMI Data Collection

To change multiple agent-based Elements to use WMI for data collection, do the following

- 1 **Ensure the global settings for WMI credentials have been set (see “Configuring Global WMI Credentials” on page 646 for more information).**
- 2 **On the up.time tool bar, click Config.**
- 3 **In the tree panel, click Bulk Element Conversion.**
- 4 **In the Windows Agent Elements section, select the check boxes that correspond to the agent-based Elements whose data collection method is to be changed to WMI.**
- 5 **Click Convert to WMI.**

When the conversion is complete, the lists of agent-based and WMI Elements will be refreshed to reflect the changes.

## Converting Multiple Elements to Agent-Based Data Collection

To change multiple WMI Elements to use the [up.time](#) Agent for data collection, do the following

- 1 **Ensure a global [up.time](#) Agent configuration exists (see “Configuring Global WMI Credentials” on page 646 for more information).**
- 2 **On the [up.time](#) tool bar, click Config.**
- 3 **In the tree panel, click Bulk Element Conversion.**
- 4 **In the WMI Elements section, select the check boxes that correspond to the WMI Elements whose data collection method is to be changed to the [up.time](#) Agent.**
- 5 **Click Convert to Agent.**

When the conversion is complete, the lists of agent-based and WMI Elements will be refreshed to reflect the changes.



For bulk WMI-to-agent conversions, the port used by all of the converted [up.time](#) Agents must match the port specified in the global agent configuration.

## 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.

### 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 Infrastructure 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 71 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 its name from the Group dropdown list.**

**6 If you want to associate this system with a Service Group, select its name Service Group dropdown list.**

**7 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:

- 8 In Novell NRM, click Profiling / Debugging.**
- 9 Check the information for server process functions.**
- 10 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 the following:

- 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.

## Adding Multiple Systems

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.

See [, , and](#) for more information.

You can find examples of entries in a hosts file in the section “Examples of Hosts File Entries” on page 99.

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 69 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: Agent Node Novell NRM Net-SNMP v2 Net-SNMP v3 pSeries LPAR Server (HMC) Virtual Node WMI Agentless
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.  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: <code>public</code> , which enables you to retrieve read-only information. <code>private</code> , which enables you to access all information

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 Element 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> <p><code>MD5</code>, a widely-used method for creating digital signatures.</p> <p><code>SHA</code>, a secure method of creating digital signatures.</p>

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: DES, an older method used to encrypt information. AES, the successor to DES, which is used with a variety of software including SSL servers.
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 <code>true</code> and <code>false</code> .
WMI Domain	The Windows domain in which WMI has been implemented.
WMI Username	The name of the account with access to WMI on the Windows domain.
WMI Password	The password for the account with access to WMI on the windows domain.

## 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\uptime\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 Display Name: prod1 Description: Main production server Type: Agent Service Group: Production Systems Port:9998 Group: Windows 2003 Servers
Node	Host Name: www.myDomain.ca Display Name: Your Domain Description: A Web site Type: Node Group: Web Sites
Novell NRM	Host Name: novell01 Display Name: dn3 Type: Novell NRM SSL: true Port: 546 Group: Unix Boxes Group: Novell System

Host Type	Sample Hosts File Entry
Net-SNMP v2	Host Name: gateway.mydomain.com Display Name: gatewaySNMP Description: snmp v2 Type: Net-SNMP v2 Read Community: myCo-pub
Net-SNMP v3	Host Name: SNMP-1 Display Name: SNMP-1 Description: Net-SNMP system Type: Net-SNMP v3 Read Community: public Username: myUsername Password: myPassword Privacy Password: myOtherPassword Group: Linux Systems
pSeries LPAR	Host Name: 10.1.2.42 Display Name: HMC Managed Server HMC Hostname: 10.1.1.255 Type: pSeries LPAR Server (HMC) Managed Server: Server-7610-31C-SN01B030K Username: hscroot Password: hscroot
Virtual Node	Host Name: router-Toronto Display Name: Toronto Router Description: Router for Toronto branch Type: Virtual Node Pingable: True Group: Routers
WMI Agentless	Host Name: Win7-Production Display Name: Windows 7 Production Description: Win7 agentless/WMI Type: WMI Agentless Group: Windows Boxes WMI Domain: windomain WMI Username: administrator WMI Password: password

## 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 Infrastructure panel, right-click the name of the Element whose profile you want to edit, then click Edit.**

The **Edit System** window appears.

- 2 **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.

- 3 **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(s)**  
One or more monitors can be used to determine the status of the Application as a whole.
- **regular service monitors**  
Other service monitors that are associated with a master service monitor, but are not used to determine the status of the Application as a whole.
- **warning and critical conditions for multiple master service monitors**

You can configure an Application to reach a warning- or critical-level state when a specific number, percentage, or all master service monitors enter those states.

This allows you to give different Applications different levels of robustness by assigning more or less “weight” to their respective groups of master service monitors. As a result, each of your Applications will provide the most accurate status possible, and fewer false positives. For example, a web server cluster of 10 servers might only cause alerts when three of them are down, whereas a mission-critical application will cause an alert when all of its master service monitors fail.

For more information on services, see “Using Service Monitors” on page 137. For information on viewing information about Applications, see “Viewing Details About Applications” on page 105.

## Adding Applications

To add an Application, do the following:

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

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

- 3** Optionally, enter a description for the Application in **Description of Application** field.
- 4** Optionally, select the group of systems in your [up.time](#) environment with which this system will be associated from the **Parent Group** dropdown list.

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

For more information on groups, see “Working with Groups” on page 108.

- 5** In the **Application Status** section, define how many master service monitors must be in a warning- or critical-level state to affect the Application as a whole.

You can include a defined number, a percentage or all master service monitors in this condition.

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

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

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

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

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

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

- 10** Click **Save**.

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

- 11 **If required, associate Alert Profiles with the Application by clicking Edit Alert Profiles when viewing the Application's details.**
- 12 **In the Alert Profile Selector pop-up window, select one or more of the Available Alert Profiles from the list, then click Save.**
- 13 **If required, associate Action Profiles with the Application by clicking Edit Action Profiles when viewing the Application's details.**
- 14 **In the Action Profile Selector pop-up window, select one or more of the Available Action Profiles from the list, then click Save.**

## Viewing Details About Applications

After you have added an Application to [up.time](#), the name of the Application appears in the **My Infrastructure** panel. The name of the Application is a hyperlink.

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

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

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

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

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

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

For information about viewing more details about Applications, see “Viewing System and Service Information” on page 52.

## Editing Applications

To edit an Application, do the following:

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

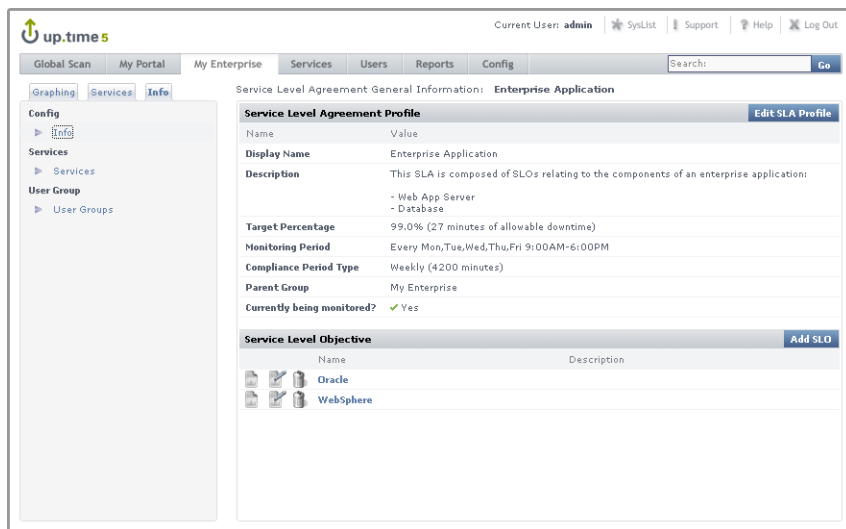
The **Edit Application** window appears.

- 2 **Edit the Application setting as described in “Adding Applications” on page 103.**

## Working with SLAs

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

From the **My Infrastructure** panel, you can view your existing SLA details by clicking the SLA name (see "Viewing SLA Details" on page 412 for more information).



The screenshot displays the up.time web interface. At the top, the user is logged in as 'admin'. The navigation menu includes 'Global Scan', 'My Portal', 'My Enterprise', 'Services', 'Users', 'Reports', and 'Config'. The 'Services' section is active, showing a list of services under 'User Group'.

The main content area displays the 'Service Level Agreement General Information: Enterprise Application'. The 'Service Level Agreement Profile' section includes the following details:

Name	Value
Display Name	Enterprise Application
Description	This SLA is composed of SLOs relating to the components of an enterprise application: - Web App Server - Database
Target Percentage	99.0% (27 minutes of allowable downtime)
Monitoring Period	Every Mon,Tue,Wed,Thu,Fri 9:00AM-6:00PM
Compliance Period Type	Weekly (4200 minutes)
Parent Group	My Enterprise
Currently being monitored?	<input checked="" type="checkbox"/> Yes

The 'Service Level Objective' section shows a table with columns for Name and Description:

Name	Description
Oracle	
WebSphere	

For information about creating and using SLAs, see "Adding and Editing SLA Definitions" on page 423.

## 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, SLAs, or Applications
- the user groups that are allowed to view the systems or Elements in a group (see “Working with User Groups” on page 393 for more information on user groups)



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

## Adding Groups

To add a group, do the following:

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



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

- 5 To give this group its own subgroups, select one or more entries from the Available Groups list, then click Add.
- 6 Select the Elements that you want to add to this group from the Available Elements list, then click Add.
- 7 Select one or more sets of users who can view this group from the Available User Groups list, then click Add.
- 8 Click Save.

## Adding Nested Groups

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

You can assign the following to nested groups:

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

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



Before you begin, ensure that you have at least one parent group defined. For more information, see “Adding Groups” on page 108.

### Adding a Nested Group

To add a nested group, do the following:

- 1 In the My Infrastructure panel, click Add Group.
- 2 Enter a descriptive name for the group in the Group Name field.
- 3 Optionally, enter a description of the group in the Group Description field.

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

## Editing Groups

To edit groups, do the following:

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

The **Edit Element Group** window appears.

- 2 Edit the group as described in “Adding Groups” on page 108.
- 3 Click Save.

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



## Working with Views

Not every user that accesses the Monitoring Station needs to view all Elements that are a part of your infrastructure. Some users may, for example, only need to 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 the Views section on the **Infrastructure** panel, as well as the the **Global Scan** panel.

### Adding Views

To add a view, do the following:

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

This name will appear when listing views in the **Infrastructure** panel.

- 3 **Optionally, enter a description in View Description field.**
- 4 **To make this view a child of an existing one, select it from the Parent View dropdown list.**



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

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

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

- 7 **Select one or more users from the Available Users for View list, then click Add.**

- 8 To add previously defined groups of users, select one or more entries from the Available User Groups list, then click Add.**
- 9 Click Save.**

## Adding Nested Views

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

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

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



Before you begin, ensure that you have at least one parent view defined. For more information, see “Adding Views” on page 111.

### Adding a Nested View

To add a nested view, do the following:

- 1 In the Infrastructure panel, click Add View.**
- 2 In the Add View window, enter a descriptive name in the View Name field.**  
This name will appear when listing views in the **Infrastructure** panel.
- 3 Optionally, enter a description in View Description field.**
- 4 In the Parent View dropdown list, select the view to which this nested view will be subordinate.**
- 5 To give this nested view its own child views, select one or more entries from the Available Element Views list, then click Add.**
- 6 Select one or more users who can view this group from the Available Users list, then click Add.**

- 7 To add previously defined groups of users, select one or more entries from the Available User Groups list, then click Add.
- 8 Click Save.

## Editing Views

To view and edit views, do the following:

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

The **Edit View** window, which contains system and user information, appears.

- 2 Edit the view as described in “Adding Views” on page 111.
- 3 Click Save.

## Deleting Elements, Applications, and Views

If you have administrator privileges, you can delete a Element, or view in the **Infrastructure** panel.

To remove an Element, Application, or View from [up.time](#), do the following:

- 1 **In the Infrastructure panel, locate the Element or up.time grouping you want to permanently remove.**
- 2 **Click the Element or grouping's gear icon.**
- 3 **In the pop-up menu, click Delete.**
- 4 **On the dialog box that appears, click OK.**



You can only delete Elements that were created in [up.time](#). You cannot manually remove Elements that represent VMware vSphere components imported into [up.time](#) via vSync.

## Acknowledging Alerts

When a problem occurs on a system that **up.time** is monitoring, the Monitoring Station sends alerts: these are notifications about the problem, sent to users who are qualified to receive them. If the user role to which they belong is configured to do so, they can also acknowledge an alert.

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

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


To acknowledge alerts, do the following:

- 1 In the Infrastructure panel, click the name of the Element that generated the alert.**

The **System General Information** subpanel appears.

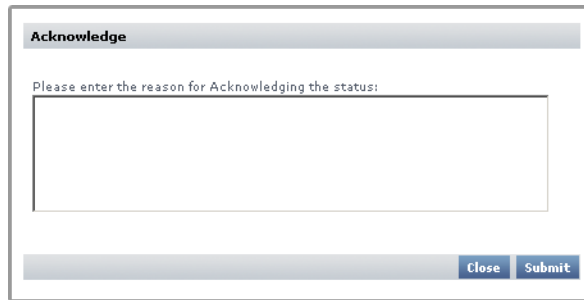
- 2 In the Tree panel, click the Services tab and then click Status.**

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

Monitor	Status	Ack	Last Check	Duration	Monitor Information
File System Capacity	CRIT		2008-04-18 17:32:05	+ 10 days 6h	C: is 100% full
UPTIME-lab-websphere51	OK		2008-04-18 17:28:06	+ 6h 6m	up.time agent running on lab-websphere51, up.time Windows-MS →
PING-lab-websphere51	OK		2008-04-18 17:31:04	+ 10m 17s	Ping completed: 1 sent, 0.0% loss, 0.0ms average round trip →
WebSphere	OK		2008-04-18 17:23:52	+ 10 days 6h	WebSphere Statistics were successfully collected
Plants Response	OK		2008-04-18 17:24:37	+ 10 days 5h	Finished playback - time: 1740 ms

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

The acknowledgement message window appears.




The screenshot shows a dialog box titled "Acknowledge". Inside the dialog, there is a text input field with the placeholder text "Please enter the reason for Acknowledging the status:". Below the input field, there are two buttons: "Close" and "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 name of the Element and service monitor involved
- a comment relating to the alert or reason for acknowledgement

The following is a sample alert acknowledgement message:

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

# CHAPTER 7

## Overseeing Your Infrastructure

---

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

<i>Overview</i> .....	118
<i>Viewing All SLAs</i> .....	126
<i>Viewing All Applications</i> .....	126
<i>Viewing All Elements</i> .....	129
<i>Viewing All Services</i> .....	131
<i>Viewing the Resource Scan Report</i> .....	132
<i>Viewing Scrutinizer Status</i> .....	135
<i>Changing Reporting Thresholds</i> .....	136

# Overview

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

The screenshot shows the up.time Global Scan interface. At the top, it displays the user 'admin' and navigation options like 'SysList', 'Support', 'Help', and 'Log Out'. Below the navigation bar, there are tabs for 'Global Scan', 'My Portal', 'My Infrastructure', 'Services', 'Users', 'Reports', and 'Config'. A search bar is also present.

The main content area is titled 'Global Scan' and includes buttons for 'View SLAs', 'View Applications', 'View All Elements', 'View Resource Scan', and 'View All Services'. The current location is 'My Enterprise'.

**Elements Table:**

Name	Service Status					Outages	CPU				Disk		Memory			
	OK	WARN	CRIT	MAINT	UNKN		ACK	1hr	12hr	24hr	USR	SYS		WIO	TOT	% Used
bogus	0	0	2	0	0	0	0	1	1							
FiltersSNMP	1	0	0	0	0	0	0	7	14	4%	5%	1%	10%	91%	n/a	0%
GingerSNMP	0	0	0	0	0	0	0	0	0	No performance data available. Is the System online?						
Google	0	0	0	0	0	0	0	0	0							
my app	2	0	0	0	0	0	0	0	0							
Novell	1	0	0	0	0	0	0	6	6	n/a	n/a	n/a	1%	0%	n/a	83%
Perl	0	0	0	0	0	0	0	0	0							

**Groups Table:**

Name	Description	Elements	Service Status by Group						
			OK	WARN	CRIT	MAINT	UNKNOWN	ACK	
Crit	-	6	12	0	0	0	0	0	0
Email Systems	-	3	6	0	0	0	0	0	0
Unix Boxes	-	11	23	2	0	0	0	0	0
VMware Boxes	-	3	2	0	1	0	0	0	0
Windows Boxes	-	9	19	0	5	0	0	1	0
Total	-		62	2	6	0	0	1	0

**Recent Outages:** A bar chart showing the number of outages over time from 16:08 to 14:08. The y-axis is labeled 'Number' and ranges from 0 to 4. The x-axis shows time intervals in 2-hour increments.

**Current Service Status:** A pie chart showing the distribution of service statuses: Okay (85.33%), Critical (10.67%), and Warning (2.67%). Unknown is listed as 1.33% but is not visible in the pie chart.

The **Elements** table displays the following information:

- the status and number of services that are associated with the Element
- the number of recent service outages



- CPU usage
- hard disk usage
- memory usage

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


The thresholds for the service status indicators are typically 70% for a warning state, and 90% for a critical state. These thresholds can be customized (see “Changing Reporting Thresholds” on page 136).

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

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

## Viewing More Information

You can view detailed information about an Element by clicking its name. To view the details of each metric (for example, CPU usage) click the number in the column for that variable to go to its Graphing page, where you will be able to generate a graph.

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:



Name	mem	rss	cpu%	mem%	pid	ppid	uid	stime	usr_sys
-bash	2.06 MB	0 KB	0	0	248	0			0s
-bash	2.00 MB	0 KB	0	0	27676	0			0s
-ksh	1.37 MB	0 KB	0	0	27709	0			0s
-sh	800 KB	0 KB	0	0	9740	0			0s
arraymon	1.13 MB	0 KB	0	0	450	0			1s

Messages  
No agent on entity.

**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

When you create groups or views (see “Working with Groups” on page 108 and “Working with Views” on page 111), they appear in their own sections in the **Global Scan** panel. The following information is displayed:

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

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

## Viewing All SLAs

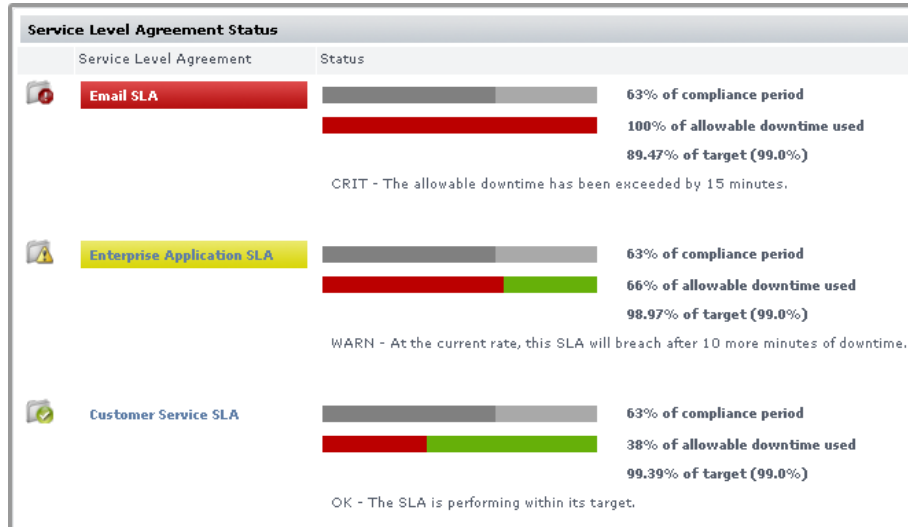
Service level agreements in the **Global Scan** panel indicate whether performance targets are being met. Although the main summary displays the status of the SLA definition as a whole, you can also expand the view to verify how well component service level objectives (SLOs) are meeting targets. (SLOs are made up of monitored services that, as a group, are used to measure a specific performance goal.)

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


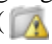
- the list of SLAs, and whether any are in a critical or warning-level state
- headway into the time period during which compliance is measured
- the percentage of allowable downtime used, after which the SLA's status becomes critical

## SLA Status Indicators

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



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

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

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

## Generating an SLA Detailed Report

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

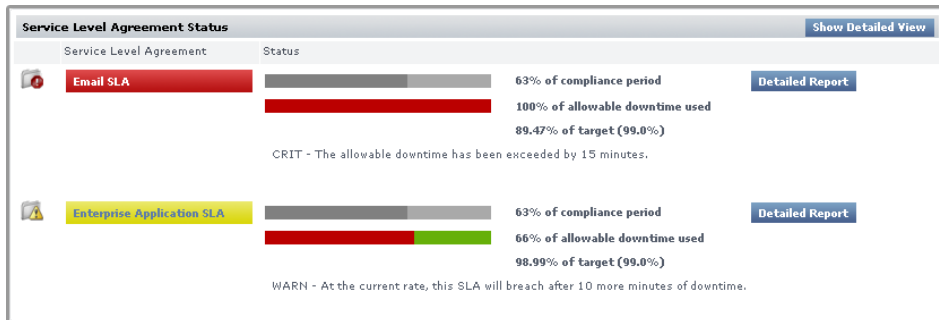
See "Reports for Service Level Agreements" on page 505 for more information.

## SLA View Types

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

### Condensed View

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



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

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

- status message

## Detailed View

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

The screenshot displays the 'Service Level Agreement Status' interface. It features two main sections: 'Email SLA' and 'Enterprise Application SLA'. Each section includes a progress bar, compliance percentage, downtime usage, and a 'Detailed Report' button. Below each SLA is a table of 'Service Level Objectives' (SLOs).

**Service Level Agreement Status** Show Condensed View

**Email SLA** Detailed Report

63% of compliance period  
100% of allowable downtime used  
89.47% of target (99.0%)

CRIT - The allowable downtime has been exceeded by 15 minutes.

Service Level Objectives			
Name	Description	Achieving	Target
Mail Server Availability	Mail server is always available (ping check).	99.43%	99.0%
Mail Server Performance	Mail server consistently passes performance checks including CPU usage, swap space usage, and network retransmit counts.	90.03%	99.0%

**Enterprise Application SLA** Detailed Report

63% of compliance period  
66% of allowable downtime used  
98.99% of target (99.0%)

WARN - At the current rate, this SLA will breach after 10 more minutes of downtime.

Service Level Objectives			
Name	Description	Achieving	Target
Application Server	Availability and performance of WebSphere	99.49%	99.0%
Database	Availability and performance of Oracle	99.49%	99.0%
Web Server	Availability and performance of the web server	100.0%	99.0%

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

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

Service Level Objective			
Name	WebSphere		
Description			
Monitoring Period	Every Mon,Tue,Wed,Thu,Fri 9:00AM-6:00PM		
Compliance Period Type	Weekly		
Target Percentage	99.0		
Member Service Monitors			
Name	Element	Status	Description
WebSphere	WebSphere (lab-websphere51)	CRIT	
Plants Response	WebSphere (lab-websphere51)	CRIT	
File System Capacity	WebSphere (lab-websphere51)	CRIT	
PING-lab-websphere51	WebSphere (lab-websphere51)	OK	Default ping check for lab-websphere51

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

For more information about viewing SLA details, and defining SLOs that help you accurately gauge the performance of your IT infrastructure, see “Working with Service Level Agreements” on page 409.


## Viewing All Applications

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


An Application consists of:

- master service monitors  
One or more monitors can be used to determine the status of the Application as a whole.
- regular service monitors  
Other service monitors that are associated with a master service monitor, but are not used to determine the status of the Application as a whole.

The status of each Application is color coded:

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

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

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

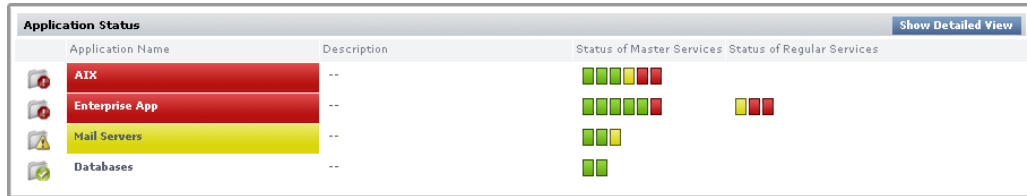
The **Applications** subpanel displays the status of each Application that you have added to [up.time](#).

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 Name	Description	Status of Master Services	Status of Regular Services
AIX	--	3 green, 1 yellow, 1 red	
Enterprise App	--	4 green, 1 yellow, 1 red	2 yellow, 1 red
Mail Servers	--	3 green	
Databases	--	2 green	

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:

The screenshot shows the 'Application Status' window with a 'Show Condensed View' button in the top right. The table below represents the data shown in the screenshot.

Application Name	Description	Element	Monitor Information
<b>AIX</b>	--		
* FSCap (member)	--	AIX QA LPAR	/usr is 53% full
* FSCap (member)	--	AIX DEV LPAR	/ is 97% full
* FSCap (member)	--	AIX5	/var is 100% full, /usr is 92% full, / is 90% full, /opt is 88% full
* PING-10.1.1.56	Default ping check for 10.1.1.56	AIX QA LPAR	Ping completed: 1 sent, 0.0% loss, 1.0ms average round trip time
* PING-10.1.1.57	Default ping check for 10.1.1.57	AIX DEV LPAR	Ping completed: 1 sent, 0.0% loss, 0.0ms average round trip time
* PING-aix5l	Default ping check for aix5l	AIX5	Ping completed: 1 sent, 0.0% loss, 1.0ms average round trip time
Application Name	Description	Element	Monitor Information
<b>Enterprise App</b>	--		
* Oracle Advanced	--	HP Integ	All requested values queried successfully
* PING-10.1.1.232	Default ping check for 10.1.1.232	lab-t1-2	Ping completed: 1 sent, 0.0% loss, 1.0ms average round trip time
* PING-10.1.1.232	Default ping check for hpointeg	HP Integ	Ping completed: 1 sent, 0.0% loss, 0.0ms

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 52.

## Viewing All Elements

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

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

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

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

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

- the percentage of the system disk that is being used and the percentage that is busy
- the amount of memory swap space that is being used

If [up.time](#) cannot contact an Element, then the following message is displayed:

```
The availability check has failed
```

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

- Click any value in the **OK, WARN, CRIT, MAINT, or UNKNOWN** columns to open the **Status** subpanel. See “Status” on page 54 for more information.
- Click any value in the **Outages** column to open the **Outages** subpanel. See “Outages” on page 55 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 547 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 574 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 570 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 (e.g., “up.time agent running on MailServer, up.time agent 3.7.2 linux”)

## Viewing the Resource Scan Report

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

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


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

### Performance Gauges

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

- CPU Usage

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



- Memory Usage

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

- Disk Busy


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

- Disk Capacity

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

## 24-Hour Performance Graphs

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

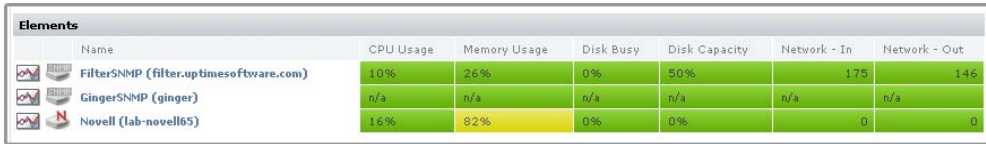





## Elements Chart

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

- CPU Usage  
The percentage of CPU resources that are being used.
- Memory Usage  
The amount of memory, expressed as a percentage of total available memory, that is being consumed by a process.
- Disk Capacity  
The percentage of storage space on the system disk that is being used.
- Network In  
The average amount of traffic coming in over the network interface.
- Network Out  
The average amount of traffic going out over the network interface.

The following image illustrates the Resource Scan chart:



Elements		CPU Usage	Memory Usage	Disk Busy	Disk Capacity	Network - In	Network - Out
	FilterSNMP (filter.uptimesoftware.com)	10%	26%	0%	50%	175	146
	GingerSNMP (ginger)	n/a	n/a	n/a	n/a	n/a	n/a
	Novell (lab-novell65)	16%	92%	0%	0%	0	0

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

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

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



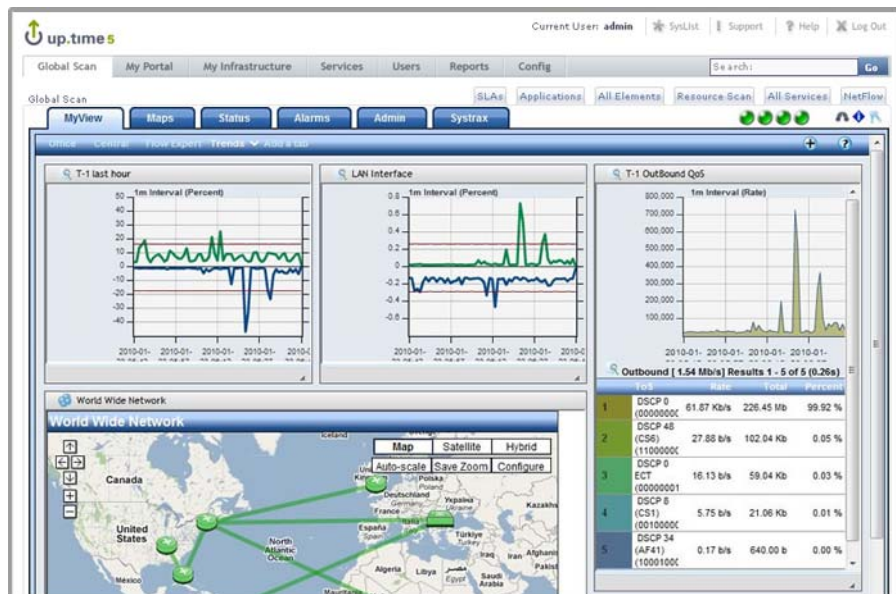


## Viewing Scrutinizer Status

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

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

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



## Changing Reporting Thresholds

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

**Global Scan** and the Resource Scan thresholds are configured by separate sets of attributes that can be changed in the **up.time Configuration** panel. By changing these attributes, you can set how large the color ranges are on resource gauges, and at what point table cells change color. See "Status Thresholds" on page 664 for more information.

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

# 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> .....	138
<i>Using Service Monitors</i> .....	139
<i>The Monitor Template</i> .....	144
<i>Cloning Service Monitors</i> .....	154
<i>Testing Service Monitors</i> .....	155
<i>Service Groups</i> .....	156
<i>Changing Host Checks</i> .....	161
<i>The Platform Performance Gatherer</i> .....	162
<i>Topological Dependencies</i> .....	164
<i>Scheduling Maintenance</i> .....	167


## 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 391 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

Agent-based service monitors require either the [up.time](#) Agent to be installed and running on the monitored system, or for Windows systems, metrics collection via WMI.

Agents or WMI 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
- Exchange 2003
- File System Capacity
- IIS
- Performance Check
- Process Count Check
- SQL Server (Advanced Metrics)
- Uptime Agent
- Windows Event Log Scanner
- Windows Service Check

Note that the [up.time](#) Agent service monitor specifically requires the [up.time](#) Agent, and cannot be used with systems whose metrics are collected via WMI.

## **Using Agentless Monitors**

Agentless monitors do not require the monitored system use WMI or the [up.time](#) Agent. 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
- DNS
- Custom Monitors
- Custom with Retained Data
- Email Delivery Monitor
- External Check
- FTP
- HTTP (Web Services)
- IMAP (Email Retrieval)
- LDAP
- Live Splunk Listener
- MySQL (Basic Checks)
- MySQL (Advanced Metrics)
- NFS
- NIS/YP
- NNTP (Network News)
- Oracle (Basic Checks)
- Oracle (Advanced Metrics)
- Oracle Tablespace Check
- Ping
- POP (Email Retrieval)
- SMTP (Email Delivery)
- SNMP
- Splunk Query
- SQL Server (Basic Checks)
- SQL Server Tablespace Check
- SSH (Secure Shell)
- Sybase
- TCP
- Web Application Transactions
- WebLogic
- WebSphere
- Windows File Shares (SMB)
- VMware Datacenter and Cluster Performance
- VMware ESX (Advanced Metrics)
- VMware ESX Server Power State
- VMware ESX Workload
- VMware Resource Pool and vApp Performance
- VMware VM Instance Power State
- VMware vSphere ESX Server Performance

## 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 373. 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 477) or a Service Metrics graph (see “Viewing System and Service Information” on page 52).

- 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 373.



## Selecting a Monitor

To select a monitor, do the following:

- 1 **Click Services on the up.time tool bar.**
- 2 **Click Add 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 144 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 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="radio"/> Service Group <input type="radio"/> Unassigned <input type="text" value="- Select a System -"/>

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 156.
- Click the **Unassigned** option.

This step is mandatory.

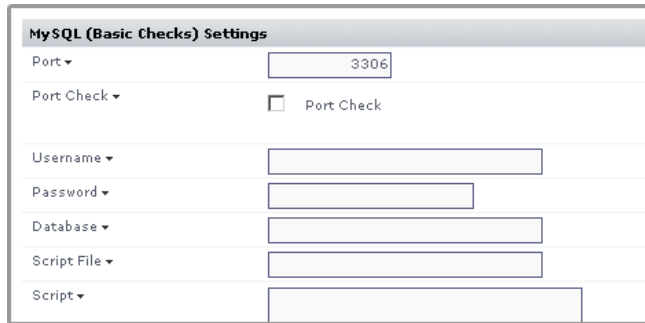
**4 Complete the following fields:**

- 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:

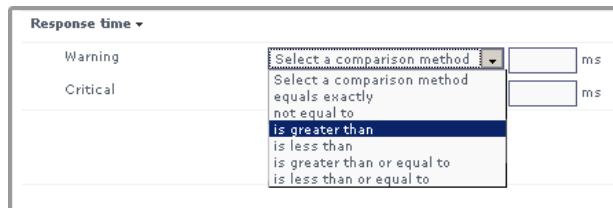


The screenshot shows a settings form titled "MySQL (Basic Checks) Settings". It contains several input fields and a checkbox:

- Port: 3306
- Port Check:  Port Check
- Username: [empty]
- Password: [empty]
- Database: [empty]
- Script File: [empty]
- Script: [empty]

### 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 method from the **Select a comparison method** dropdown list, as shown below:



The screenshot shows a "Response time" dropdown menu. The dropdown is open, showing a list of comparison methods. The "is greater than" option is highlighted. To the right of the dropdown are two input fields, each with "ms" next to it.

Warning	Select a comparison method	ms
Critical	Select a comparison method	ms

- Select a comparison method
- equals exactly
- not equal to
- is greater than
- is less than
- is greater than or equal to
- is less than or equal to

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 any patterns that do not correspond to the regular expression you define.

For example, if creating a service monitor for your Leech and Microsoft IIS FTP servers, you may want to ensure any message from them includes the FTP server name as part of the standard response. In this case, you can enter the following expression:

```
Leech|Microsoft
```

A missing name means a server may have been compromised or is not working correctly, in which case **up.time** would generate a critical alert.

- 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

The Response Time setting 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 window 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", a text input field, and a unit label "ms".

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

For information on configuring Warning thresholds, see “Configuring Warning and Critical Thresholds” on page 147.

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

For information on configuring Warning thresholds, see “Configuring Warning and Critical Thresholds” on page 147.



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

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 Elements, 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 Element. 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 defined Response Time.

- 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, see “Alerts and Actions” on page 429.

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 () beside the name of the service monitor.

A copy of the monitor template for the service monitor appears.

- 3 Enter information in the fields of the monitor template.

As a minimum, you must:

- enter a new name for the monitor in the **Service Name** field
- select a system to which you want to apply the monitor from the **Host** dropdown list

- 4 Click **Save**.


## 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.**

## 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 of 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 21.

## Creating Service Groups

To create a service group that can be applied to physical systems and network devices being monitored by [up.time](#), do the following:

- 1 On the [up.time](#) tool bar, click Services.**
- 2 In the Tree panel, click Add Service Group.**  
The **Add Service Group** window appears.
- 3 Acknowledge the creation of a “regular” service group by clicking Continue.**
- 4 On the second Add Service Group screen, enter a descriptive name for this group in the Name of Service Group field.**
- 5 Optionally, enter a description of the group in the Description field.**
- 6 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 From the Available Element Groups list, select one or more existing groups to immediately associate with the service group, then click Add.**

Select the **Include subgroups** check box to ensure any nested groups are also included. (For more information, see “Adding Nested Groups” on page 109.)

- 9 Select one of the following options from the Available Elements 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 Infrastructure** 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 108 for more information about grouping hosts.

- 10 Select one or more hosts from the list to immediately associate with the service group, then click Add.**
- 11 Click Finish.**

## Creating VMware vSphere Service Groups

To create a service group that will be used exclusively for VMware vCenter components, do the following:

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

The **Add Service Group** window appears.
- 3 Select the vSphere service group type, and click Continue.**

- 4 On the second Add Service Group screen, enter a descriptive name for this group in the Name of Service Group field.
- 5 Optionally, enter a description of the group in the Description field.
- 6 In the Select Master Services section, choose the existing service monitors that you would like to be in this vSphere service group.
- 7 In the Select vCenter Server section, choose the monitored VMware vCenter to which the service group will apply.
- 8 Using the subsequent sections for each type of VMware vCenter component, indicate how extensively and dynamically the service group will be applied:
  - none: the service group will not be applied to the indicated VMware vCenter component
  - any discovered: the service group will be unconditionally applied the VMware vCenter component type; this includes current existing components, as well as new ones that are detected through vSync
  - existing: the service group will only be applied to the existing datacenters, clusters, ESX hosts, resource pools, vApps, or VMs; it will not be applied to anything added to the [up.time](#) inventory via vSync after configuration


If you select existing VMware vCenter components, in the selection tool that appears, choose the components to which the service group will apply.
- 9 Click Finish.

## Editing Service Groups

To edit a service group used for physical infrastructure assets, 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** 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 Element Groups assigned to the group, do the following:
  - Add Element Groups by clicking on one or more entries in the **Available Element Groups** list, and then clicking **Add**.
  - Modify whether an Element Group's nested groups are included by selected or clearing the **Include subgroups** check box.
  - Remove systems by clicking on one or more entries in the **Selected Element Groups** list, and then clicking **Remove**.
  - Click **Save**.
- 7 To edit the Elements in the group, do the following:
  - Add systems by clicking on one or more systems in the **Available Elements** list, and then clicking **Add**.
  - Remove systems by clicking on one or more systems in the **Selected Elements** list, and then clicking **Remove**.
  - Click **Save**.

## Editing VMware vSphere Service Groups

To edit a service group used VMware vCenter components, 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 **Change the relevant service group details:**
  - In the **Service Group** section, enter a new name or description.
  - In the **Select vCenter Server** section, choose another VMware vCenter server whose components the service group will be applied.
  - In the relevant **Select Member** sections, modify how extensively or dynamically the service group will be applied to the VMware vCenter:
    - none: the service group will not be applied to the indicated VMware vCenter component
    - any discovered: the service group will be unconditionally applied the VMware vCenter component type; this includes current existing components, as well as new ones that are detected through vSync
    - existing: the service group will only be applied to the existing datacenters, clusters, ESX hosts, resource pools, vApps, or VMs; it will not be applied to anything added to the [up.time](#) inventory via vSync after configuration
    - If you select existing VMware vCenter components, in the selection tool that appears, choose the components to which the service group will apply.
- 5 **Click Finish.**

## 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 Infrastructure 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:**



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 Elements or Element 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, `up.time` will generate alerts for all the Elements behind the routers because the service monitors cannot communicate with them.

Topological dependencies help eliminate these kinds of unnecessary alerts by allowing administrators to create parent-child relationships between Elements. Both Elements and Element groups can be dependent on a parent system or node. With these relationships, topological dependencies work in two ways:

**Shared status:** If a topological parent is experiencing downtime, the child Elements in the topology will share the status (i.e., an Element's dependencies will automatically switch to its status). A service monitor will know that Elements 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 Elements and services that are children of that parent.

**Parent checks:** An outage with an Element will initiate a host check on its topological parent. By looking “upward,” [up.time](#) can find the root of the problem.



This parent host check behavior also applies to service monitor and host relationships, outside of topological dependencies.

## 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 Element 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

Normally, [up.time](#) will notify you that an Element or service is unavailable when systems or services are not online. During a maintenance period, the Monitoring Station assumes an Element cannot be contacted, thus will not generate any alerts for it.

Typically, maintenance is configured as a scheduled event, whether regular (e.g., a system back-up that occurs at a specific time each day or week), or planned (e.g., a system that will be taken offline on a Friday night for an upgrade).

Additionally, in cases where work needs to be done on an Element outside of a pre-defined, scheduled period, maintenance status can be assigned to an Element ad hoc in the **My Infrastructure** panel. An Element group can also be put into temporary maintenance mode, affecting its Elements and subgroups. Any Elements added to, or removed from, the Element group during the temporary maintenance period will inherit the appropriate state.

## 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 Add Maintenance Profiles.**
- 3 **Enter a descriptive name for the profile in the Profile Name field.**
- 4 **Enter time period expressions in the Definition field that together make up the maintenance window.**

See “Time Period Definitions” on page 677 for information on the types of time period expressions that are valid in [up.time](#).

- 5 **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**.

## Putting an Element or Group into Temporary Maintenance Mode

To put an Element or Element group into temporary maintenance mode, do the following:

- 1 On the **up.time** tool bar, click **My Infrastructure**.
- 2 Locate the **Element** or **Element group** whose status is to be temporarily changed to **MAINT**.
- 3 Click the **Element** or **group's gear icon**.
- 4 In the pop-up menu, click **Put into Temporary Maintenance**.

The **Element** or **group's** status is immediately reflected in **Global Scan** with an **In Maintenance** status icon.

When work is complete, restore the **Element** or **group** status by using the **Take Out of Temporary Maintenance** command in **My Infrastructure**.



# CHAPTER 9

## Monitoring VMware vSphere

---

This chapter covers monitoring your VMware vSphere virtual infrastructure in [up.time](#), including concepts and configuration, in the following sections:

<i>Overview</i> .....	172
<i>VMware vSphere Monitoring Concepts</i> .....	173
<i>Managing vSync</i> .....	178
<i>Standalone Monitoring for vCenter VMs</i> .....	184
<i>vCenter Configuration Mirroring</i> .....	188
<i>vSphere Element Configurations</i> .....	192

### Overview

If your organization has a virtual infrastructure that is managed by VMware vCenter, you can use [up.time](#) to monitor, diagnose, and report on your virtual machines and management servers. From an administrative perspective, seamlessly integrating virtual and physical assets in [up.time](#) has several key benefits:

- you can save time by working with one dashboard and management tool for monitoring, alerting, and alert acknowledgement
- you can generate diagnosis graphs from the same monitoring tool, allowing you to view both virtual and physical machines in the same views
- in addition to [up.time](#)'s dedicated VMware vSphere reports, you can extend [up.time](#)'s general resource-focused reports to VMware vCenter servers, ESX servers, and VMs

## VMware vSphere Monitoring Concepts

There are three main conceptual components to integrated virtual and physical monitoring:

- metrics for VMware vSphere components, mainly pertaining to resource usage, but also including power states
- **up.time**'s vSync feature, the syncing engine that ensures VMware vSphere components and inventories are accurately mirrored in **up.time**, as well as their basic, non-agent-based metric data
- vSphere components that are monitored and managed as **up.time** Elements; these include VMware vCenter servers, ESX servers, and VMs

In addition to understanding these concepts, it is important that you also understand the way other **up.time** functions change in a VMware vSphere monitoring context.

### About up.time vSync

The core of integrated monitoring is **up.time**'s vSync, whose key functions are replicating metrics and mirroring topologies, from VMware vSphere to **up.time**.

vSync takes performance metrics gathered by VMware vSphere, normally for use in VMware tools such as the vSphere Client, and represents them in **up.time**. VMware vSphere metrics that are available in **up.time** include performance data for VMs, the ESX servers that host them, and the VMware vCenter servers that manage your configurations (including but not limited to datacenters, clusters, and resource pools).

vSync also regularly monitors your VMware vSphere datacenter's dynamic environment (including both physical and virtual assets), ensuring the VMware vSphere inventory that **up.time** is using for monitoring and reporting is always current. For more information, see "Managing vSync" on page 178.

## vSphere Components as up.time Elements

Some VMware vSphere components that act as logical groupings, including datacenters, clusters, resource pools, and vApps, are hierarchically mirrored in [up.time](#)'s **My Infrastructure** inventory, and are represented through their reported resource metrics. On the other hand, VMware vSphere components that are actual hosts, whether virtual or physical (i.e., a VMware vCenter server, its component ESX servers, or their respective VMs) are represented in [up.time](#) as Elements.

When you add a VMware vCenter server as an Element, all of the VMware vSphere components, whether organizational or compute resources, as defined through the vSphere Client, are imported to [up.time](#), and their VMware-vSphere-collected metric data is migrated to the [up.time](#) DataStore. Additionally, all ESX servers and VMs also become Elements.

There will be cases where administrators will want to actively manage which ESX server Elements (and subsequently, which VMs) are monitored by [up.time](#). For example, licensing constraints may prevent you from monitoring every ESX server in [up.time](#), or the performance of particular portions of your virtual infrastructure may not be considered mission critical, and demand the same level of uptime.

In these cases, administrators can include or exclude specific [up.time](#) Elements from being monitored. Exclusions can be made on a per-VM basis, or by a logical grouping at the VMware vSphere level (e.g., by cluster).

Any VMware vCenter component is by default represented in **My Infrastructure** as either a known host or a known VM, and is grouped as such in the inventory. (By default, the **My Inventory** group names are **Discovered Hosts** and **Discovered VM Hosts**). When Elements are ignored, they are removed from **My Infrastructure**. If new hosts or VMs are discovered during vSync, they are added to the appropriate **My Infrastructure** group.

For more information, see “Managing vCenter Inventories in up.time” on page 178.

## vSphere Components and Topological Dependencies

When a VMware vCenter server is added to [up.time](#), the hierarchical structure of its components is retained (and used in **My Infrastructure**).



**up.time** observes particular monitoring rules to these existing VMware vSphere topologies that supplement the behavior of physical topologies that have been defined in **up.time**. (See “Topological Dependencies” on page 164 for more information.)

When you define a physical topology or VMware vSphere topology in a Topological Dependency, the following behaviors are commonly observed:

- if a topological parent is experiencing downtime, the child Elements in the topology will share the status (i.e., an Element's dependencies will automatically switch to its status)
- an outage with an Element, whether actual or topological, will initiate a host check on its parent (e.g., a service monitor and its host, or a host and its topological parent)

However, there are behaviors unique to Topological Dependencies based on a VMware vSphere topology:

- as VMs migrate, their links to their ESX hosts is maintained
- **up.time** will be aware of power states in the virtual infrastructure, such that parent Elements that are powered down will not spawn alerts with its child Elements (e.g., all of a VMware vCenter servers many ESX servers and VMs)

## vSphere Components and Service Groups

In **up.time**, vSphere service groups exist to allow you to group similar Elements to perform common host checks, not unlike groupings created in the vSphere Client (e.g., a cluster or datacenter). The key difference with vSphere service groups from regular ones is once established, vSphere service groups will be managed by the vSync process. Any changes detected with the VMware vCenter server's topology will automatically be reflected in the **up.time** service group.

For information on creating a service group, see “Creating Service Groups” on page 156.

## Reports and Graphs for vCenter Server Metrics

As one of vSync's main benefits is to enable you to combine virtual and physical monitoring, the following [up.time](#) tools are available assist with VMware vSphere analysis and reporting.

Service Monitors consist of the following:

- Datacenter and Cluster Performance
- Resource Pool and vApp Performance
- vSphere ESX Server Performance
- ESX (Advanced Metrics)
- ESX Server Power State
- VM Instance Power State

Diagnosis graphs consist of the following:

- CPU Workload
- CPU Usage
- Multi-CPU Usage
- Wait and Ready Time
- Memory Workload
- Memory Usage
- Memory Profile
- Memory: Guest Paging
- Memory Swap
- Process Workload
- Detailed Process Information
- Number of Processes
- Network Workload
- Network I/O
- Network Errors
- Disk Workload

- Disk I/O
- Disk Latency
- Disk Errors
- Disk Storage Capacity
- Disk Storage Profile
- Power Consumption
- VM Instance Motion
- Power States

Reports consist of the following:

- vSphere Workload Report
- VM Sprawl Report
- VMware Workload Report
- VMware Infrastructure Density Report

## Managing vSync

To take full advantage of integrated virtual and physical monitoring, it is important to understand how to manage VMware vSphere inventories in **up.time**, and how vSync keeps **up.time** abreast of them.

When a VMware vCenter server is added as an Element to **up.time**, its component ESX servers and VMs are also added as Elements. It is at the vCenter Element level that you can tweak what is monitored. Additionally, you can “promote” VM Elements to be monitored with greater depth using the **up.time** agent, or WMI.

The dynamic nature of a VMware vCenter server’s configuration is easily managed by **up.time** through automated profile monitoring; whatever you have included in your inventory as parent-level items, **up.time** vSync will maintain an accurate mirror. See “vCenter Configuration Mirroring” on page 188 for more information.

## Managing vCenter Inventories in up.time

After a VMware vCenter server has been added as an Element to **up.time**, you can manually manage all of its inventory. By default, all of a VMware vCenter’s components are monitored by **up.time**. As a central monitoring configuration point, the VMware vCenter Element gives you the ability to determine which VMs are monitored in **up.time** and as a consequence, which VMware vSphere metrics are kept in the DataStore.

You can select VMs to be “ignored” by **up.time**. When VMs are ignored, their topological relationship to other VMware vCenter components is retained, but their metrics are no longer recorded to the DataStore.

As the topology of your virtual infrastructure changes, you can make changes to the ignore list. Historical data for VM Elements are retained in the DataStore, regardless if they are currently being monitored, allowing administrators to continue to perform diagnosis on a continuously changing virtual environment.



If graphs or reports are generated that include time or date ranges where no data were collected for a VM Element, those ranges in the charts will show no data.

Although there are no resource-related reasons why you would not want [up.time](#) to monitor all ESX servers in your virtual infrastructure, other factors may dictate your monitoring policies:

- as each ESX server counts toward your [up.time](#) license, you may have resource monitoring restrictions that require prioritization
- you are not interested in keeping performance metrics for some parts of your virtual infrastructure (e.g., some VMs are being used for testing, and are not considered mission-critical)
- similarly, as ignored VMs do not appear in live monitoring pages in [up.time](#) such as **Global Scan** or the **Resource Scan** panel, you may want to exclude non-essential VMs to avoid cluttering them

On the VMware vCenter Element's **Info** page, you can view its complete inventory, broken down by logical groupings created in the vSphere Client (e.g., by datacenter). It is here you can add or remove VMs (as individual machines, or in groups at the datacenter or cluster level) from the ignore list.

## Viewing vCenter Inventory Summaries

To see an overview of your VMware vCenter Element's inventory, do the following:

- 1 In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**
- 2 On the VMware vCenter Element profile page, ensure the Info tab is selected.**
- 3 Click Inventory Summary.**

On this page, you will see your VMware vCenter server's total available resource counts, further broken down by datacenter.

## Removing Elements from Monitored Inventories

To stop VMware vCenter server Elements from being monitored by [up.time](#), do the following:

- 1 In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**

- 2 **On the VMware vCenter Element profile page, ensure the Info tab is selected.**
- 3 **Click Inventory Detail.**
- 4 **In the vCenter Inventory Detail page, select the checkboxes that correspond to the individual VMs or clusters that you want to stop monitoring.**
- 5 **Click Add Selected Elements to Ignore.**

The inventory page will refresh, with your selections appearing in the **Ignored Elements** section.

### Re-adding Ignored Elements to Monitored Inventories

To resume monitoring VMware vCenter server Elements in [up.time](#), do the following:

- 1 **In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**
- 2 **On the VMware vCenter Element profile page, ensure the Info tab is selected.**
- 3 **Click Inventory Detail.**
- 4 **In the vCenter Inventory Detail page, from the Ignored Elements list, select the checkboxes that correspond to the individual VMs or clusters that you want to start monitoring.**



Note that to streamline the process, you can only select or remove VMs or datacenters as they were added. For example, if you added an entire datacenter to the ignore list, you cannot select individual VMs.

- 5 **Click Remove Selected Elements from Ignore List.**

The inventory page will refresh, with your selections appearing in the active monitoring section.

## vCenter vSync Settings

From the VMware vCenter server Element's **Info** page, you can configure how [up.time](#) polls the server's topology for changes, and what [up.time](#) does when changes are encountered.

Using alert and action profiles, you can configure an extensive set of behaviors depending on what changes occur with the VMware vCenter topology. This means you can bundle VMware vSphere events with the those you normally track in your physical infrastructure.

For example, if your VMware vSphere environment is configured to dynamically allocate more compute resources to a specific department at a specific time of year (e.g., the end of the fiscal year for the accounting department), you can use [up.time](#) to detect CPU changes to the appropriate resource pool, and confirm this action in a log.

See “vCenter Configuration Mirroring” on page 188 for more information on the types of VMware vSphere events that [up.time](#) will watch for.

### vSync Configuration Settings

The following settings are used to configure a VMware vCenter server's vSync behavior:

Alert Profiles	A list of alert profiles (listed in alphabetical order) that are triggered when new VMs or ESX hosts belonging to the VMware vCenter server are discovered when vSync scans the VMware topology.
Action Profiles	A list of action profiles (listed in alphabetical order) that are triggered when new VMs or ESX hosts belonging to the VMware vCenter server are discovered when vSync scans the VMware topology.

Default Basic Performance Data Polling Interval	<p>The amount of time, in minutes, that <a href="#">up.time</a> waits between metric data retrieval operations for newly discovered ESX servers.</p> <p>Select the <b>Apply to Existing ESX Servers</b> check box to have this modified polling interval applied to all ESX servers that are part of the VMware vCenter server's topology, and the <a href="#">up.time</a> inventory. Any ESX server whose polling rates have been individually changed will conform to this new global setting.</p>
Default Advanced Performance Data Polling Interval	<p>The advanced data settings perform the same functions as those above, but are for advanced VMware vSphere metrics.</p>

## Changing vSync Settings

To modify how vSync mirrors a VMware vCenter server Element's configuration in [up.time](#), do the following:

- 1 **In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**
- 2 **On the VMware vCenter Element profile page, ensure the Info tab is selected.**
- 3 **In the vSync section, click Edit Settings.**
- 4 **Make changes to vSync settings as needed:**
  - modify the alert profile or action profile lists for newly discovered VMs in the specified My Infrastructure group
  - modify the alert profile or action profile lists for newly discovered ESX servers in the specified My Infrastructure group



- modify the ESX server polling intervals, and indicate whether this change is to be retroactively applied to all known ESX servers



The polling intervals modified here will apply only to newly discovered ESX servers via vSync. Whether individual configured, or from a previous global setting, existing ESX server Elements in [up.time](#) will use older intervals, unless you use the **Apply to Existing ESX Servers** option.

**5 Click Save.**

## Manually Resyncing up.time with a vCenter Topology

To force [up.time](#) to resync with a VMware vCenter server (i.e., search for new VMs and ESX servers, metric data for monitored Elements, and changes to the topology), do the following:

- 1 In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**
- 2 On the VMware vCenter Element profile page, ensure the Info tab is selected.**
- 3 In the vSync section, click Resync Now.**

A **Resyncing** status update will persist until [up.time](#) has completed the vSync process. Once the sync operation has been completed, the polling intervals configured for this VMware vCenter server will be reset.

## Standalone Monitoring for vCenter VMs

To facilitate simpler virtual asset management, VMs that belong to a VMware vCenter server topology are always added to [up.time](#) when that vCenter is added as an Element. The metrics used to monitor these VMs are collected by VMware vSphere. However, there may be cases where you would like an Action Profile to perform self-healing steps on a mission-critical VM.

In these scenarios, you can, from [up.time](#)'s configuration perspective, “detach” a VM from the VMware vCenter hierarchy that was originally imported, in order to control it independently of the vCenter server. Once performed, this VM is, in a sense, promoted to receive additional performance data delivered by the [up.time](#) Agent, or for Windows systems, via WMI.

Although the topology remains intact on the VMware vSphere side, from [up.time](#)'s perspective, the VM is a standalone Element.

If the VMware vCenter Server is configured for additional data collection, any VM that is part of a VMware vCenter inventory is automatically “promoted” to be independently monitored when either the [up.time](#) agent has been installed on it, or it is configured to use WMI for metrics retrieval.

This behavior is controlled by the additional data collection settings for the vCenter, which were first configured (or skipped) when the VMware vCenter server was first added to [up.time](#). These additional data collection settings can be modified at any time, after which [up.time](#) will search for appropriate VMs.

The frequency at which [up.time](#) checks for agent-enabled VMs is based on the **Rescan Time Interval** setting, which by default, is every 180 minutes.

### Configuring Standalone VM Candidacy Scanning

To configure the method with which [up.time](#) scans VMware vCenter server VMs for eligibility for standalone monitoring, do the following:

- 1 In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**
- 2 On the VMware vCenter Element profile page, ensure the Info tab is selected.**

- 3 On the main profile page, locate the **Additional Agent Monitoring** section, then click **Edit Settings**.
- 4 If desired, modify the **Rescan Time Interval** setting.



When you change this interval, the previous one will first be completed before adjusting to the new one.

- 5 **Configure the data collection details with which candidate VMs will be configured to collect metrics:**
  - i For the **up.time** agent, indicate the port on which it is listening, and whether it is securely communicating with **up.time** using **SSL**.
  - i For data collection via **WMI**, indicate the host and domain on which **WMI** has been implemented, and the username and password required for access.



If a VM is detected to have both the **up.time** agent as well as **WMI** credentials, **up.time** will always mark the VM for standalone monitoring via the agent.

- 2 **Click Save.**

## Modifying Data Collection Details for an Existing Standalone VM

For individual VMs that were scanned and discovered using VMware vCenter standalone VM settings, you can make changes to their data collection method by doing the following:

- 1 In the **My Infrastructure** panel, click the gear icon beside the **WMI-** or **agent-based VM** you wish to modify.
- 2 In the pop-menu, select **View**.
- 3 On the VM's profile page, ensure the **Info** tab is selected.
- 4 On the main profile page, in the **Data Collection** section, click **Edit Collection Method**.
- 5 In the pop-up configuration window, modify the details for the existing data collection method for the VM:

- For the up.time agent, indicate the port on which it is listening, and whether it is securely communicating with up.time using SSL.
  - For data collection via WMI, indicate the host and domain on which WMI has been implemented, and the username and password required for access.
- 6 **Click Save.**
  - 7 **When the pop-up window closes, you can make sure the settings you have provided are correct by clicking the Test button.**

## Detaching Ignored VMs

Any VM whose parent VMware vSphere object has been marked to be ignored cannot individually be added to the monitored list. (See “Managing vCenter Inventories in up.time” on page 178 for more information.) In these cases, any ignored VM that is a candidate for standalone monitoring will be shown with an accompanying **Enable Monitoring** button, found on the VMware vCenter server’s **Info** page:



## Enabling Standalone Monitoring on Ignored VMs

To remove an ignored VM from up.time’s mirrored inventory of the VMware vCenter server, do the following:

- 1 **Ensure the up.time agent has been installed on the VM Element, or that up.time is configured to use WMI as a data collection method for Windows-based Elements.**

For more information, see “Configuring Global Data Collection Methods” on page 646.

- 2 **In the My Infrastructure panel, click the gear icon beside the VMware vCenter server, then in the pop-menu, select View.**
- 3 **On the VMware vCenter Element profile page, ensure the Info tab is selected.**
- 4 **Click Inventory Detail.**

- 5 **On the vCenter Inventory Detail page, from the Ignored Elements list, locate the VM Element you wish to enable for standalone monitoring.**

All VMs that are part of the VMware vCenter server's inventory that are found to be eligible for standalone monitoring will be listed with a corresponding **Enable Monitoring** button.

- 6 **Click Enable Monitoring.**



Enabling standalone monitoring for a VM permanently detaches it from the VMware vCenter server's configuration in [up.time](#).

- 7 **Repeat these steps for each VM.**

All VMs that have been detached for standalone monitoring will no longer appear in [up.time](#) as part of the VMware vCenter server's inventory, or as a virtual machine, and will instead appear and report as a physical Element.

## vCenter Configuration Mirroring

In order for [up.time](#) to successfully integrate virtual and physical monitoring, it needs to be completely in step with a highly dynamic virtual infrastructure. For administrators, understanding the level at which [up.time](#) mirrors changes helps you figure out what you can integrate into an alerting policy, and how to avoid false positives.

### vSync Change Events

The following VMware vCenter event types are detectable by vSync, along with the action performed in [up.time](#):

Datacenters	
datacenter found	<ul style="list-style-type: none"> <li>• automatically added to any vSphere service groups that have a master datacenter service that unconditionally takes newly discovered datacenters as child services</li> <li>• association is established with a VMware vCenter server</li> </ul>
datacenter deleted	all references to the datacenter are removed from up.time
ESX Servers	
host found	<ul style="list-style-type: none"> <li>• automatically added to any vSphere service groups that have a master ESX host service that unconditionally takes newly discovered ESX hosts as child services</li> <li>• association is established with a VMware vCenter server</li> <li>• automatically moved to the list of ignored Elements for a VMware vCenter server if its addition will violate your license</li> </ul>

host deleted	<ul style="list-style-type: none"> <li>• monitoring is disabled, and automatically switched to list of ignored Elements for a VMware vCenter server</li> <li>• all references to the ESX server in up.time are removed</li> </ul>
power state change	all regular procedures associated with power state changes are performed
host added to or removed from cluster	VMware vCenter server topology is updated
configuration change	configuration details reflected in host's up.time profile unless it is on the list of ignored Elements
<b>VMs</b>	
VM licensing change	VMs do not affect up.time licensing if they are managed by an ESX server that is being monitored by up.time
VM found	<ul style="list-style-type: none"> <li>• automatically added to any vSphere service groups that have a master VM service that unconditionally takes newly discovered VMs as child services</li> <li>• association is established with an ESX server</li> </ul>
VM deleted	<ul style="list-style-type: none"> <li>• monitoring is disabled, and automatically switched to list of ignored Elements for a VMware vCenter server</li> <li>• all references to the VM in up.time are removed</li> </ul>
power state change	all regular procedures associated with power state changes are performed
VM vMotion	automatically update topological dependencies to reflect the new host for the VM, and the previous one
configuration change	configuration details reflected in VM's up.time profile unless it is on the list of ignored Elements

<b>vApps</b>	
vApp found	<ul style="list-style-type: none"> <li>• automatically added to any vSphere service groups that have a master vApp that unconditionally takes newly discovered vApps as child services</li> <li>• association is established with a VMware vCenter server</li> </ul>
vApp deleted	<ul style="list-style-type: none"> <li>• automatically removed from the hierarchical listing in <b>My Infrastructure</b></li> <li>• all service monitor references to the VMware vCenter are removed</li> </ul>
power state change	all regular procedures associated with power state changes are performed
configuration change	configuration details reflected in vApp's up.time profile
<b>Clusters</b>	
cluster found	<ul style="list-style-type: none"> <li>• automatically added to any vSphere service groups that have a master cluster service that unconditionally takes newly discovered clusters as child services</li> <li>• association is established with a VMware vCenter server</li> <li>• automatically moved to the list of ignored Elements for a VMware vCenter server if its addition will violate your license</li> </ul>
cluster deleted	<ul style="list-style-type: none"> <li>• automatically removed from the hierarchical listing in <b>My Infrastructure</b></li> <li>• all service monitor references to the VMware vCenter are removed</li> </ul>
configuration change	configuration details reflected in cluster's up.time profile



Resource Pools	
resource pool found	<ul style="list-style-type: none"> <li>• automatically added to any vSphere service groups that have a master resource pool that unconditionally takes newly discovered resource pools as child services</li> <li>• association is established with a VMware vCenter server</li> </ul>
resource pool deleted	<ul style="list-style-type: none"> <li>• automatically removed from the hierarchical listing in <b>My Infrastructure</b></li> <li>• all service monitor references to the VMware vCenter are removed</li> </ul>
configuration change	configuration details reflected in resource pool's up.time profile

## vSphere Element Configurations

Outside of vSync-related management, you may want to change the way vSphere Elements (i.e., VMware vCenter servers, ESX servers, and VMs) are displayed in [up.time](#), as well as test that their data collection is functioning correctly.

### vCenter Element Profiles

The **System Profile** section on a VMware vCenter server’s **Info** page allows you to modify its descriptive information in [up.time](#), change how [up.time](#) communicates with it, and how often it is polled for metric data.

#### vCenter Server Profile Settings

The following settings define a VMware vCenter server profile in [up.time](#):

Display Name	The name used for the VMware vCenter server Element throughout the up.time interface (e.g., <b>My Infrastructure</b> ).
Host Name	The actual host name of the server.
Description	An optional description for the server that is used in different parts of the up.time interface (e.g., in tool tips in <b>Global Scan</b> and <b>My Infrastructure</b> )
Monitored Status	A switch that enables and disables monitoring for the server in up.time. When disabled, the Element will no longer appear in <b>Global Scan</b> , and its child ESX servers and their respective VMs will also cease data collection.
Storage Polling Interval	The frequency that the VMware vCenter server will collect storage-related metric data.
Performance Data Polling Interval	The frequency that the VMware vCenter server will collect performance data for all of its components (datacenters, clusters, resource pools, vApps).

Username	The VMware administrator username used to access the VMware vCenter server.
Password	The password required for administrator access.
Port	The port through which up.time communicates with the VMware vCenter server.
Connection Status	A standard up.time status indicator for the VMware vCenter server.

## Viewing and Editing a vCenter Server Profile

To modify the display details for a VMware vCenter server Element in [up.time](#), or modify the rate at which it gathers metric data for its component resources, do the following:

- 1 In the **My Infrastructure** panel, click the gear icon beside the **VMware vCenter** server, then in the pop-menu, select **View**.
- 2 On the **VMware vCenter Element** profile page, click **Edit System Profile**.
- 3 In the **Edit System Profile** pop-up window, make all required changes.
- 4 Click **Save**.

## ESX Server Element Profiles

The **Info** tab for an ESX server allows you to modify its descriptive information in [up.time](#), test its data collection status, and view its VM inventory.

For each ESX server, [up.time](#) also performs regular configuration checks to ensure all aspects of the hypervisor’s managed virtual assets, as well as its own hardware details are up to date. By default, this information is polled by the vSphere ESX Server Configuration Gatherer service monitor, and includes the following configuration details:

- virtual and network datastore profiles and capacities
- NICs and virtual networks

- physical CPU, disk, and file system information
- for managed VMs, vNICs and virtual disks

## ESX Server Profile Details and Settings

The following details and settings can be viewed or configured for an ESX server profile in [up.time](#):

Display Name	The name used for the ESX server throughout the up.time interface (e.g., <b>My Infrastructure</b> ).
Host Name	The actual host name of the server.
Description	An optional description for the server that is used in different parts of the up.time interface (e.g., in tool tips in <b>Global Scan</b> and <b>My Infrastructure</b> ).
Parent Group	The group the ESX server is assigned to, for organization in <b>My Infrastructure</b> .
vSphere Parent Object	The VMware vSphere component (e.g., cluster) to which the ESX server belongs. If the user has access rights, this object will link to that Element's <b>Quick Snapshot</b> page.
vCenter Server	The VMware vCenter server that manages this ESX server. If the user has access rights, this will link to that Element's <b>Quick Snapshot</b> page.
Monitored Status	A switch that enables and disables monitoring for the server in up.time. When disabled, the Element will no longer appear in <b>Global Scan</b> , and its VMs will cease data collection.
Basic Performance Data Polling Interval	The amount of time, in minutes, that up.time waits between metric data retrieval operations.
Advanced Performance Data Polling Interval	The advanced data settings perform the same functions as those above, but are for advanced VMware vSphere metrics.

## Viewing and Editing an ESX Server Profile

To modify the display details for an ESX server Element in [up.time](#), or modify the rate at which it gathers metric data for its component resources, do the following:

- 1 In the **My Infrastructure** panel, click the **ESX** server name.
- 2 On the **ESX** server profile page, click **Edit System Profile**.
- 3 In the **Edit System Profile** pop-up window, make all required changes.
- 4 Click **Save**.

## Testing and Modifying Data Collection for an ESX Server

To ensure [up.time](#) can successfully draw metric data collected by the VMware vSphere platform for an ESX server, do the following:

- 1 In the **My Infrastructure** panel, click the **ESX** server name.
- 2 On the **ESX** server profile page, in the **Data Collection** section, click **Test for Basic Platform Performance Status**.
- 3 For the **Advanced Platform Performance Status**, click **Test**.
- 4 To make any changes, click **View Settings**, then click **Edit Info**.

See “Editing the Platform Performance Gatherer” on page 162 for more information on making changes to data polling for an Element.

## VM Element Profiles

The **System Profile** section on a VMware vCenter server’s **Info** page allows you to modify its descriptive information in [up.time](#), change how [up.time](#) communicates with it, and how often it is polled for metric data.

## VM Profile Details and Settings

The following details and settings can be viewed or configured for an ESX server profile in [up.time](#):

Display Name	The name used for the VM throughout the up.time interface (e.g., <b>My Infrastructure</b> ).
Host Name	The actual host name of the VM.
Description	An optional description for the VM that is used in different parts of the up.time interface (e.g., in tool tips in <b>Global Scan</b> and <b>My Infrastructure</b> )
Parent Group	The group the VM is assigned to, for organization in <b>My Infrastructure</b> .
Monitored Status	A switch that enables and disables monitoring for the VM in up.time. When disabled, the VM does not appear in <b>Global Scan</b> , and no data is collected for it.
Basic Performance Data Polling Interval	The amount of time, in minutes, that up.time waits between metric data retrieval operations for the VM. This option appears only if the VM is receiving metric data via WMI or the <a href="#">up.time</a> agent.
Number of Processes to Retrieve	The number of VM processes, based on the most usage, those metrics are collected for the VM. This option appears only if the VM is receiving metric data via WMI or the <a href="#">up.time</a> agent.

## Viewing and Editing a VM Profile

To modify the display details for a VM Element in [up.time](#), do the following:

- 1 **In the My Infrastructure panel, click the VM name.**
- 2 **On the VM profile page, click Edit System Profile.**
- 3 **In the Edit System Profile pop-up window, make all required changes.**

#### 4 Click Save.

If the VM receives its metric data via WMI or the [up.time](#) agent, you can also modify the collection method if desired. See “Modifying Data Collection Details for an Existing Standalone VM” on page 185 for more information.





# CHAPTER 10

## Agent Monitors

---

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

<i>File System Capacity</i> .....	201
<i>Performance Check</i> .....	204
<i>Process Count Check</i> .....	208

## 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 42.

## 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 all local drives; on UNIX and Linux servers, [up.time](#) looks at all local file systems (e.g., /var, /export, /usr).

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.

Windows Volume Mount Points can be monitored when the host Element is monitored through WMI, not the [up.time](#) agent (see “Working with Systems” on page 69 for more information). Note that the level of detail for mounted volumes on Windows XP and 2000, when reported through WMI, is limited: the mounted volume name and exact location are not always accurate, but other pertinent information, such as volume capacity and usage, are correct.



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 144.

### 2 Complete the following fields:

- Global Warning Threshold (Mandatory)

Create a threshold that generates a warning. This threshold can be an actual amount (in MB, GB, or TB), or percentage of disk space that is used or is free.

- Global Critical Threshold (Mandatory)

Create a threshold that generates a critical alert, whether it is an actual amount, or percentage of disk space used or free.

- 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`.  
Case sensitivity is not taken into account when monitor-defined mount points are matched with those on the file system.
- The **Warning** threshold, which is a percentage or amount of space used or free on the mount point that, when exceeded, generates a warning.
- The **Critical** threshold, which is a percentage or actual amount of space used or free on the mount point that, when exceeded, generates a critical alert.

The thresholds that you set for each mount point will be calculated separately from the thresholds that you specified in step 2.

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

For more information, see “Response Time” on page 148.

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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 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), averaged over the number of seconds that you specify, that is being used
- 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 144.

**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 formatting the packets for transmission, or transmitting the packets.

- 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 151 for more information).
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).



- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 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, file extension, or any parameters.

For example, on UNIX systems, the process “`/usr/bin/vmstat -p`” is checked as “`vmstat`”, and on Windows systems, “`process.exe`” should be entered as “`process`”.

- 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 147.

- Response Time

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 3.**
- 4 **Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 **Click Finish.**



# CHAPTER 11

## VMware Monitors

---

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

<i>Datacenter and Cluster Performance</i> .....	213
<i>Resource Pool and vApp Performance</i> .....	215
<i>vSphere ESX Server Performance</i> .....	218
<i>ESX (Advanced Metrics)</i> .....	222
<i>ESX Workload</i> .....	225
<i>ESX Server Power State</i> .....	229
<i>VM Instance Power State</i> .....	231

## Overview

The VMware service monitors allow you to monitor and alert based on the performance and status your virtual resources. These monitors can watch for threshold violations with computing resources for VMs, ESX servers, and changes in power states.

Most of these service monitors use VMware vCenter-collected metrics that are made available to up.time through vSync. These VMware vCenter-monitored components, combined with more granular agent-based server monitoring, provide you with choice between breadth and depth.

## vSphere Performance Monitors

The vSphere performance monitors allow you to monitor and alert on specific VMware vSphere components: datacenters, clusters, resource pools, and vApps.

The metrics collected through VMware vCenter servers can be used by [up.time](#) through vSync, and subsequently be used to trigger [up.time](#)'s own alerts and actions, allowing you to integrate both your vSphere-managed and non-virtual resources.

These performance monitors can answer questions such as the following:

- Is the CPU usage of VMs in a vApp, resource pool, cluster, or datacenter passing an acceptable level?
- Is the memory consumed by VMs in a vApp, resource pool, cluster, or datacenter passing an acceptable level?
- Is the number of ESX servers that are a part of the cluster or datacenter exceeding an acceptable number and threatening performance?

### Datacenter and Cluster Performance

The Datacenter Performance and Cluster Performance monitors can trigger alerts on metrics collected through vSync.

#### Datacenter Performance and Cluster Performance Monitor Metrics

The following VMware vSphere metric types for datacenter or cluster performance can be used to configure thresholds in [up.time](#):

Time Interval	A positive integer indicating the number of minutes' worth of performance data samples to average, then compare against threshold definitions (default: 30)
---------------	---

Number of Running VMs: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the average number of VMs powered on during the time interval.
Number of Running Hosts: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the average number of vSphere ESX servers powered on during the time interval.
CPU Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of CPU cycles consumed by VMs belonging to this datacenter or cluster.
Memory Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of memory consumed by VMs belonging to this datacenter or cluster.

## Configuring Datacenter Performance or Cluster Performance Monitors

To configure a Datacenter Performance or Cluster Performance monitor, do the following:

- 1 Select the monitor from the Add Service Monitor window, in the VMware Monitors section.**
- 2 Click Continue to begin configuring the service monitor.**
- 3 Complete the monitor information fields.**

See “Monitor Identification” on page 144 for more information on configuring service monitor information fields.

- 4 In the Cluster Performance Settings section, configure the monitor’s warning- and critical-level threshold values:**
  - Time Interval
  - Number of Running VMs
  - Number of Running Hosts
  - CPU Consumed



- Memory Consumed

For more information about setting thresholds and response time, see “Configuring Warning and Critical Thresholds” on page 147.

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**6 Click Finish.**

## Resource Pool and vApp Performance

The Resource Pool Performance and vApp Performance monitors can trigger alerts on metrics collected through vSync.

### Resource Pool Performance and vApp Performance Monitor Metrics

The following VMware vSphere metric types for resource pool and vApp performance can be used to configure thresholds in [up.time](#):

Time Interval	A positive integer indicating the number of minutes' worth of performance data samples to average, then compare against threshold definitions (default: 30)
---------------	---

Number of Running VMs: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the average number of VMs powered on during the time interval.
CPU Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of CPU cycles consumed by VMs belonging to this resource pool or vApp.
Memory Consumed: warning threshold and critical threshold	Warning- and critical-level thresholds can be set, using positive integers, for the total percentage of memory consumed by VMs belonging to this resource pool or vApp.

## Configuring Resource Pool Performance or vApp Performance Monitors

To configure a Resource Pool Performance or vApp Performance monitor, do the following:

- 1 Select the monitor from the Add Service Monitor window, in the VMware Monitors section.**
- 2 Click Continue to begin configuring the service monitor.**
- 3 Complete the monitor information fields.**

See “Monitor Identification” on page 144 for more information on configuring service monitor information fields.

- 4 In the Resource Pool Performance Settings section, configure the monitor’s warning- and critical-level threshold values:**

- Time Interval
- Number of Running VMs
- CPU Consumed
- Memory Consumed

For more information about setting thresholds and response time, see “Configuring Warning and Critical Thresholds” on page 147.

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**6 Click Finish.**

## ESX Server Monitors

ESX Server monitors focus on the ESX server host, as a physical computing resource, for monitoring and alerting.

There are currently three ESX-related monitors: ESX (Advanced Metrics), a monitor that uses an [up.time](#) agent on the ESX server; the vSphere ESX Server Performance monitor, which uses metrics transferred to [up.time](#) using vSync, and the legacy ESX Workload monitor, which is a legacy monitor that can no longer be added to [up.time](#).

The metrics collected for these ESX server monitors can be used to trigger [up.time](#) alerts and actions. These performance monitors can answer questions such as the following:

- Are CPU or memory usage on the host too high?
- Are network and disk I/O usage or latency within acceptable limits?
- Are disk and network error rates too high?
- Are memory ballooning targets being exceeded?

## vSphere ESX Server Performance

The vSphere ESX Server Performance monitor allows you to alert based on performance checks on ESX server Elements managed by VMware vSphere, but monitored in [up.time](#) via vSync.

### vSphere ESX Server Performance Monitor Metrics

The following vSphere metric types for ESX server performance can be used to configure thresholds in [up.time](#):

Time Interval	A positive integer indicating the number of minutes' worth of performance data samples to average, then compare against threshold definitions (default: 30)
---------------	---

<p>CPU Check: value type, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for average CPU usage as either percentage usage, or average MHz usage</p>
<p>Memory Check: value type, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for one of the following value types:</p> <ul style="list-style-type: none"> <li>• Usage (%) – percentage of total configured or available memory used</li> <li>• Memory Consumed (MB) – amount of memory consumed by VMs on this host</li> <li>• Memory Active (MB) – amount of memory actively used by VMs on this host</li> <li>• Balloon Memory (MB) – amount of memory allocated by <code>vmmemctl</code> across all VMs on this host</li> <li>• Zero Memory (KB) – memory that only contains 0s allocated to VMs</li> </ul>
<p>Swap Check: value type, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for either swap space used (in MB), or swap rate (the combined swap-in rate and swap-out rate, in KBps, across all VMs on this host).</p>
<p>Disk Device I/O: coverage, value type, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for one of the following value types:</p> <ul style="list-style-type: none"> <li>• Usage (KBps) – aggregate disk I/O rate across all VMs on the host</li> <li>• Physical Device Command Latency (ms) – average time to process a read and write from the physical device</li> <li>• Queue Command Latency (ms) – average time spent in the VMkernel queue per SCSI command</li> <li>• Command Latency (ms) – average time taken to process a SCSI command issued by the Guest OS to the VM</li> </ul> <p>Checks are made against the average for all detected disk devices, or any individual device that is violating the threshold.</p>

<p>Disk Device Errors Check: coverage, value type, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for either the number of SCSI command aborts per minute, or the number of bus resets per minute. Checks are made against the average for all detected disk devices, or any individual device that is violating the threshold.</p>
<p>Network I/O: coverage, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for the aggregate received and transmitted rate (in KBps). Checks are made against the average for all detected network interfaces, or any individual network interface that is violating the threshold.</p>
<p>Network Errors Check: coverage, value type, warning threshold, and critical threshold</p>	<p>Warning- and critical-level thresholds can be set, using positive integers, for the aggregate received and transmitted packets dropped per minute. Checks are made against the average for all detected network interfaces, or any individual network interface that is violating the threshold.</p>

## Configuring vSphere ESX Server Performance Monitors

To configure a vSphere ESX Server Performance monitor, do the following:

- 1 Select the monitor from the Add Service Monitor window, in the VMware Monitors section.**
- 2 Click Continue to begin configuring the service monitor.**
- 3 Complete the monitor information fields.**

See “Monitor Identification” on page 144 for more information on configuring service monitor information fields.



When selecting an Element associated with this service monitor, only ESX servers monitored in [up.time](#) via vSync will appear in the **Single System** list.

- 4 In the vSphere ESX Server Performance Settings section, in the Time Interval sub-section, enter the number of minutes' worth of time samples that will be used to compare thresholds.
- 5 For the following metric categories, select the metric unit of measurement, then configure the monitor's warning- or critical-level threshold values:

- CPU Usage
- Memory
- Swap
- Disk Device I/O
- Disk Errors
- Network I/O
- Network Errors

For more information on these metrics, see “vSphere ESX Server Performance Monitor Metrics” on page 218.

For more information about setting thresholds, see “Configuring Warning and Critical Thresholds” on page 147.

- 6 **Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 7 **Click Finish.**

## ESX (Advanced Metrics)

The ESX (Advanced Metrics) monitor offers greater visibility into your ESX environment by expanding on the high level usage metrics for a virtual machine’s CPU, memory, and disk activity.

### ESX Advanced Metrics Monitor Metrics

The following ESX server metrics can be used to configure thresholds:

Percent Wait	Guest metric - The percentage of time that a virtual CPU was not running. A non-running CPU could be idle (halted) or waiting for an external event such as I/O.
Memory Balloon (Avg)	Guest metric - The average amount of memory, in KB, held by memory control for ballooning.
Memory Balloon Target	Guest metric - The total amount of memory, in KB, that can be used by memory control for ballooning.
Memory Overhead (Avg)	Guest metric - The average amount of additional host memory, in KB, allocated to the virtual machine.
Memory Swap In (Avg)	Guest metric - The average amount of memory, in KB, that was swapped in.
Memory Swap Out (Avg)	Guest metric - The average amount of memory, in KB, that was swapped out.
Memory Zero (Avg)	Guest metric - The average amount of memory, in KB, that was zeroed out.
Memory Swap Used (Avg)	Host metric - The average amount of memory, in KB, that was used by the swap file.
Memory Swap Target	Guest metric - The total amount of memory, in KB, that can be swapped.



Disk Total Latency	Host metric - The average time, in milliseconds, taken for disk commands by a guest OS. This is the sum of <code>kernelCommandLatency</code> and <code>physical deviceCommandLatency</code> .
Disk Kernel Latency	Host metric - The average time, in milliseconds, spent in the ESX Server <code>VMkernel</code> per command.
Disk Device Latency	Host metric - The average time, in milliseconds, taken to complete a command from the physical device.
Disk Queue Latency	Host metric - The average time, in milliseconds, spent in the ESX Server <code>VMkernel</code> queue per write.
Disk Commands Aborted	Host metric - The number of disk commands aborted during the defined interval.
Disk Commands Issued	Host metric - The number of disk commands issued during the defined interval.
Disk Bus Resets	Host metric - The number of bus resets during the defined interval.

## Configuring ESX (Advanced Metrics) Monitors

To configure an ESX (Advanced Metrics) monitor, do the following:

- 1 **Select the monitor from the Add Service Monitor window, in the VMware Monitors section.**
- 2 **Click Continue to begin configuring the service monitor.**
- 3 **Complete the monitor information fields.**  
See “Monitor Identification” on page 144 for more information on configuring service monitor information fields.
- 4 **In the ESX (Advanced Metrics) Settings section, configure the monitor’s warning- and critical-level alerting thresholds by completing the following fields:**
  - Percent Wait

- Memory Balloon
- Memory Balloon Target
- Memory Overhead
- Memory Swap In
- Memory Swap Out
- Memory Zero
- Memory Swap Used
- Memory Swap Target
- Disk Total Latency
- Disk Kernel Latency
- Disk Device Latency
- Disk Queue Latency
- Disk Commands Aborted
- Disk Commands Issued
- Disk Bus Resets
- Response time

For more information on these metrics, see “ESX Advanced Metrics Monitor Metrics” on page 222.

For more information about setting thresholds and response time, see “Configuring Warning and Critical Thresholds” on page 147.

**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**6 Click Finish.**

## ESX Workload

The ESX Workload monitor collects a set of metrics from all of the instances that are running on an ESX v3 or v4 server over a specified time period. This monitor is a legacy monitor that cannot be added to your [up.time](#) configuration as a new service monitor. It exists in upgraded configurations that included it.

The monitor 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 server that is running three instances. You configured the ESX 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 Workload monitor focuses on the value of 315 MB and, since it exceeds the warning threshold, issues an alert.

## ESX Workload Monitor Metrics

The following metrics are used by the ESX Workload monitor:

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 CPU resources, measured in megahertz (MHz), that the instances on the ESX server must consume before <a href="#">up.time</a> issues a warning.
CPU Critical Threshold	The amount of CPU resources, measured in megahertz MHz, that the instances on the ESX server must consume before <a href="#">up.time</a> issues a critical alert.
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 <a href="#">up.time</a> 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 <a href="#">up.time</a> 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 <a href="#">up.time</a> 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 <a href="#">up.time</a> issues a critical alert.
Memory Usage Warning Threshold	The amount of overall system memory, measured in megabytes (MB), that must be exceeded before <a href="#">up.time</a> issues a warning.
Memory Usage Critical Threshold	The amount of overall system memory, measured in megabytes (MB), that must be exceeded before <a href="#">up.time</a> issues a critical alert.

Percent Ready Warning Threshold	The percentage of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server. If the valued returned from the server exceeds this threshold, then <a href="#">up.time</a> issues a warning.
Percent Ready Critical Threshold	The percentage of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server. If the valued returned from the server exceeds this threshold, then <a href="#">up.time</a> issues a critical alert.
Percent Used Warning Threshold	The percentage of CPU time that an instance running on an ESX server is using. If the valued returned from the server exceeds this threshold, then <a href="#">up.time</a> issues a warning.
Percent Used Critical Threshold	The percentage of CPU time that an instance running on an ESX server is using. If the valued returned from the server exceeds this threshold, then <a href="#">up.time</a> issues a critical alert.

## Modifying an ESX Workload Monitor Configuration

To modify the configuration of a legacy ESX Workload monitor, do the following:

- 1 **If required, change the monitor information fields.**  
See “Monitor Identification” on page 144 for more information.
- 2 **In the ESX Workload Settings section, modify any of the monitor’s existing warning- or critical-level threshold values:**
  - Time Interval
  - CPU Usage
  - Network Bandwidth Usage
  - Disk Usage
  - Memory Usage
  - Percent Ready

- Percent Used

For more information on these metrics, see “ESX Workload Monitor Metrics” on page 226.

For more information about setting thresholds, see “Configuring Warning and Critical Thresholds” on page 147.

### **3 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

### **4 Click Finish.**

## Power State Monitors

The power state monitors help you manage both available computing resources within your clusters, resource pools, and other logical divisions in your vSphere-managed infrastructure, as well as power consumption in your physical datacenters. Power state changes to your hosts, and the VMs running on them, can be alerted and acted on.

The power state monitors can answer questions such as the following:

- Has a mission-critical VM been powered off?
- Did a routine maintenance procedure start and complete properly?
- Are enough expected VMs powering down during the weekend, indicating vSphere's Distributed Power Management is functioning correctly?

### ESX Server Power State

The ESX Server Power State monitor watches for changes to the power states of an ESX server that is managed by VMware vSphere, and can run alert or action profiles based on the change.

#### ESX Server Power State Monitor Status Types

In *up.time*, vSphere hosts will be in one of the following states:

Powered On	The host is running.
Powered Off	The host was powered off by an administrator through the vSphere Client.
Put on Standby	The host was put in standby mode either explicitly by an administrator, or automatically by vSphere Distributed Power Management (DPM).

Put in Maintenance	The host state is determined to be “unknown” if it is disconnected or not responding, implying it is in maintenance.
--------------------	--

## Configuring ESX Server Power State Monitors

To configure an ESX Server Power State monitor, do the following:

- 1 Select the monitor from the Add Service Monitor window, in the VMware Monitors section.**
- 2 Click Continue to begin configuring the service monitor.**
- 3 Complete the monitor information fields.**

See “Monitor Identification” on page 144 for more information on configuring service monitor information fields.



When selecting an Element associated with this service monitor, only ESX servers monitored in [up.time](#) via vSync will appear in the **Single System** list.

- 4 In the main ESX Server Power State Settings section, in the Powered On sub-section, do the following:**
  - i In the Set Status to drop-down box, indicate what the monitor’s [up.time](#) state will be when the ESX server’s state is Powered On.**
  - ii From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-on state.**
  - iii From the list select which (if any) Action Profiles will be triggered when the host enters a powered-on state.**
- 4 In the Powered Off sub-section, do the following:**
  - i In the Set Status to drop-down box, indicate what the monitor’s [up.time](#) state will be when the ESX server’s state is Powered Off.**
  - ii From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-off state.**
  - iii From the list select which (if any) Action Profiles will be triggered when the host enters a powered-off state.**



- 4 In the Put on Standby sub-section, do the following:
  - i In the Set Status to drop-down box, indicate what the monitor's up.time state will be when the ESX server's state is Standby.
  - ii From the list, select which (if any) Alert Profiles are triggered when the host enters a standby state.
  - iii From the list select which (if any) Action Profiles will be triggered when the host enters a standby state.
- 4 In the Put in Maintenance sub-section, do the following:
  - i In the Set Status to drop-down box, indicate what the monitor's up.time state will be when the ESX server's state is Unknown.
  - ii From the list, select which (if any) Alert Profiles are triggered when the host enters an unknown state.
  - iii From the list select which (if any) Action Profiles will be triggered when the host enters an unknown state.

For more information on these power states, see “ESX Server Power State Monitor Status Types” on page 229.

- 4 Complete the following settings:
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- 5 Click Finish.

## VM Instance Power State

The VM Instance Power State monitor watches for changes to the power states of a VM running on an ESX server that is managed by vSphere, and can run alert or action profiles based on the change.

## VM Instance Power State Monitor Status Types

A virtual machine’s three basic power states are as follows:

Powered On	The virtual machine is running.
Powered Off	The virtual machine is not running.
Suspended	The virtual machine is not running, but a snapshot of its running applications and processes is retained.

## Configuring VM Instance Power State Monitors

To configure a VM Instance Power State monitor, do the following:

- 1 Select the monitor from the Add Service Monitor window, in the VMware Monitors section.**
- 2 Click Continue to begin configuring the service monitor.**
- 3 Complete the monitor information fields.**

See “Monitor Identification” on page 144 for more information on configuring service monitor information fields.



When selecting a VM associated with this service monitor, only VMs monitored in [up.time](#) via vSync will appear in the **Single System** list.

- 4 In the main VM Instance Power State Settings section, in the Powered On sub-section, do the following:**
  - i In the Set Status to drop-down box, indicate what the monitor’s [up.time](#) state will be when the VM’s state is “powered on”.**
  - ii From the list, select which (if any) Alert Profiles are triggered when the host enters a powered-on state.**
  - iii From the list select which (if any) Action Profiles will be triggered when the host enters a powered-on state.**
- 4 In the Powered Off sub-section, do the following:**

- i** In the **Set Status** to drop-down box, indicate what the monitor's **up.time** state will be when the VM's state is "powered off".
    - ii** From the list, select which (if any) **Alert Profiles** are triggered when the host enters a powered-off state.
    - iii** From the list select which (if any) **Action Profiles** will be triggered when the host enters a powered-off state.
  - 4** In the **Suspended** sub-section, do the following:
    - i** In the **Set Status** to drop-down box, indicate what the monitor's **up.time** state will be when the VM's state is "suspended".
    - ii** From the list, select which (if any) **Alert Profiles** are triggered when the host enters a suspended state.
    - iii** From the list select which (if any) **Action Profiles** will be triggered when the host enters a suspended state.

For more information on these VM power states, see "ESX Server Power State Monitor Status Types" on page 229.

- 4** **Complete the following settings:**
  - Timing Settings (see "Adding Monitor Timing Settings Information" on page 151 for more information)
  - Monitoring Period settings (see "Monitor Timing Settings" on page 149 for more information)
- 5** **Click Finish.**



# CHAPTER 12

## Microsoft Windows Monitors

---

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

<i>Windows Event Log Scanner</i> .....	236
<i>Windows Service Check</i> .....	240
<i>Windows File Shares (SMB)</i> .....	243
<i>Active Directory</i> .....	245

## Windows Event Log Scanner

The Windows Event Log Scanner alerts on specific entries in a Windows log file. This monitor searches through events based on text strings, as well as the log and error type. When the monitor runs, with WMI-based collection, events are retrieved in 15-minute batches; with agent-based collection, the number of events retrieved is user-defined.

To prevent false positives, the monitor ignores log entries that are older than when it was last run. To avoid performance degradation, maximum number of log entries (which has a default 1,000) is 10,000 lines.

## 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 Complete the following fields:**

- Event Log Type

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.

- Other

A custom or external log whose name will be defined in the next step.

- Other Windows Log to Search

When the “Other” event log type is selected in the previous step, this field appears. Enter the name of an additional Windows event log that you want this service monitor to use. This log may accompany an application platform you are running, or could be a custom log; regardless, the name you provide should match the name that appears in the Windows Event Viewer.

- 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.

- Number of Lines

The number of lines in the log file that [up.time](#) will scan, using the criteria specified in the monitor template. The default is 1000 and the maximum is 10000.

- Match source with

The application, system component, or application module that triggered the event.

- 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 event ID with  
A number that identifies the type of event.
- Match user name with  
The name of the user associated with a logged event.
- Match computer name with  
The name of the computer on which the event occurred.
- 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.
- Response Time  
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 147.  
  
To save the data from the thresholds for graphing or reporting, click the **Save for Graphing** checkbox beside the Response Time metrics.

**3 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)



- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 4 **Click Finish.**

## Windows Service Check

The Windows Service Check monitor alerts you to changes in the status of Windows services. Windows services are processes that extend the features of Windows by providing support to other programs; they are controlled in the Microsoft Management Console. 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 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  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 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 3.**
- 4 Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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.

### Configuring Windows File Shares (SMB) Monitors

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

- 1 In the Windows Files Shares (SMB) monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete the following fields:**

- Username

The user name that is required to login to the file share. The value entered can include the file share domain if input with the following formats: `<domain>\<username>` or `<domain>;<username>`

- 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

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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**

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 3** **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.**



- 4 **Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 **Click Finish.**



# CHAPTER 13

## Application Monitors

---

The application monitors track the performance and health of following:

<i>Uptime Agent</i> .....	250
<i>Exchange</i> .....	252
<i>IIS</i> .....	258
<i>WebLogic</i> .....	261
<i>WebSphere</i> .....	269
<i>Web Application Transactions</i> .....	275
<i>Email Delivery Monitor</i> .....	282
<i>Splunk Query</i> .....	288
<i>Live Splunk Listener</i> .....	290

## 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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 147.

- Major

The major version number of the agent. For more information, see “Understanding Major and Minor Versions” on page 14.

- Platform

The operating system on which the agent is installed and running.

- Response Time

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 147.

- 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 3.**

- 4 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information).

- Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**5 Click Finish.**

## Exchange

The Exchange 2003 and Exchange monitors identify 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.

Use [up.time](#)'s Exchange 2003 monitor if you are using and monitoring Microsoft Exchange 2000 or 2003; use the Exchange monitor for later versions (e.g., Microsoft Exchange 2007 and 2010).

### Configuring Exchange 2003 Monitors

To configure an Exchange 2003 monitor for your Microsoft Exchange 2000 or 2003 server, do the following:

**1 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**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 147.

- 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

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information).
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information).
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 Click Finish.**

## Configuring Exchange Monitors

To configure an Exchange monitor for your Microsoft Exchange 2007 or 2010 server, do the following:

- 1 Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.
- 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 147.

- 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 Messages Sent Per Second**  
The maximum number of messages sent per second allowed by the SMTP server.
- **SMTP Messages Received Per Second**  
The maximum number of messages received per second allowed by the SMTP server.
- **SMTP Average Bytes Per Message**  
The average number of message bytes per inbound message received, indicating the size of messages received through an SMTP receive connector.
- **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.
- **Average Delivery Time**  
The average time, in milliseconds, between an Exchange server receiving a message from the client, and an Exchange server delivering the message to an Inbox.
- **Active Connections**  
The number of connections to the Exchange store that have shown activity in the last 10 minutes.
- **Active Client Logons**  
The number of clients that performed any action within the last 10-minute time interval.

- **Active User Count**  
The number of unique user connections that have logged on to the server and shown activity in the last 10-minute time interval.
- **Current Webmail Users**  
The number of unique users currently logged in to Outlook Web Access. This counter decreases when users manually log out or their sessions time out.
- **Webmail User Logons Per Second**  
The number of Outlook Web Access logins or login attempts per second.
- **RPC Averaged Latency**  
The average time, in milliseconds, it takes for the last 1,024 packets to be processed.
- **RPC Operations Per Second**  
The rate that RPC operations occur, and implicitly, how many RPC requests are outstanding.
- **RPC Requests**  
The number of client requests that are currently being processed by the Exchange store.
- **Response Time**  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

**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 151 for more information).
- **Alert Settings** (see “Monitor Alert Settings” on page 151 for more information).

- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information).
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 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 147.

- 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

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 147.

- 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 3.**

- 4 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

- 5 Click Finish.**

## WebLogic

The WebLogic 8 and WebLogic monitors collect 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 515 for more information).

The WebLogic monitors collect the following metrics from a WebLogic server:

<b>Variables</b>	<b>Metrics</b>
<p>Connection Pools</p>	<ul style="list-style-type: none"> <li>• <b>FailuresToReconnectCount</b> The number of times that the connection pool failed to reconnect to a data store.</li> <li>• <b>ConnectionDelayTime</b> The average time that was required to connect to a connection pool.</li> <li>• <b>ActiveConnectionsCurrentCount</b> The current number of active connections in a JDBC connection pool.</li> <li>• <b>ActiveConnectionsHighCount</b> The highest number of active connections in a JDBC connection pool.</li> <li>• <b>LeakedConnectionsCount</b> The total number of connections that have been checked out of, but not returned to, the connection pool.</li> <li>• <b>CurrCapacity</b> The current number of database connections in the JDBC connection pool.</li> <li>• <b>NumAvailable</b> The number of available sessions in the session pool that are not currently being used.</li> <li>• <b>WaitingForConnectionCurrentCount</b> 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>• <b>AccessTotalCount</b> The total number of times an attempt was made to get an EJB instance from the free pool.</li> <li>• <b>BeansInCurrentUseCount</b> The number of EJB instances in the free pool which are currently in use.</li> <li>• <b>CachedBeansCurrentCount</b> The total number of EJBs that are in the execution cache.</li> <li>• <b>ActivationCount</b> The number of EJBs that have been activated.</li> </ul>
Other	<ul style="list-style-type: none"> <li>• <b>HeapSizeCurrent</b> The amount of memory, in bytes, that is in the WebLogic server's JVM heap.</li> <li>• <b>HeapFreeCurrent</b> The current amount of free memory, in bytes, that is in the WebLogic server's JVM heap.</li> <li>• <b>OpenSocketsCurrentCount</b> The current number sockets on the server that are open and receiving requests.</li> <li>• <b>AcceptBacklog</b> The number of requests that are waiting for a TCP connection.</li> <li>• <b>ExecuteThreadCurrentIdleCount</b> 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 data-bbox="591 274 1186 366">•PendingRequestCurrentCount The number of pending requests that are in the server’s execution queue.</li> <li data-bbox="591 387 1215 479">•TransactionCommittedTotalCount The total number of transactions that have been processed by the WebLogic server.</li> <li data-bbox="591 499 1172 591">•TransactionRolledBackTotalCount The total number of transactions that have been rolled back.</li> <li data-bbox="591 612 1179 704">•InvocationTotalCount The total number of times that a servlet running on the WebLogic server was invoked.</li> </ul>

Before you can use the WebLogic monitors, you must perform additional steps outside of [up.time](#). The steps performed depend on the version of your WebLogic server: WebLogic 8 monitoring requires that you deploy the `weblogic.jar` file on the [up.time](#) Monitoring Station; WebLogic 9 or 10 monitoring requires that you enable the Internet Inter-Orb Protocol (IIOP) on your WebLogic server.

## Monitoring WebLogic 8

In order for [up.time](#) to collect information from a WebLogic 8.1 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\weblogic81\server\lib
```

**2 Copy the file to the `externaljar` directory on the Monitoring Station.**

For example, on Windows, copy the file to the following directory:

C:\Program Files\uptime software\uptime\externaljar\



Users who deployed “WebLogic” monitors from [up.time](#) 5.0 or earlier for their WebLogic 8.1 server applications should note that the monitor was renamed to “WebLogic 8” starting with [up.time](#) 5.1. The “WebLogic” monitor is used to monitor WebLogic 9, 10, or 11.

## Configuring WebLogic 8 Monitors

To configure WebLogic 8 monitors, do the following:

- 1 In the WebLogic 8 monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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.

- 3 Specify a warning and critical threshold for the following:**

- **the appropriate WebLogic metrics**  
For more information about each metric, see page 262.
- **Response Time**  
This is the length of time a service check takes to complete.

For more information on using thresholds to set alerts, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 the previous step.**
- 5 **Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 6 **Click Finish.**

## Monitoring WebLogic 9–11

In order for [up.time](#) to collect information from a WebLogic 9, 10, or 11 server, the the Internet Inter-Orb Protocol (IIOP) must be enabled on your WebLogic server.

To enable prepare your WebLogic server for monitoring, do the following:

- 1 **Enable IIOP on your WebLogic server.**

For example, on WebLogic 10, select the **Protocols** tab when configuring server settings, then select the **Enable IIOP** checkbox.
- 2 **Enter an IIOP user name.**
- 3 **Enter an IIOP user password.**
- 4 **If possible, restart the WebLogic server.**

The user name and password created here are used when configuring a WebLogic 10 monitor in [up.time](#).

## Configuring WebLogic Monitors

To configure monitors for WebLogic 9–11, do the following:

**1 In the WebLogic monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 Complete the following fields:**

- Username

The IIOP user name you created when you first enabled IIOP on the WebLogic server.

- Password

The IIOP password you created when you first enabled IIOP on the WebLogic server.

- WebLogic Port

The number of the port number on which the WebLogic server is listening. The default is 7001.

**3 Limit the returned results of a specific resource type by completing some of the following fields:**

- Number of Results

A limit on the number of matching application resources, whose metrics are collected.

- EJB Name Regex Filter

A regular expression used to limit metrics collection to a specific EJB or set of EJBs.

- Servlet Name Regex Filter

A regular expression used to limit metrics collection to a specific servlet.

- JDBC Resource Name Regex Filter

A regular expression used to limit metrics collection to a specific JDBC resource.

**4 Specify a warning and critical threshold for the following:**

- the appropriate WebLogic metrics  
For more information about each metric, see page 262.
- Response Time  
This is the length of time a service check takes to complete.

For more information on using thresholds to set alerts, see “Configuring Warning and Critical Thresholds” on page 147.

**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 the previous step.**

**6 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**7 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 515 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 following table lists the counters the WebSphere monitor collects from a WebSphere Application Server.

Variable	Counters
Connection pools	<ul style="list-style-type: none"> <li>•PoolSize The size of the connection pool to the data source.</li> <li>•FreePoolSize The number of free connections in the pool.</li> <li>•PercentUsed The percentage of the connection pool that is currently in use.</li> <li>•WaitTime 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 The total number of connections that were created.</li> <li>•CloseCount The total number of connections that were closed.</li> <li>•WaitingThreadCount The number of threads that are currently waiting for a connection.</li> <li>•UseTime 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 The number of times that the Enterprise JavaBeans that are running on the server were created.</li> <li>•RemoveCount The number of times that the EJBs were removed.</li> <li>•PassivateCount 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 The total number of method calls that were made to the EJBs.</li> <li>•MethodResponseTime 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 The percent of CPU resources that were used since the last query.</li> <li>•HeapSize The total amount of memory that is available for the JVM.</li> <li>•UsedMemory The amount of memory that is being used by the JVM.</li> </ul>
Other	<ul style="list-style-type: none"> <li>•ActiveCount The number of global transactions which are concurrently active.</li> <li>•CommittedCount The total number of global transactions that have been committed.</li> <li>•RolledBackCount The total number of global transactions that have been rolled back.</li> <li>•LiveCount The number of servlet sessions that are currently cached in memory.</li> <li>•PoolSize The average number of threads in the servlet connection thread pool.</li> <li>•TimeSinceLastActivated 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, enable the **Application Security** option.
  - Under Enterprise Applications - perfServletApp - Security role to user/group mapping, disable the **Everyone** option, and enable the **All authenticated** option.
- 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete the following fields:**

- WebSphere Port

The number of the port number on which WebSphere is listening. The default is 9080.

- Response Time

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 3 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.**

- 4 Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 Click Finish.**

## Web Application Transactions

A Web transaction is a series of Web pages that together fulfill a specific function for end users. A common Web transaction example is the checkout process on an e-commerce site, during which end users select a shipping option, pay for their items, and have their credit card verified. During this transaction, many calls are made to the application and data layers as the end-user provides, and the servers process, information.

Although the type of Web application that is monitored by up.time users is typically different (e.g., intranet applications), the structure of the transaction is the same: an end user steps through a sequence of Web pages that take inputted information and initiate appropriate actions with application or database servers.

The up.time Web Application Transaction monitor tests the speed and availability of an end-user Web transaction. Specifically, the Web Application Transaction monitor performs two roles:

- it confirms the general availability of an end-user Web transaction by executing a previously recorded script then reporting whether all pages that make up the web transaction were successfully processed
- it reports on the speed of the Web transaction both as a whole, and broken down by previously defined stages

Both the availability and speed of Web transactions can be used in reports and as triggers for alerts.

### Using the Web Application Transaction Monitor

Use the Web Application Transaction monitor to record a series of URLs that together make up a transaction. This recording should be of a transaction that acts as a suitable test of your Web application delivery infrastructure.

During the recording process, declare checkpoints that demarcate significant stages in the Web transaction. Isolating the different stages in an end-user transaction allows you to view stage-specific speed tests in reports, which ultimately helps you identify where problem areas exist.

For example, if a transaction relies on processing on the application layer, makes multiple calls to the data layer, and is accessible worldwide, creating

checkpoints during the recording phase helps you ascertain whether the application server, database management server, or network may be the reason behind a poorly performing transaction.

The following sample checkpoints could be created for an e-commerce transaction:

- Browse Catalog
- Add to Shopping Cart
- Checkout
- Credit Card Validation

The following sample checkpoints could be created for an internal office transaction:

- Login
- Browse Orders
- View Order Details

## Configuring Web Application Transaction Monitors

You can define Web application transactions by manually stepping through one and declaring checkpoints at key stages:

### **1 Open a Web browser, and configure its proxy settings so that you can record a transaction:**

- Open the dialog where network and connection settings are made.
- Configure the browser's proxy to "localhost" on port 8001.
- Ensure these settings have also been applied to SSL or secure communications.
- Set the proxy to bypass the Monitoring Station.

This step will require you to select an option such as "no proxy" or "bypass proxy server," and may also require you to manually enter your Monitoring Station URL or IP address.


Using the monitor as a proxy will allow it to intercept Web traffic as you generate it.

- 2 **In the browser, navigate to the starting point of the Web application whose performance you will be monitoring.**
- 3 **In the [up.time](#) Add Service window, select the Web Application Transaction monitor, then click Continue.**

The Web Application Transaction Recorder is displayed, and the monitor is now listening on port 8001 for traffic.


- 4 **Begin stepping through the Web transaction as an end user, providing the required data or actions.**

Every URL visited during the transaction is logged and displayed in the recorder.

 The Web Application Transaction monitor records all data inputted during recording: this includes any login information. It is recommended that you use a test account for the Web application, otherwise any user data will be visible in the recorded script.


- 5 **At each major step in the Web transaction that signals a new analysis point, enter a checkpoint name in the text box at the top of the window, then click Mark Checkpoint.**

For example, create a checkpoint at a transaction step where the application takes user-inputted data and makes database calls.

 You will later set Warning and Critical thresholds that apply to every segment declared in your recording. It is recommended that the divisions between your checkpoint intervals are reasonably consistent.

- 6 **Continue to repeat steps 4 and 5 until you have completed enough of the Web transaction to test it, then click Next.**
- 7 **Complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

 Even though the Web application performance is not directly tied to an Element’s performance, making this selection is still required: the service based on this monitor needs to be associated with an Element in order to be viewed in areas such as **Global Scan** or **My Infrastructure**.

- 8 **Configure the Web Application Transaction Settings:**

- Script to play back

If desired, optimize the playback script (e.g., remove extraneous URLs such as image downloads).

- Text that must appear

Enter a text string that can be used to confirm the script playback was successful (e.g., a phrase that appears on the final page of the application). If the monitor does not find this text, its status changes to Critical. By providing mandatory text, you can ensure an alert is triggered in cases where a Web application is malfunctioning, but checkpoint-to-checkpoint times are fast enough to fulfill response time requirements.

- Text that must not appear

Enter a text string that should *not* appear at any point during the script playback (e.g., a client- or server-error HTTP status code). If the monitor finds this text, its status changes to Critical. Use this feature, as you would use mandatory text, to ensure a malfunctioning application triggers an alert.

- User Agent String

Select the Web browser and version used to record the script. This selection determines the user agent string used in the HTTP requests to the application server, and should be provided in case the application blocks access by scripts.

- Checkpoint Times

Enter the Warning and Critical Checkpoint Time thresholds. An alert is generated with these thresholds if any of the recorded Web transaction's checkpoint times exceeds the supplied values.

- Response Time

Enter the Warning and Critical Response Time thresholds. An alert is generated with this threshold if the entire transaction playback time exceeds the supplied values. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 9 **Enter Warning- and Critical-level thresholds for the overall response time of the monitor.**



Most of the monitor's Response Time is comprised of the Delivery Time and the Retrieve Time. Ensure the values provided for the Response Time thresholds roughly correspond with those provided for the other thresholds. For more information, see "Configuring Warning and Critical Thresholds" on page 147.

**10 Complete the following settings:**

- Timing Settings (see "Adding Monitor Timing Settings Information" on page 151 for more information)
- Alert Settings (see "Monitor Alert Settings" on page 151 for more information)
- Monitoring Period settings (see "Monitor Timing Settings" on page 149 for more information)
- Alert Profile settings (see "Alert Profiles" on page 433 for more information)
- Action Profile settings (see "Action Profiles" on page 441 for more information)


**11 Click Finish.**

## Viewing and Diagnosing Web Transaction Performance

To view Web transaction performance via playback, create a Service Metrics graph for the Web Application Transaction monitor's system. To generate a Service Metrics graph, either select the system to which the Web Application Transaction monitor is associated in **My Infrastructure**, or the monitor itself in the main **Services** panel. Click the **Graphics** tab, then click **Service Metrics**.

The Service Metrics graph shows how long each transaction segment took to complete during playback, and in doing so, provides an end-to-end performance snapshot of the components of your infrastructure that deliver applications to users. For example, the following metrics graph shows that

the execution of the commands found in checkpoint 3 took excessively long to complete:



Since other checkpoints performed well, the poor performance of a single checkpoint indicates possible issues with a particular server, and not the network infrastructure. This theory can be further investigated by looking at the performance metrics for the server in question.

Use the Web Application Transaction monitor’s playback script to verify which servers are being used during a problem checkpoint. In the **Service Instances** panel, click the monitor to view the script, then locate the system that is being accessed (e.g., with GET and POST commands). Use this as an investigative starting point: although an application or Web server is often referenced in the script, the problem may be found deeper in the application stack (e.g., a database server to which the referenced Web server makes calls during the checkpoint).

## Using Web Transaction Performance in SLA Reports

Your Web applications will typically call on systems on the application and database tiers, as well as make use of internal- and external-facing network devices. Since the Web Application Transaction monitor directly reports on the performance of a Web transaction, it in effect indirectly reports on the health of your IT infrastructure as a whole.

This broad reporting coverage makes the Web Application Transaction monitor an ideal monitor to include in service level agreement reports.

For more information on SLA reports, see “Reports for Service Level Agreements” on page 505.

## Email Delivery Monitor

Although specific [up.time](#) monitors are available for your POP, IMAP, and SMTP servers, their monitoring duties focus on availability and response time. To test your IT infrastructure's ability to send or receive emails within a reasonable amount of time, use the Email Delivery monitor.

Typically, email delivery tests include a server that is part of your IT infrastructure and monitored by [up.time](#). In these cases, you will test either incoming mail delivery times by supplying information about a monitored POP3 or IMAP server, or test outgoing mail delivery times by supplying information about a monitored SMTP server.

The Email Delivery executes several steps in order to calculate mail delivery and retrieval time:

- the monitor requests an internal or external SMTP server to send a generated test mail (when the monitor asks the SMTP server to send the mail, the monitor records the delivery time)
- the monitor waits for five seconds, then logs in to and checks an internal or external POP3 or IMAP mail server to verify the mail was received
- if the test mail is not found, the monitor waits another five seconds and checks again (and continues to check until the process has either timed out or the mail is found)
- the monitor confirms the mail was received and reports both the delivery and retrieval times

## Configuring Email Delivery Monitors

Define the Email Delivery monitor by providing information about the outgoing and incoming mail servers:

## 1 Complete the monitor information fields.

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.



Once created, the Email Delivery monitor service can be included with status reports for the system or group you select. If this monitor is reporting outgoing mail delivery times, the system should be a monitored SMTP server; if incoming mail delivery times are being measured, the system should be a monitored POP3/IMAP mail server.

## 2 Complete the Outgoing Email Settings:

- **SMTP Hostname**  
Provide the name or IP address of the SMTP server.
- **SMTP Port**  
Provide the port used to communicate with the SMTP server. Leave this field blank to use the default SMTP port (25).
- **SMTP Username**  
Provide the authenticated SMTP user name.
- **SMTP Password**  
Provide the authenticated SMTP user password.
- **SMTP Uses SSL**  
Specify whether the SMTP server sends and receives encrypted communication using SSL.
- **Destination Email Address**  
Enter the test email address used by the monitor. The monitor sends an email to this address, and this address is checked for receipt of the test email.



Although the Email Delivery monitor attempts to promptly find and delete test emails, network issues may prevent timely cleanups. To avoid potential Inbox clutter, it is recommended that you create a dedicated test email account as the destination address.

- **Delivery Time**

Enter the Warning and Critical Delivery Time thresholds. The smallest unit of time used for these thresholds is seconds. Given the speed at which SMTP servers should finish processing an outgoing email, is it recommended that you set the Warning threshold to one second.

### **3 Complete the Incoming Email Settings:**

- **POP3/IMAP Hostname**

Provide the name or IP address of the mail server.

- **POP3/IMAP Port**

Provide the port used to communicate with the mail server. Leave this field blank to use the default POP3 or IMAP port (110 and 143, respectively).

- **POP3/IMAP Username**

Provide the login name for the destination email account.

- **POP3/IMAP Password**

Provide the password for the destination email account.

- **POP3/IMAP Uses SSL**

Specify whether the mail server sends and receives encrypted communication using SSL.

- **Retrieve Time**

Enter the Warning and Critical retrieval time thresholds. The smallest unit of time used for these thresholds is seconds, and the monitor checks for receipt of the test mail in five-second intervals. Enter values in multiples of five.

### **4 Enter Warning- and Critical-level thresholds for the overall response time of the monitor.**

Enter the Warning and Critical Response Time thresholds. An alert is generated with this threshold if the combined email delivery and response time exceeds the supplied values. For more information, see “Configuring Warning and Critical Thresholds” on page 147.


**5 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**6 Click Finish.**

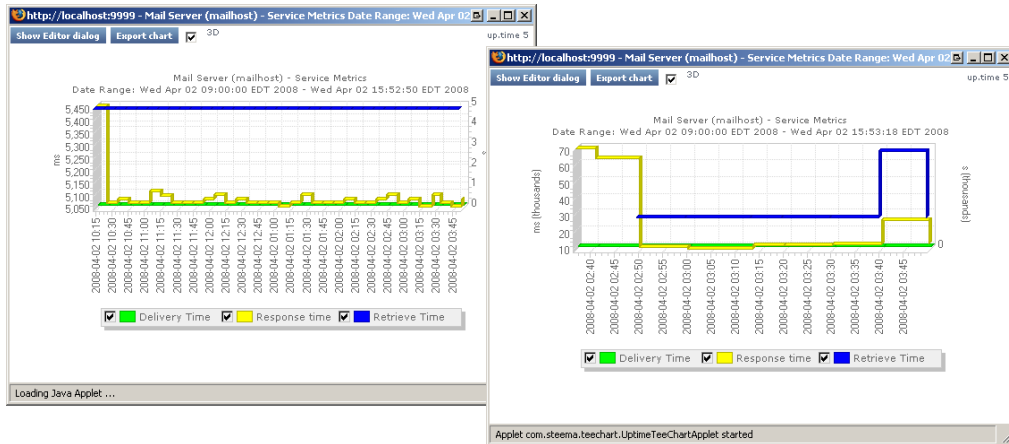
## Diagnosing and Reporting Email Delivery Problems

If the Email Delivery monitor reaches a Critical state, the first investigation step is to review the message produced by [up.time](#). In the **System Status** panel, view the message belonging to the system to which the monitor is attached, which should point you in the right direction. For example, the status message below indicates the monitor reached a critical state because the retrieval time from an external POP3 server exceeded the defined threshold; your SMTP server is most likely not responsible for the delay:



Speculation based on the status message can be confirmed using a Service Metrics graph for the Email Delivery monitor’s system. This graph

indicates whether the delivery and retrieval time are within acceptable limits (below left), or if one or both are unusually long (below right): .




To generate a Service Metrics graph, either select the system to which the Email Delivery monitors are associated in **My Infrastructure**, or the monitor itself in the main **Services** panel. Click the **Graphics** tab, then click **Service Metrics**.

Even if the Service Metrics graph indicates delivery and retrieval times are not exceeding defined thresholds (and **up.time** is not sending out critical alerts), it is still an ideal investigative starting point if you are getting critical feedback from your users about email delivery times.

If the Email Delivery monitor's Service Metrics graph confirms that there are delays somewhere within your network infrastructure, you can investigate further by using the service monitor you created for your mail server. Co-ordinate your Email Delivery monitor's metrics graphs or reports with those from a service monitor you have assigned to your mail server (e.g. Exchange) while focusing on metrics that may be related outgoing or incoming mail time delays. For example, in the Exchange service monitor metrics graph below, the mail server experienced a high



SMTP Local Queue Length that did not always coincide with the SMTP Messages Per Second count:



## 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 653 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 144.

**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 defined warning and critical thresholds.

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 ▾		
Warning	is greater than or equal to ▾	5 results
Critical	is greater than or equal to ▾	10 results

- **Response Time**

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

**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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 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 (from Splunk 4.x only). This monitor is very similar to the External Check monitor (see page 380), and uses scripts that are bundled with [up.time](#) (found in the `/scripts` subdirectory) to return Live Splunk information to the Monitoring Station.

To use this monitor, you must first modify the two Splunk scripts that are included with [up.time](#):

- `alertUptimeStatusHandler.sh`
- `alertUptime.py`

This pair of scripts take 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.
- --monitor  
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:

```
alertUptimeStatusHandler.sh --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 22). 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 generated on a Splunk server, you must do the following:

- 1 **Edit the `alertUptime.py` script to point to the **up.time** Monitoring Station:**
  - Navigate to the `/scripts` directory on the Monitoring Station.
  - Open the file `alertUptime.py` in a text editor.
  - Find the following entry in the file:

```
host = "uptime-host"  
port = "9996"
```
  - Change the values for `host` and `port` to the host name and port of the Monitoring Station.
  - Save and close the file.
- 2 **Edit the `alertUptimeStatusHandler.sh` script to configure how the Live Splunk is reported on the Monitoring Station:**

- Open `alertUptimeStatusHandler.sh` in a text editor (found in the `/scripts` directory on the Monitoring Station).
  - For the `message` option, enter a diagnostic message that accompanies a Live Splunk captured by the `up.time` service monitor.
  - For the `status` option, enter the status of the service being monitored.
  - For the `monitorName` option, enter the name of the service monitor that is listening to the Live Splunk.
  - Save and close the file.
- 3 Copy the `alertUptimeStatusHandler.sh` and `alertUptime.py` scripts from the Monitoring Station's `/scripts` directory to the `/data/splunk/bin/scripts` directory on the Splunk server.**
  - 4 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 `alertUptimeStatusHandler.sh` in the field.

## Configuring the Live Splunk Listener Monitor

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 144.

**2 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**3 Click Finish.**





# CHAPTER 14

## Database Monitors

---

The database monitors track the performance and health of following:

<i>MySQL (Advanced Metrics)</i> .....	296
<i>MySQL (Basic Checks)</i> .....	303
<i>Oracle (Advanced Metrics)</i> .....	305
<i>Oracle (Basic Checks)</i> .....	308
<i>Oracle Tablespace Check</i> .....	311
<i>SQL Server (Basic Checks)</i> .....	314
<i>SQL Server (Advanced Metrics)</i> .....	318
<i>SQL Server Tablespace Check</i> .....	322
<i>Sybase</i> .....	327

## 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 behavior 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 144.

**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 147.

- **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 any of the following occurs:

- 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
  - “`INSERT DELAYED`” can only be used for “`INSERT`” statements that specify value lists, as 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 the Knowledge Base article “SQL Server Locks.”
- **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 the Knowledge Base article “SQL Server Locks.”
- **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 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 Click Finish.**



## MySQL (Basic Checks)

The MySQL (Basic Checks) monitor does the following:

- 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, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 **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**  
Type or copy the script that you want **up.time** to match against the database. Use this option if your script is short or will not regularly change. This option is required if you do not have access to the file system on the Monitoring Station.
  - **Script File**  
As an alternative to directly entering a script, enter the full path on the Monitoring Station to the script that this monitor will run against the database.
  - **Match**  
Enter a string that you want to match against the return value from the script.
  - **Response Time**  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 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 151 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
    - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
    - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
    - Action Profile settings (see “Action Profiles” on page 441 for more information)
  - 5 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 **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**  
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 147.
- 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 151 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
    - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
    - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
    - Action Profile settings (see “Action Profiles” on page 441 for more information)
  - 5 **Click Finish.**

## Oracle (Basic Checks)

The Oracle (Basic Checks) monitor does the following:

- 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 311) 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, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 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**  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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 327.

## Configuring Oracle Tablespace Check Monitors

To configure Oracle Tablespace Check monitors, do the following:

- 1 In the Oracle Tablespace Check monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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  
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 147.
- 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 151 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
    - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
    - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
    - Action Profile settings (see “Action Profiles” on page 441 for more information)
  - 5 **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 does the following:

- 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, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 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 the Knowledge Base article “SQL Server Ports.”

- 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 146. 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

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

**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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**5 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, consider a host configured to have the following:

- 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 the Knowledge Base article “SQL Server Locks.”

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, complete the monitor information fields.**  
To learn how to configure monitor information fields, see “Monitor Identification” on page 144.
- 2 In the Instance field, type the name of the SQL server instance to which you want to connect.**
- 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 147.

- **Lock Wait / Sec.**  
The amount of time, in seconds, to wait for a database lock. For more information about locks, see the Knowledge Base article “SQL Server Locks.”
- **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 the Knowledge Base article “SQL Server Locks.”
- **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**  
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.
- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 6 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

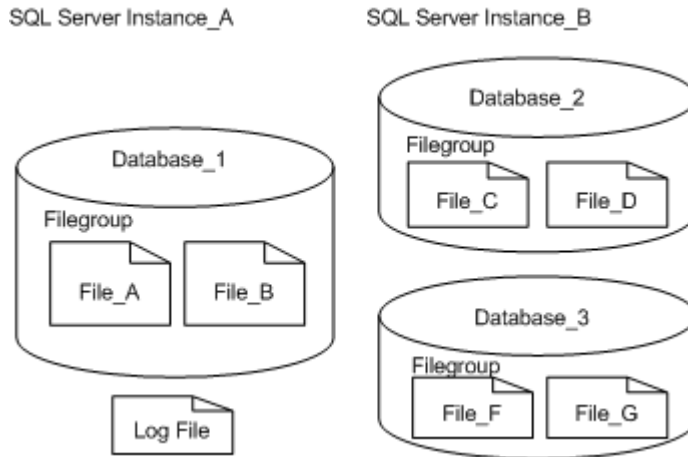
When you configure your SQL Server databases, you can the maximum size of data files to prevent disk drives from running out of space. If you do not

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 file groups 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

### 2 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 the Knowledge Base article “Configuring SQL Server Ports.”

- 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 synonyms 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  
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 147.
- 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 151 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
    - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
    - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
    - Action Profile settings (see “Action Profiles” on page 441 for more information)
  - 5 Click Finish.**



## Sybase

The Sybase monitor does the following:

- 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 sampledbl; 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 **Click Finish.**



# CHAPTER 15

## Network Service Monitors

---

The network service monitors track the health and performance of the following:

<i>DNS</i> .....	332
<i>FTP</i> .....	335
<i>HTTP (Web Services)</i> .....	337
<i>IMAP (Email Retrieval)</i> .....	341
<i>LDAP</i> .....	343
<i>NFS</i> .....	347
<i>NIS/YP</i> .....	349
<i>NNTP (Network News)</i> .....	351
<i>Ping</i> .....	355
<i>POP (Email Retrieval)</i> .....	357
<i>SSH (Secure Shell)</i> .....	359
<i>SMTP (Email Delivery)</i> .....	361
<i>SNMP</i> .....	363
<i>TCP</i> .....	370

## 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, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 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

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

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 147.

**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 151 for more information)

- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)

- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)

- Alert Profile settings (see “Alert Profiles” on page 433 for more information)

- Action Profile settings (see “Action Profiles” on page 441 for more information)

**5 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete the following fields:**

- Port

The number of the port number on which the FTP server is listening. The default is 21.

- Server Response

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 147.

- Response Time  
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 147.
- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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 verify several things:

- 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, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 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**

The user ID and password, in the form `userid:password`. For example:

```
jlamport:bluefrog5
```

- **Virtual Host**

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**

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 147.

- **Follow Re-Direct Actions**

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**

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” 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**

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**

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  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 Click Finish.**

## IMAP (Email Retrieval)

The IMAP (Email Retrieval) monitor confirms whether an IMAP server is doing the following:

- listening on a defined port
- running on a defined system or on a group of systems
- using a particular version

### Configuring IMAP (Email Retrieval) Monitors

To configure IMAP (Email Retrieval) monitors, do the following:

- 1 In the IMAP (Email Retrieval) monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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 156.

- Server Response

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 147.

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**

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 147.

**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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**5 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
- one of the many freely-available LDAP browsing and editing tools
- 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 Complete the following fields:**

- Port

The number of the port number on which the LDAP server is listening. The default is 389.

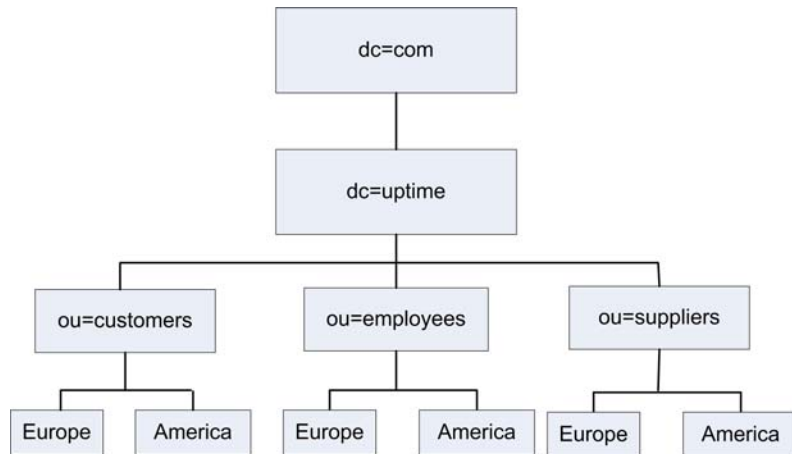
- Password

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

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 Complete the following fields:**

- **Mounts**

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 147.

- **Response Time**

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 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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, enabling 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 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

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

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 147.

**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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**5 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete the following fields:**

- Port

The number of the port on which the NNTP server is listening. The default is 119.

- Server Response

The server response according to the value that you want to measure.

For information on command implementation, see “Command Implementation” on page 351.

For information on response categories, see “Response Category” on page 352.

For information on general response, see “Response Codes” on page 352.

- Response Time

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

### **5 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 the following:

- 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, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 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

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

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 147.

**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 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

**5 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 verify whether a POP server is doing the following:

- listening on a defined port
- running on a defined system
- running on a group of systems
- running a particular version of POP

## Configuring POP (Email Retrieval) Monitors

To configure POP (Email Retrieval) monitors, do the following:

- 1 In the POP (Email Retrieval) monitor template, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete the following fields:**

- Expected Server Response

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 147.

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

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 Click Finish.**



## SSH (Secure Shell)

The SSH (Secure Shell) monitor determines if the secure shell utility (SSH) is available and is running on the defined port. SSH is both a program and a network protocol for securely logging into and executing commands on a remote computer. It provides secure encrypted communications between two untrusted hosts over an insecure network.

### Configuring SSH (Secure Shell) Monitors

To configure SSH (Secure Shell) monitors, do the following:

- 1 In the SSH (Secure Shell) monitor template, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete Secure Shell monitor settings by entering the appropriate Warning and Critical thresholds.**

For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- Port

The number of the port on which SSH is listening. The default is 22.

- Major

The major version number of SSH. This is the number immediately to the left of the decimal in the version number. In the following example, the major version number is 2:

```
SSH_2.0_SUN_SSH1.0
```

- Minor

The minor version number of SSH. This is the number immediately to the right of the decimal in the version number. In the following example the major version number is 0:

```
SSH_2.0_SUN_SSH1.0
```

- SSH Server Version

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

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 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 Complete the following fields:**

- Port

The number of the port on which the SMTP server is listening. The default is 25.

- Expected Server Response

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 147.

- Response Time

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

- 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 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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 do the following:

- 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 Element to **up.time**. For more information, see “Adding Systems or Network Devices” on page 71.

## 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:

- 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 to do the following:

- 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 Service Instance**.
- 3 In the **Add Service Monitor** 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 367 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 367 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 the version number of an SNMP implementation from the SNMP Version dropdown list.**
- 2 **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).

- 3 **Enter the number of the port on which SNMP is listening in the SNMP Port field.**
- 4 **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](#).

### 5 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 147.

Each OID has one or more settings associated with it, as shown in the following image:



.1.3.6.1.6.3.1.1.4.1 ▾	
<b>snmpTrapOID ▾</b>	
Warning	Select a comparison method ▾
Critical	Select a comparison method ▾
.1.3.6.1.6.3.1.1.6.1 ▾	
<b>snmpSetSerialNo ▾</b>	
Warning	Select a comparison method ▾
Critical	Select a comparison method ▾
.1.3.6.1.6.3.1.2.2.5 ▾	
<b>snmpSetGroup ▾</b>	

- Response Time
 

Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 6 **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.**
  - 7 **Complete the following settings:**
    - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
    - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
    - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
    - Action Profile settings (see “Action Profiles” on page 441 for more information)
  - 8 **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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

**2 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  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 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 151 for more information)
    - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
    - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
    - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
    - Action Profile settings (see “Action Profiles” on page 441 for more information)
  - 5 **Click Finish.**



# CHAPTER 16

## Advanced Monitors

---

You can configure advanced monitors to collect performance information. Advanced monitors are described in the following sections:

<i>Overview</i> .....	374
<i>Custom Monitors</i> .....	376
<i>Custom with Retained Data</i> .....	378
<i>External Check</i> .....	380
<i>Plug-In Monitors</i> .....	382

## Overview

In some cases, the standard [up.time](#) service monitors may not fully enable you to monitor all of the systems, applications, and proprietary devices in your environment; in some cases, you may need to capture unique metrics. To do this, you can configure advanced service monitors, or download and install customized plug-in monitors.

These 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 22). 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 (e.g., `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 140.
- For a description of monitor identification information fields, see “Monitor Identification” on page 144.
- For a description of monitor timing settings, see “Monitor Timing Settings” on page 149.
- For a description of alert settings, see “Monitor Alert Settings” on page 151.
- For a description of Alert Profiles, see “Alert Profiles” on page 433.
- For a description of Action Profile, see “Action Profiles” on page 441.

## 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 477) or a Service Metrics graph (see “Viewing System and Service Information” on page 52).

## Configuring Custom Monitors

To configure Custom monitors, do the following:

**1 In the Custom monitor template, complete the monitor information fields.**

To learn about monitor information fields, see “Monitor Identification” on page 144.

**2 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 146.
- **Response Time**

Optionally, enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.
- 3 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.**
- 4 Complete the following settings:**
  - Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
  - Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
  - Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
  - Alert Profile settings (see “Alert Profiles” on page 433 for more information)
  - Action Profile settings (see “Action Profiles” on page 441 for more information)
- 5 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 374 for more information)

You can specify whether 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 477) or a Service Metrics graph (see “Viewing System and Service Information” on page 52).

You can also specify whether the monitor will use retained values as arguments in the custom monitor script. The values used are from the last instance the monitor was run. This most recent data sample can be used in cases where you would like to set alerts based on value comparisons.

## 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, complete the monitor information fields.**

To learn how to configure monitor information fields, see “Monitor Identification” on page 144.

- 2 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.

- Pass retained variables from previous collection as environment variables  
Select this checkbox if you would like the monitor to be able to use the most recently collected variables as arguments in the script.
- 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  
Enter the Warning and Critical Response Time thresholds. For more information, see “Configuring Warning and Critical Thresholds” on page 147.

### 3 Complete the following settings:

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

### 4 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 374 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 144.

- 2 Complete the following settings:**

- Timing Settings (see “Adding Monitor Timing Settings Information” on page 151 for more information)
- Alert Settings (see “Monitor Alert Settings” on page 151 for more information)
- Monitoring Period settings (see “Monitor Timing Settings” on page 149 for more information)
- Alert Profile settings (see “Alert Profiles” on page 433 for more information)
- Action Profile settings (see “Action Profiles” on page 441 for more information)

- 3 Click Finish.**

## Plug-In Monitors

[up.time](#) can be integrated with plug-in monitors that are not part of the standard distribution. Plug-in monitors are custom service monitors that have been created by uptime software, or other [up.time](#) users.

The benefit of sharing plug-in monitors is that uptime customers with relatively unique, but not exclusive, monitoring needs can share the results of their efforts with each other. Additionally, if uptime software creates a custom plug-in monitor for a customer's environment, this monitor would then be available to all customers.

The uptime Support Portal is the host to all plug-in monitors. There, you can find and download a plug-in monitor archive before installing it on your Monitoring Station. All plug-in monitors that have been installed will always appear in the **Add Service Monitor** window, ready to be configured as would any pre-packaged system monitor:



## Installing Plug-In Monitors

To use a plug-in monitor with [up.time](#), do the following:

- 1 **Download the plug-in monitor from the uptime Support Portal.**
- 2 **Locate the `loadpluginmonitor` script, which is found in your [up.time](#) `scripts` directory.**
- 3 **In a command line shell, change to the `[UP.TIME_HOME]/scripts/` directory, and locate the `loadpluginmonitor` script.**
- 4 **Run the `loadpluginmonitor` script with a single argument that points to the location and name of the plug-in monitor you downloaded.**



The plug-in monitor will be installed in a subdirectory under the `/scripts` directory. The installation directory is determined by the plug-in monitor's XML file.

**5** Run the `up.time` GUI.

**6** Click **Services** on the `up.time` tool bar.

**7** Click **Add Service Instance** in the **Tree Panel**.

The **Add Service Monitor** window appears.

**8** In the **Advanced Monitors** section, you will see the plug-in monitor you added to `up.time`.

You can now select and configure the plug-in monitor.



# CHAPTER 17

## Configuring Users

---

This chapter describes the [up.time](#) user management functions in the following sections:

<i>Working with User Roles</i> .....	386
<i>Working with Users</i> .....	389
<i>Working with User Groups</i> .....	393
<i>Managing Distribution Lists</i> .....	396
<i>Working with Notification Groups</i> .....	399
<i>Changing How Users Are Authenticated</i> .....	401

## Working with User Roles

User roles define the following:

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

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

## Adding User Roles

To add user roles, do the following:

- 1 On the [up.time](#) tool bar, click Users.**
- 2 In the Tree panel, click Add New User Role.**  
The **Add User Role** window appears.
- 3 Type a name for this role in the Name of User Role field.**  
This name will appear in the [up.time](#) Web interface.
- 4 Optionally, type a short description in the Description of User Role field.**
- 5 In the first Permissions area of the Add User Role window, you assign the user permissions to View, Add, Edit, or Delete the following items by clicking the checkbox beside each item:**
  - Users
  - Elements
  - Services
  - Element Groups
  - Action Profiles
  - Alert Profiles

- Time Periods
  - Service Level Agreements
  - Element Views
- 6** **Optionally, in the second Permissions area enable one or more of the following options by clicking the Allowed checkbox:**
- Administrator  
The user can perform all [up.time](#) administration tasks.
  - Acknowledge Alerts  
The user can acknowledge an alert. See “Understanding Alerts” on page 430 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 456 for more information.
- 7** **Click Save.**

## Viewing User Roles

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

To view user roles, do the following:

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

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

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

Info				
Permission	View	Add	Edit	Delete
Users	✓	-	-	-
Elements	✓	-	-	-
Services	✓	-	-	-
Element Groups	✓	-	-	-
Action Profiles	✓	-	-	-
Alert Profiles	✓	-	-	-
Time Periods	✓	-	-	-
Service Level Agreements	✓	-	-	-
Element Views	✓	-	-	-
Permission	Allowed			
Administrator	-			
Acknowledge Alerts	-			
Save Reports	✓			

## Editing User Roles

To edit user roles, do the following:

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

The **Edit User Roles** window appears.

- 3 Edit the user role information as described in the section “Adding User Roles” on page 386.**

## Working with Users

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

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

### Adding Users

To add users, do the following:

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

The **Add User** window appears.

- 2 Type a name for the user, which will be used to log into [up.time](#), in the Username field.**

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

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

If using an AD or LDAP directory to authenticate users, [up.time](#) will refer to the directory for password information during user login. For more information, see “Changing How Users Are Authenticated” on page 401.

- 4 If you have set a user password, re-enter it in the Confirm Password field.**
- 5 Enter the full name of the user in the First Name and Last Name fields.**
- 6 Optionally, enter the user's geographical location or department in the Location field.**
- 7 If the user will be receiving alerts via email, enter the user's email address in the Email Address field.**

- 8 Select one of the following options from the Time Period for Emailing dropdown list:**
  - 24x7
  - 9am to 5pm weekdays
  - another Monitoring Period that you have previously created
- 9 If the user will receive alerts on their cell phone or pager, enter the email address of the user's cell phone or pager in the Pager/Cellphone Address field.**

The email address takes the following format:

```
<number>@mobile_provider_domain
```

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

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

The options are the same as the ones listed in Step 8.
- 11 If the user will receive alerts via the Windows messaging service, enter the name of the user's computer in User's Windows Desktop Hostname field.**



To receive popup alerts, you must enable the Windows messaging service on the user's computer. See "Enabling the Windows Messaging Service" on page 433 for information.

- 12 Enter the workgroup or domain to which the user's computer belongs in the User's Windows Desktop Workgroup field.**
- 13 Select an option from the Time Period for Windows Popups dropdown list**

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



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



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

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

- Alert on Critical

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

- Alert on Warning

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

- Alert on Unknown

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

- Alert on Recovery

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

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



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

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

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

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

For more information on user roles, see the section “Working with User Roles” on page 386.

- 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 393.

- 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 389.

## Working with User Groups

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

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

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



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

## Adding User Groups

To add user groups, do the following:

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


## Viewing User Groups

To view user groups, do the following:

- 1 In the **Tree panel**, click **View User Groups**.  
A list of user groups appears in the **User Groups** subpanel.

## Editing User Groups

To edit user groups, do the following:

- 1 In the **Tree panel**, click **View User Groups**.
- 1 Do one of the following:
  - Click the **Edit** icon (  ) beside the name of the user group.
  - Click the name of the user group whose information you want to edit, and then click **Edit User Group** in the **User Group** subpanel.

The **Edit User Group** window appears.

- 2 **Edit the information as described in the section “Adding User Groups” on page 394.**

## Deleting User Groups

To delete user groups, do the following:

- 1 **In the Tree panel, click View User Groups.**
- 2 **Click the Delete icon (  ) beside the name of the user group that you want to delete.**

You cannot delete the SysAdmin user group.

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

## Managing Distribution Lists

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

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

### Adding Distribution Lists

To add Distribution Lists, do the following:

- 1 Click Users on the [up.time](#) tool bar.**
- 2 In the Tree panel, click Add New Distribution List.**
- 3 Type a descriptive name in the Display Name field.**  
You will select this name when defining a Notification Group.
- 4 Select a Monitoring Period from the Time Period for Emailing list:**
  - 24x7
  - 9am to 5pm weekdays
  - another Monitoring Period that you have previously created
- 5 Select the Should the Distribution List receive alerts? check box.**
- 6 Configure the type of alerts those on the Distribution List will receive by selecting one or more of the following check boxes:**
  - Alert on Critical  
The user receives an alert when [up.time](#) detects a critical problem with one or more monitored services.

- **Alert on Warning**  
The user receives an alert when **up.time** detects a potential problem with one or more monitored services.
- **Alert on Unknown**  
The user receives an alert when **up.time** detects an error in the configuration of the monitor, or if **up.time** cannot execute the service check.
- **Alert on Recovery**  
The user receives an alert when the service recovers from an error – for example, an application, process or service restarts, or a server reboots.

**7 Click Save.**

## Viewing Distribution Lists

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


To view Distribution Lists, do the following:

- 1 Click Users on the up.time tool bar.**
- 2 In the Tree panel, click View Distribution Lists.**  
A list of Distribution Lists appears in the **Distribution Lists** subpanel.
- 3 Click the name of the Distribution List that you want to view.**  
The details of the group appear in the **Distribution Lists** subpanel.

## Editing Distribution Lists

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

To edit Distribution Lists, do the following:

- 1 Do one of the following:**
  - Click the **Edit** icon (  ) beside the name of the Distribution List.

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

The **Edit Distribution List** window appears.

- 2 **Edit the group as described in “Adding Distribution Lists” on page 396.**



## 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 433 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 433.
- 4 **Optionally, type a description of the group in the Description of Notification Group field.**
- 5 **Select one or more Alert Profiles to apply to the group from the Available Alert Profiles list, then click Add.**
- 6 **Select one or more users to add to the group from the Available Users list, then click Add.**
- 7 **Select one or more Distribution Lists to add to the group from the Available Distribution Lists, then click Add.**
- 8 **Click Save.**

## Viewing Notification Groups

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

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

To view Notification Groups, do the following:

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

A list of Notification Groups appears in the **Notification Groups** subpanel.

- 3 **Click the name of the Notification Group that you want to view.**

The details of the group appear in the **Notification Groups** subpanel.


- 4 **To view the details of an Alert Profile, click the name of the profile.**

## Editing Notification Groups

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

To edit Notification Groups, do the following:

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

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

The **Edit Notification Group** window appears.

- 2 **Edit the group as described in “Adding Notification Groups” on page 399.**

## Changing How Users Are Authenticated

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

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

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



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

## Active Directory Authentication

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

### Enabling Active Directory for Authentication

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

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

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

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

Clicking **Save** switches the authentication source to Active Directory. Administrators still need to create profiles for all [up.time](#) users, but will not need to set a password for each one. See “Adding Users” on page 389 for more information.

### Defining Active Directory Synchronization Mapping

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

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

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

**1 Click Edit Configuration to open the User Authentication Configuration pop-up window.**

**2 Select the Synchronization Enabled check box.**

All user synchronization configuration options appear.

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

By default, synchronization occurs every hour.

**4 In the AD Group Distinguished Name field, enter the name of the AD group of up.time users (e.g., CN=uptime users, CN=Groups, DC=yourdomain, DC=com).**

**5 If required, enter an appropriate administrative AD Username and AD Password required to access the directory.**

**6 In the User Name field, provide the name attribute used to retrieve the user name (e.g., sAMAccountName).**

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

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

- i First Name (e.g., givenName)**
- ii Last Name (e.g., sn)**
- iii Location (e.g., physicalDeliveryOfficeName)**
- iv Email Address (e.g., userPrincipalName)**
- v Pager/Cellphone**
- vi User's Windows Desktop Host Name**
- vii User's Windows Desktop Workgroup**



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

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

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

## LDAP Authentication

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

### Enabling LDAP for User Authentication

To configure [up.time](#) to check an LDAP listing for user passwords, do the following:

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

You will next need to provide access details for the Active Directory server.

- 5 In the LDAP URL field, enter the address for the LDAP server.**

If directory communication occurs through secure channels, such as TLS or SSL, ensure this is reflected in the server address (e.g., “`ldaps://`” instead of “`ldap://`”).

- 6 Enter the LDAP Query that [up.time](#) will use on the LDAP server to look up a user's name.**
- 7 Continue to the next section to enable and configure synchronization from the Active Directory listing to [up.time](#) user profiles. If you do not wish to synchronize users, click Save.**

Clicking **Save** switches the authentication source to the LDAP directory. Administrators still need to create profiles for all [up.time](#) users, but will not need to set a password for each one. See “Adding Users” on page 389 for more information.

## Defining LDAP Synchronization Mapping

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

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

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

- 1 Click Edit Configuration to open the User Authentication Configuration pop-up window.**

- 2 Select the Synchronization Enabled check box.**

All user synchronization configuration options appear.

- 3 In the Synchronize Users field, enter the frequency at which [up.time](#) user information will be synchronized with the LDAP listing.**

By default, synchronization occurs every hour.

- 4 In the LDAP Group Distinguished Name field, enter the name of the LDAP group of [up.time](#) users (e.g., CN=uptime users, CN=Groups, DC=yourdomain, DC=com).**

- 5 If required, enter an appropriate administrative LDAP Username and LDAP Password required to access the directory.**

- 6 In the User Name field, provide the attribute used to retrieve the user name.**

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

- 7 For the remaining Field Mappings, provide attributes for other user details you would like to synchronize with the [up.time](#) user profile:
  - i First Name
  - ii Last Name
  - iii Location
  - iv Email Address
  - v Pager/Cellphone
  - vi User's Windows Desktop Host Name
  - vii User's Windows Desktop Workgroup



Any user attributes chosen to be synchronized with the directory will not be editable in [up.time](#).

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

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

## up.time DataStore Authentication

By default, [up.time](#) uses its own database for password storage and look-up.

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



## Enabling the DataStore for User Authentication

To use [up.time](#) DataStore to store passwords for user authentication, do the following:

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



# CHAPTER 18

## Working with Service Level Agreements

---

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

<i>Overview</i> .....	410
<i>SLAs, Service Monitors, and SLOs</i> .....	411
<i>Viewing Service Level Agreements</i> .....	412
<i>SLA Compliance Calculation</i> .....	415
<i>SLA-Creation Strategies</i> .....	418
<i>Working with SLA Reports</i> .....	422
<i>Adding and Editing SLA Definitions</i> .....	423

## Overview

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

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

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

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

- translate IT infrastructure demands into quantifiable and reportable goals

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

- use trends to anticipate new infrastructure requirements

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

- generate SLA reports that demonstrate compliance and break down objectives

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

## SLAs, Service Monitors, and SLOs

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

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

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

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

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

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

## Viewing Service Level Agreements

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

### Viewing SLA Status

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

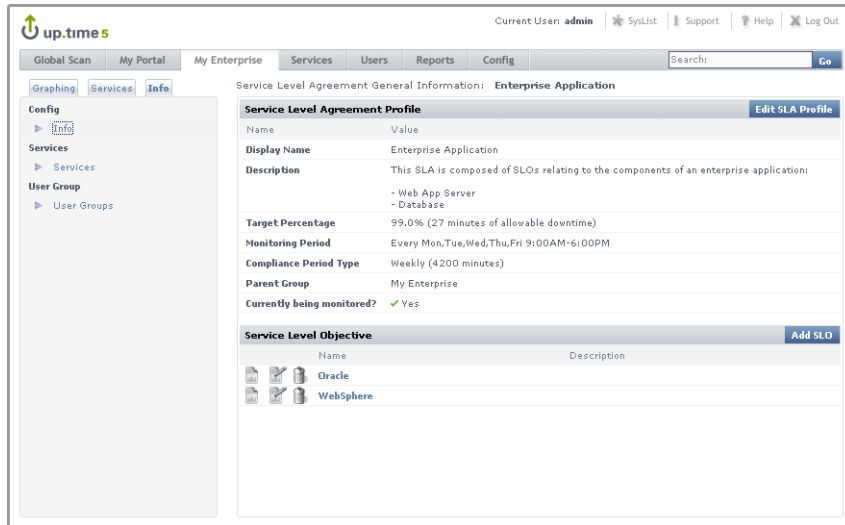


For more information about what kind of SLA information you can view in the **Global Scan** panel, see “Viewing All SLAs” on page 121.

### Viewing SLA Details

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

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



The screenshot shows the up.time 5 web interface. The top navigation bar includes 'Global Scan', 'My Portal', 'My Enterprise', 'Services', 'Users', 'Reports', and 'Config'. The left sidebar has 'Graphing', 'Services', and 'Info' tabs. The main content area is titled 'Service Level Agreement General Information: Enterprise Application'. It features a 'Service Level Agreement Profile' section with the following details:

Name	Value
Display Name	Enterprise Application
Description	This SLA is composed of SLOs relating to the components of an enterprise application: - Web App Server - Database
Target Percentage	99.0% (27 minutes of allowable downtime)
Monitoring Period	Every Mon,Tue,Wed,Thu,Fri 9:00AM-6:00PM
Compliance Period Type	Weekly (4200 minutes)
Parent Group	My Enterprise
Currently being monitored?	<input checked="" type="checkbox"/> Yes

Below this is a 'Service Level Objective' section with a table listing SLOs:

Name	Description
Oracle	
WebSphere	

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

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

For more information about system information in general, see “Viewing System Information” on page 52.

You can view information about the services that make up the SLA by clicking the **Services** tab in the Tree panel. The options available in the Tree panel are summarized in “Viewing Service Information” on page 54.

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

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

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

See “A Note About SLOs and Compliance” on page 417 for more information about SLOs and the Achieving statistic.



## SLA Compliance Calculation

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

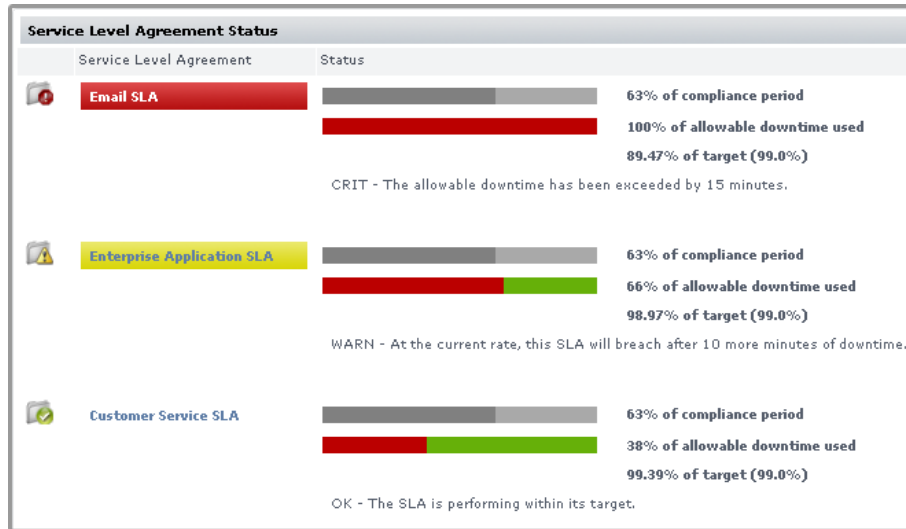
For example, consider an SLA whose compliance period type is weekly and its Monitoring Period is Monday through Friday, 9 p.m. to 5 p.m. The Monitoring Period consists of five eight-hour days—in other words, 40 hours, or 2400 minutes. If the SLA's target is 95%, it has 120 minutes of allowable downtime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

### Reporting SLA Status

An SLA's reported status in the **Global Scan** panel includes the following in the form of progress bars: the percentage of the Monitoring Period that has expired, and the percentage of allowable downtime consumed during the Monitoring Period. (See “Viewing All SLAs” on page 121 for information about SLA information in the **Global Scan** panel.)

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


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



## Handling Simultaneous Service Downtime

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

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



## A Note About SLOs and Compliance

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

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

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



See “Viewing SLA Details” on page 412 for information on how to find information such as the Achieving statistic in an SLA summary.

## SLA-Creation Strategies

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

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

### Setting Up and Gathering Data for Monitors

Determine which service monitors will best reflect the end-user experience, based on the aspect of your infrastructure that your SLA will cover. See “SLAs, Service Monitors, and SLOs” on page 411 for some sample SLAs and objectives.

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


### Identifying Outages and Improvable Performance

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

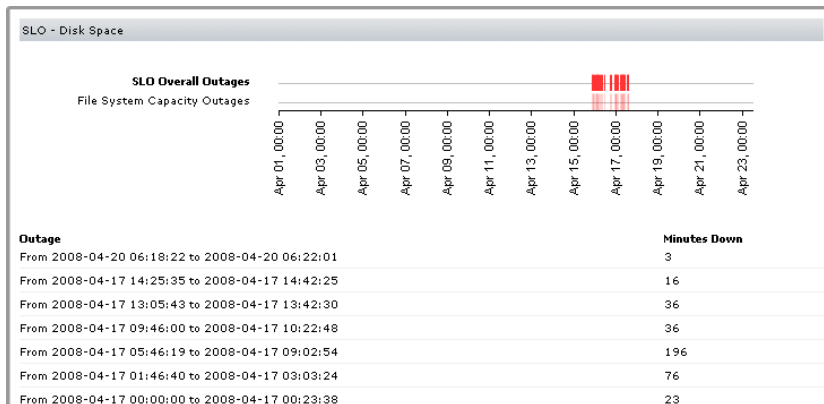
Consider a sample SLA called System Performance that is meant to ensure your application servers are not experiencing excessive loads; this can be indicated by CPU usage and disk space. The first service level objective is

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


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



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



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





## Developing Baselines

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

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

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



## **Working with SLA Reports**


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

See “Reports for Service Level Agreements” on page 505 for more information.



## Adding and Editing SLA Definitions

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

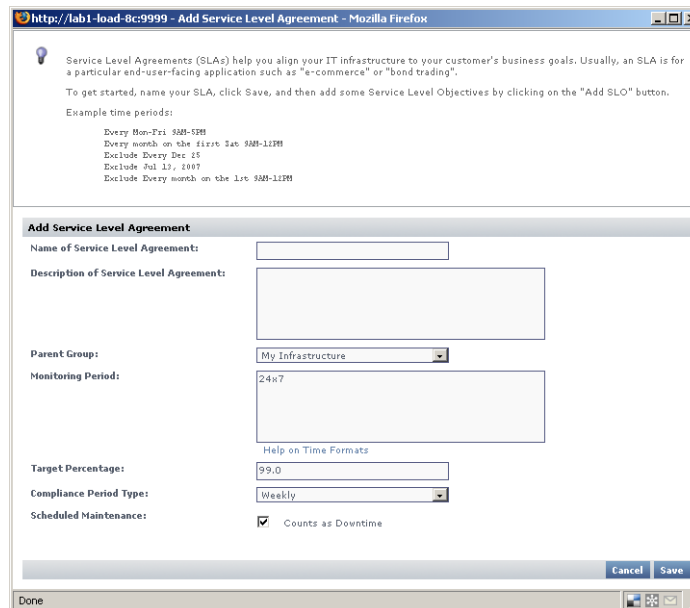
 When you create an SLA, it will be inserted into the current compliance period. For example, a newly created SLA that reports over a monthly compliance period will, if created on the 15th of the month, already be around 50% through the period.

### Adding a Service Level Agreement

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

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

The Add Service Level Agreement window appears:



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

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

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

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

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

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

For more information on groups, see “Working with Groups” on page 108.

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

You will need to create a time period definition (e.g., “Every Mon-Sat 8AM-6PM”). See “Monitoring Periods” on page 449 and “Time Period Definitions” on page 677 for more information.

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

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

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

- 9 Click Save.**

Once saved, the SLA’s **Service Level Agreement General Information** subpanel is displayed (see “Viewing SLA Details” on page 412 for more information). From this page, you can add SLOs, as well as associate Alert Profiles and Action Profiles to the SLA.

## Adding Service Level Objectives to an SLA

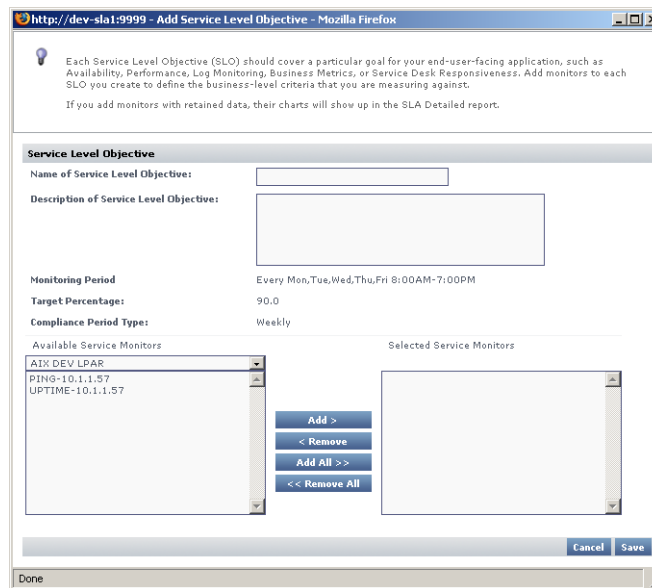
To add a service level objective to an SLA, do the following:

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

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

- 2 Click **Add SLO**.

The **Add Service Level Objective** window appears:



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

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

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

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

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

## **Associating Alert and Action Profiles to an SLA**

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

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


## **Editing SLA and SLO Definitions**

To edit a service level agreement, do the following:

- 1 **In the My Infrastructure panel, right-click the name of the Service Level Agreement that you want to modify, then click Edit.**  
The **Edit Service Level Agreement** window appears.
- 2 **Edit the SLA as described in the previous section.**  
See “Adding a Service Level Agreement” on page 423 for information.

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

- Monitoring Period
- target percentage
- compliance period type

 Any changes made are immediately reflected in any SLA reporting.

To edit a service level objective, do the following:

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


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

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

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

See “Adding Service Level Objectives to an SLA” on page 425 for information.

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

 Any changes made are immediately reflected in any SLA reporting.



# CHAPTER 19

## Alerts and Actions

---

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

<i>Understanding Alerts</i> .....	430
<i>Alert Profiles</i> .....	433
<i>Working with Custom Alert Formats</i> .....	437
<i>Action Profiles</i> .....	441
<i>Monitoring Periods</i> .....	449

## Understanding Alerts

When a problem occurs at a Datacenter, Application, or SLA, 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 22 for more information)
- the name of the service that is experiencing the problem
- the current state of the service
- any output from the monitor

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

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

`up.time` can send alerts via:

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

The following is a sample email alert:

```
Notification type: Problem
1/12/2008 10:52
Host: 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 151.

## Understanding the Alert Flow

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

Considering the following example:


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

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

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


## Alerts and Actions *Understanding Alerts*

This alert flow is illustrated in the following diagram:



## Alert Profiles

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



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

## Enabling the Windows Messaging Service

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

To enable the Windows messaging service, do the following:

- 1 **In Windows, select Start > Control Panel.**
- 2 **In the Control Panel, double click Administrative Tools, and then double click Services.**  
The **Services** window appears.
- 3 **Find and then double click Messenger in the list of services.**  
The **Messenger Properties** dialog box appears.
- 4 **In the Messenger Properties dialog box, select Automatic from the Startup type dropdown list.**
- 5 **Click Apply.**

## Creating Alert Profiles

To create Alert Profiles, do the following:

- 1 On the [up.time](#) tool bar, click **Services**.
- 2 In the Tree panel, click **Add Alert Profile**.  
The **Add Alert Profile** window appears.
- 3 Type a descriptive name for the profile in the **Name of Alert Profile** field.
- 4 In the **Start alerting on notification number** field, enter the number of times an error must occur before [up.time](#) sends an alert notification.
- 5 Enter the number of times to re-send the notification in the **End alerting on notification number** field.

Optionally, click the **Never Stop Notifying** option to have [up.time](#) continually send notifications.

- 6 **Select one of the following notification options:**

- **Email Alert**  
Sends the alert to the email addresses of the members of a Notification Group.
- **Pager Alert**  
Sends the alert to the pagers of the members of a Notification Group.
- **Script Alert**  
Uses a script to send the alert via SMS to the mobile phones of the members of a Notification Group.

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

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

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

- Windows Popup Alert

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

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

## Viewing Alert Profiles

To view Alert Profiles, do the following:

- 1 **On the up.time tool bar, click 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 434.**

## Associating Alert Profiles to Elements

You can associate an Alert Profile to any Service Monitor, Application, or SLA if their state changes from OK to Warning or Critical. Alert Profiles are normally associated with any of these monitored items at the time of their configuration; Alert Profile associations can also be modified with existing service monitor definitions.

See Chapter 8, “Using Service Monitors”, “Working with Applications” on page 103, and “Adding and Editing SLA Definitions” on page 423 for more information about configuring Service Monitors, Applications, and SLAs, respectively.

## Working with Custom Alert Formats

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

To define a custom alert format, do the following:


- 1 Define an Alert Profile, as described on page 434.**
- 2 In the Custom Format Options section, click Custom Formats.**
- 3 From the dropdown list, select one of the following options:**
  - **Small Template**

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

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

Contains the information in the medium template, as well as the status of the host.
- 4 Click Fill.**

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



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

## Custom Alert Format Variables

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

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

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

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



Variable	Definition
\$DATETIME\$	The date and time at which the alert was generated. This appears in the subject line of the message.
\$SERVICENAME\$	The name of the service, along with the name of the host for which the alert was generated. For example, if the alert was generated by the ping check for the server MailHub, then PING-MailHub appears in the alert. This appears in the subject line of the message.
\$SERVICESTATE\$	One of the following: <ul style="list-style-type: none"> <li>• OK</li> <li>• WARN</li> <li>• CRIT</li> <li>• MAINT</li> <li>• UNKNOWN</li> </ul> This appears in the subject line of the message.
\$DATE\$	The date on which the alert was generated.
\$TIME\$	The time at which the alert was generated.
\$HOSTNAME\$	The name of the host (as saved in <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>

<b>Variable</b>	<b>Definition</b>
\$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

## Action Profiles

Action Profiles are templates that direct **up.time** when it encounters a problem on a monitored system. You can associate an Action Profile to any Service Monitor, Application, or SLA if their state changes from OK to Warning or Critical. Action Profiles are normally associated with any of these monitored Elements at the time of their configuration; Action Profile associations can also be changed when you are modifying existing service monitor definitions.

See Chapter 8, “Using Service Monitors”, “Working with Applications” on page 103, and “Adding and Editing SLA Definitions” on page 423 for more information about configuring Service Monitors, Applications, and SLAs, respectively.

Actions include one of the following tasks:

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

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

## VMware vCenter Orchestrator Workflow Actions

If an administrator has integrated **up.time** with VMware vCenter Orchestrator (see “VMware vCenter Orchestrator Integration” on page 649, you can configure Action Profiles to initiate Orchestrator workflows.

Orchestrator is a VMware vCenter Server add-on that allows its administrators to create workflows that automate vCenter management tasks. These Orchestrator workflows are open ended: all vCenter actions are available for automation through the processing of parameters and

runtime arguments. **up.time** Action Profiles can be configured to provide input parameters to specific workflows, thus integrating vCenter management with **up.time**'s monitoring and alerting capabilities.

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

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

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

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

### Orchestrator Input Parameter Variables

When configuring a VMware vCenter Orchestrator workflow, you have at your disposal a set of **up.time**-specific variables that can be entered as parameter variables, and whose ensuing runtime values will be passed to the Orchestrator workflow during execution. The variables available to you are those that are used when creating a custom alert format. See "Custom Alert Format Variables" on page 438 for information.

## SNMP Trap Actions

You can also configure an Action Profile to send an SNMP trap to a particular host. An SNMP trap is notification that is issued by a system that is running SNMP when a problem occurs. The host to which the SNMP trap is sent must be running an SNMP trap listener.

If you use SNMP traps, the trap message will be sent in the format specified by the [up.time](#) MIB. This MIB is found in the `scripts` directory. The uptime software enterprise OID is `.1.3.6.1.4.1.24216`.

## Creating Action Profiles

To create Action Profiles, do the following:

- 1 On the [up.time](#) tool bar, click Services.**
- 2 In the Tree panel, click Add Action Profile.**  
The **Add Action Profile** window appears.
- 3 Enter a name for this profile in the Name of Action Profile field.**
- 4 Specify the number of times an error must occur before [up.time](#) sends a notification in the Start action on notification number field.**
- 5 Specify the number of times action will be carried out in the End action on notification number field.**

Optionally, select the **Never Stop Notifying** option to continually carry out the action in this profile until the problem is resolved.

- 6 If VMware vCenter Orchestrator integration has been enabled, and you would like the Action Profile to drive an Orchestrator workflow, do the following:**

- i In the Select Workflow field, input a workflow to configure.**

You can either scroll through and select the workflow from the drop-down list, or begin typing the workflow's name.

- ii Click Get Parameters.**

[up.time](#) will retrieve information from the Orchestrator server and dynamically display configuration fields for the chosen workflow's input parameters.

- iii Configure the input parameter fields for the workflow.**

For information on the specific configuration parameters available for the chosen workflow, consult the appropriate developer's documentation.

- 4 If you would like the Action Profile to write to a log, in the Log File field, enter the name and path to a log file on the Monitoring Station to which error information will be written.**
- 5 If you would like the Action Profile to run a recovery script, in the Recovery Script field, enter the name and path to a script that will reboot a server, or restart an application, process, or service.**

The recovery script will also have the following information appended to it:

- the date and time on which the error occurred
- the type of error notification that was sent
- the name of the host on which the error occurred
- the state of the host
- the name of the service that threw the error
- the state of the service
- the output that was generated by the error

For example:

```
"/usr/local/uptime/recover.sh" "24/12/2007 5:01:05"  
"Problem" "printserver" "null" "WinSrv-Print Spooler"  
"CRIT/threshold error" "servicestatus: Not Running does  
not match Running (Service 'Print Spooler' found, status:  
Not Running, took 12ms)"
```



You can also use the recovery script to file trouble tickets with a system like Remedy, or to interact with third party software packages.

- 6 If you are setting up an Action Profile for a Windows server, you can also leave the Windows Service as Agent, and complete the following fields:**
  - Windows Host  
The name of the host on which the service is running.
  - Agent Port

The port on which the **up.time** agent that is installed on the system is listening. The default is 9998.

- **Use SSL**  
Select this option if **up.time** will securely communicate with the host using SSL (Secure Sockets Layer).
- **Agent Password**  
Enter the password that is required to access the agent that is running on the system that is being monitored. For information on setting the agent password, see the uptime software Knowledge Base article entitled, “What is the password for the Windows agent?”
- **Windows Service**  
The name of the specific Windows service to which the Action Profile will apply.
- **Action**  
Select one of the following actions:
  - None
  - Start
  - Stop
  - Restart

**7 If you are setting up an Action Profile for a Windows server that is using a WMI implementation, you can also select the Windows Service as WMI, and complete the following fields:**

- **WMI Host:**  
The name of the host on which the service is running.
- **Windows Domain:**  
The Windows domain in which WMI has been implemented.
- **Username:**  
The name of the account with access to WMI on the Windows domain.

- Password:  
The password for the account with access to WMI on the windows domain.
- Windows Service  
The name of the specific Windows service to which the Action Profile will apply.
- Action  
Select one of the following actions:
  - None
  - Start
  - Stop
  - Restart

**8 If you want to send SNMP traps to a particular host, complete the following fields:**

- SNMP Trap Host  
The name of the host that monitors SNMP traps.
- SNMP Trap Port  
The port number on the trap host to which the SNMP trap is sent.
- SNMP Trap Community  
The name which acts as a password for sending trap notifications to the trap host.
- SNMP Trap OID (optional)  
The object identifier (OID) that identifies the SNMP trap – for example, `.1.3.6.1.2.1.34.4.1.7`.

**9 If Splunk integration has been enabled, and you would like the Action Profile to write to the Splunk log, complete the following fields:**

- Splunk Hostname  
The host name of the server on which Splunk is running.



- Logging Port

The port on which the Splunk server is listening for logging requests. This port is configured in Splunk, and you will need to contact the Splunk administrator for this information.

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

For more information on Splunk integration, see “Splunk Settings” on page 653.

- 10 **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, displaying the settings that you configured when you created the profile, as well as a list of the services that are attached to the profile.

- 3 **To test whether or not the profile works, click the Test Action Profile button.**

A popup window appears, and the 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 443.**

## 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 Services.**
- 2 In the Tree panel, click Add Monitoring Period.**  
The **Add Monitoring Periods** window appears.
- 3 Type a name in the Monitoring Period Name field.**
- 4 In the Definition section, enter one or more time period expressions that combine to create a full Monitoring Period definition.**

See “Time Period Definitions” on page 677 for information on the types of time period expressions that are valid in up.time.

- 5 Click Save.**



# CHAPTER 20

## 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> .....	452
<i>Generating Reports</i> .....	453
<i>Saving Reports</i> .....	456
<i>Scheduling Reports</i> .....	459
<i>The Report Log</i> .....	462

### 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 465.

## 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 459.



If you do not receive a scheduled report, check the Report Log (see “The Report Log” on page 462) or contact your system administrator.

## Report Generation Options

up.time can generate reports in four ways:

- **Print to Screen**  
Displays the report in a new window. This is the default option.
- **PDF to Screen**  
Converts the report to a PDF document, and displays it in a new window. You can save the PDF document to a local or network drive, or print it.
- **XML to Screen**  
Displays the report, as an unformatted XML document, in a new window.
- **Email Address**  
Enables you to email the report, as a PDF document attached to an email message, to:
  - A specific up.time user, for example a system administrator.  
Click **User** and then select the name of an up.time user to whom you want to send the report from the dropdown list.
  - The members of one or more up.time user groups.  
Click **Group** and then select the name of an up.time user group to which you want to send the report from the dropdown list.
  - One or more email addresses.  
Click the **Email Address** option, and then type the email address of the person to whom you want to send the report in the field. To send the report to multiple recipients, type their email addresses in the field separated by commas or semi-colons. For example:

The screenshot shows a dialog box titled "Generate Now" with four buttons: "Email", "Print to Screen", "PDF to Screen", and "XML to Screen". Below the buttons are three radio button options: "User:", "Group:", and "E-mail Address:". The "E-mail Address:" option is selected. Below these options is a text input field labeled "Email address:" containing the text "headadmin@myorganization.com, fixit@myorganization.com,".



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 459 for more information.

To save reports, do the following:

- 1 In the Save Report area of the Report subpanel, select one of the following options:**
  - HTML
  - PDF
  - XML
  - Email
- 2 If you selected Email in step 1, specify one of the email options.**
- 3 Type a name for the report in the Save to My Portal As field.**
- 4 Optionally, type a description for the report in the Report Description field.**
- 5 Click Save Report.**

## Saving Reports to the File System

You can save reports to the file system of a server in your environment so others in your organization can view the reports. You can, for example, save a report to a Web server for viewing on your Intranet. The reports are saved as either PDF or HTML files. The system administrator can specify the

directory on the server in which reports will be saved by adding the following entry to the file `uptime.conf`:

```
publishedReportRoot=<directory_name>
```

Where `<directory_name>` the directory into which **up.time** will write reports – for example, `C:/Program Files/uptime software/uptime/`. The report files are saved to a subdirectory named `GUI/published`. You need permissions to write to the `published` directory.

**up.time** automatically names each report file. The file name contains the following information:

- name of the report, taken from the **My Portal** panel
- date on which the report was run
- user name of the person who ran the report

The following is an example of a report file name:

```
Service Outages_2006-01-24_rfripp.pdf
```

To save reports to a file system, do the following:

- 1 **In the Save Report area of the Report subpanel, enter a name for the report in the Save to My Portal As field.**
- 2 **Optionally, enter a description of the report in the Description field.**
- 3 **Select either HTML or PDF from the list of options.**
- 4 **Click the Publish Report option.**
- 5 **Click the Scheduled Report option, and then select a date and time for the report to run.**

For more information on scheduling reports, see “Scheduling Reports” on page 459.

- 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

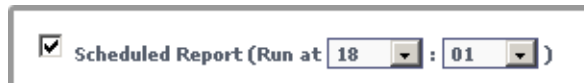
If you need to run a report at a particular interval – for example, daily or weekly – you can schedule when the report should be generated. [up.time](#) generates the report and emails it to a user or group of users.

For example, you generate a File System Capacity Growth Report – which charts the amount of disk usage for a system. However, the system for which you are generating the report schedules backups from midnight to 4:00 a.m. Due to the gap caused by the backup, the CPU usage and disk activity statistics are not indicative of the overall system load. You can specify that the report does not cover the periods of time over which the backups occur.

To schedule reports, do the following:

- 1 In the Reports subpanel, select the Email option in the Save Report section of the subpanel, and then select one of the following options:**
  - User
  - Group
  - E-mail Address
- 2 Type a name for the report in the Save to My Portal As field.**
- 3 Optionally, type a description for the report in the Report Description field.**
- 4 Click the Scheduled Reports checkbox, and then select the time at which to run the report from the dropdown lists.**

For example, to run the report at 3:30 p.m., select 15 from the first dropdown list and 30 from the second dropdown list, as shown below:



Scheduled Report (Run at 18 : 01 )

5 Select one of the following options:

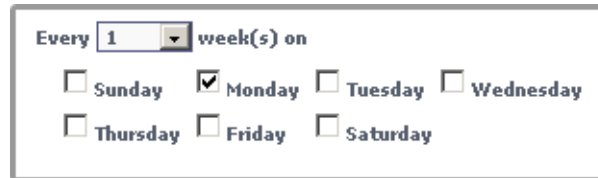
- Daily



Every 1 day(s)  
 Every Weekday

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



Every 1 week(s) on  
 Sunday  Monday  Tuesday  Wednesday  
 Thursday  Friday  Saturday

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



Day 1 of every 1 month(s)  
 The first Sunday of every 1 month(s)

Do one of the following:

- Select the **Day** option. From the first dropdown list, select the day (from 1 to 31) on which to run the report. Then, select the month (from 1 to 12) during which to run the report.

For example, if you select 3 and 7 from the dropdown lists, the report will be run on the third day of every seventh month.

- Select the second option, then do the following:
  - select first, second, third, fourth, or last from the first dropdown list
  - select a day of the week on which the report will run from the second dropdown list
  - select a number from 1 to 12 from the third dropdown list

For example, if you select *second*, *Tuesday*, and *9* from the dropdown lists, the report will be run on the second Tuesday of every ninth month.



If you are saving an existing report after editing it or saving a new report with the name of an existing one, *up.time* displays a warning dialog box. Click **OK** on the dialog box to overwrite the report. Or, click **Cancel** on dialog box to give the report a different name.

## The Report Log

The Report Log tracks the progress and status of scheduled reports, or reports that are running in the background. Using the Report Log, you can quickly determine whether or not reports have been successfully generated. If they have not, then you can use the log to determine why report generation failed.

The **Report Log** subpanel tracks the status of reports in the following sections:

- Pending Reports

Reports that are in the report queue, and are waiting to run. This section contains the following information:


- the name of the report
- the description of the report, if available
- whether or not the report is scheduled
- the date and time on which the report will be run

The following image illustrates the **Pending Reports** section:

Pending Reports			
Report Name	Report Description	Scheduled?	Next Run Time
 ReportEnterpriseCPUUtilization		✗ No	Apr 17, 2008 16:48
 ReportCPUUtilizationSummary		✗ No	Apr 17, 2008 16:49
 ReportFileSystemServiceTime		✗ No	Apr 17, 2008 16:50

- 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
 ReportResourceUsage		✗ No	Apr 17, 2008 16:48

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
 ReportSlaDetailed	2008-04-03 17:07:36.0	2008-04-03 17:09:14.0	finished	Executed successfully
 ReportSlaSummary	2008-04-03 17:09:14.0	2008-04-03 17:09:22.0	finished	Executed successfully
 ReportSlaSummary	2008-04-03 17:09:22.0	2008-04-03 17:09:32.0	finished	Executed successfully
 ReportVmwareWorkload	2008-04-17 15:27:02.0	2008-04-17 15:27:22.0	finished	Executed successfully

## 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 21

## Using Reports

---

This chapter describes the reporting features of [up.time](#) in the following sections:

<i>Reports for Performance and Analysis</i> .....	466
<i>Reports for Capacity Planning</i> .....	480
<i>Reports for Service Level Agreements</i> .....	480
<i>Reports for Availability</i> .....	508
<i>Reports for J2EE Applications</i> .....	515
<i>Reports for Virtual Environments</i> .....	522

## 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 23.

### 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.



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

- CPU Performance Graph

Tracks the performance of a system's CPU over a specified time period. This information appears as a graph in the report.

- TCP Retransmits

Any network services that may not be completing properly because of undue network or system load. This information appears as a graph in the report.

- Disk Statistics

The following statistics for each disk on a system:

- percentage of the disk that is busy
- average queue length
- number of reads and writes per second
- number of blocks being accessed per second
- average wait time, in seconds
- average service time, in seconds



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 111 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 454 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 456 and “Scheduling Reports” on page 459 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.

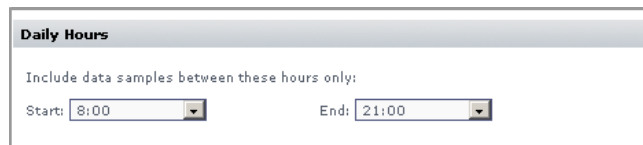
### Creating a Multi-System CPU Report

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 23.

- 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 panel titled "Daily Hours" with a subtitle "Include data samples between these hours only:". Below the subtitle are two dropdown menus. The "Start:" dropdown is set to "8:00" and the "End:" dropdown is set to "21:00".

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 111 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 454 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 456 and “Scheduling Reports” on page 459 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 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 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  $USR \%$

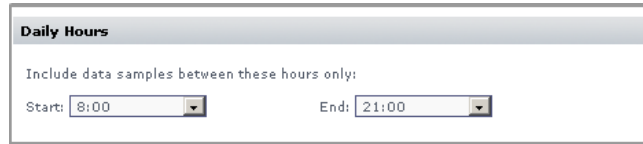
### Creating a CPU Utilization Ratio Report

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 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 for more information.

## Wait I/O Report

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

The Wait I/O report contains the following information:

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

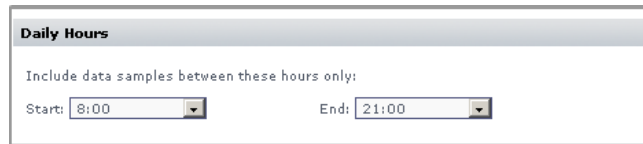
### Creating a Wait I/O Report

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

- 1 In the Reports Tree panel, click Wait I/O.**
- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 window titled "Daily Hours". Inside the window, there is a label "Include data samples between these hours only:". Below this label, there are two dropdown menus. The first dropdown menu is labeled "Start:" and has "8:00" selected. The second dropdown menu is labeled "End:" and has "21:00" selected.

For example, if you want to report to cover the hours from 1:00 a.m. to 1:00 p.m., select 1:00 from the **Start** dropdown list, and 13:00 from the **End** dropdown list.

- 4 Optionally, enter a value in the Highlight average WIO over threshold field.**

Any system with an average Wait I/O percentage that exceeds the value that you enter in this field will be highlighted in red in the report. As well, the following text appears in the header of the report:

```
Systems with an Average Wait I/O over x.x% are highlighted
```

Where x.x is the percentage that you entered in this field.

- 5 If you want to generate reports for systems in specific groups, select the groups from the List of Groups area.**

- 6 To generate reports for one or more views, select the groups from the List of Views area.**  
See “Working with Views” on page 111 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 454 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 459 for more information on configuring a scheduled report.

## Service Monitor Metrics Report

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

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

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

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

## Creating Service Monitor Metrics Reports

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

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

For more information, see “Understanding Dates and Times” on page 23.

- 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 111 for more information about views.

- 5 If you are generating reports for specific systems in your environment, select them from the **List of Elements**.
- 6 Click **Go** to page 2.

A table containing the current retained service metrics appears in the **Service Metrics** subpanel.

- 7 Click the checkboxes in the **Select** column to select the variables on which you want to report as shown below:

Current Retained Service Metrics					
Host	Instance Name	Instance Description	<input type="checkbox"/> Select	Variable	Units
WebSphere(lab-websphere51)	WebSphere		<input type="checkbox"/>	Connection Pool Average Time ▾	ms
			<input type="checkbox"/>	Connection Pool Closed ▾	
			<input type="checkbox"/>	Connection Pool Created ▾	
			<input type="checkbox"/>	Connection Pool Percent Used ▾	%
			<input type="checkbox"/>	Connection Pool Size ▾	

- 8 **Optionally, select one of the following:**

- Show all non-ranged metrics on one chart

This option combines all of the variables you selected in one chart. Any ranged metrics will appear in their own charts.



- Display charts as stacked area

Each chart in the report will have two or more data series stacked on top of each other, rather than the line graph that usually appears in the report.

**9 To save the report, do the following:**

- Enter a name for the report in the **Save to My Portal As** field.
- Optionally, enter text in the **Description** field.
- Click **Save Report**.

The report parameters are saved to the **My Portal** panel. Doing this does not generate the report.

**10 To schedule the saved report to run at a specific time or interval, click the Scheduled checkbox.**

See “Saving Reports” on page 456 and “Scheduling Reports” on page 459 for more information.

## Reports for Capacity Planning

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

- Enterprise CPU Utilization Report
- File System Capacity Growth Report
- Server Virtualization Report
- Solaris Mutex Exception Report
- Network Bandwidth Report
- Disk I/O Bandwidth Report
- CPU Run Queue Threshold Report
- File System Service Time Summary Report

### Enterprise CPU Utilization Report

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

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

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

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

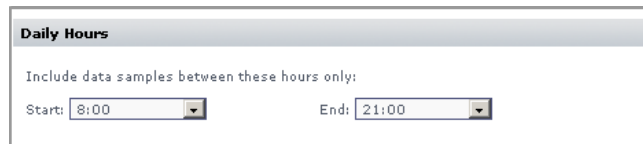
## Creating an Enterprise CPU Utilization Report

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

- 1 **In the Reports Tree panel, click Enterprise CPU Utilization.**
- 2 **In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 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 23.

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 111 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 454 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 456 and “Scheduling Reports” on page 459 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 23.

- 3 **Click the Display Element custom fields option to insert the content of the custom fields in the system profile into the report.**

The custom fields contain additional information about the system – for example, the types of reports that should be run on this system or when maintenance is scheduled. For more information, see page 101.

- 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 111 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 454 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 459 for more information on configuring a scheduled report.

## Using the Server Virtualization Report

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

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

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

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

## Solaris Mutex Exception Report

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

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

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

### Creating a Solaris Mutex Exception Report

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

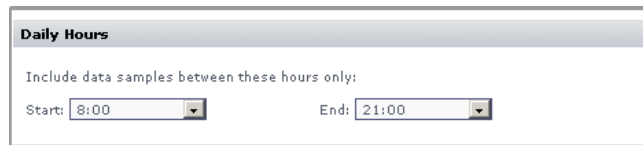
- 1 In the Reports Tree panel, click Solaris Mutex Exception.**

- 2 **In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

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

- 3 **If you want the report to only include data from certain hours during the day, select those hours from the dropdown lists in the Daily Hours section, as shown below:**



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

- 4 **Optionally, enter a value in the Highlight average SMTX over threshold field.**

If the number of mutex stalls for a system, averaged for all of its CPUs over the defined reporting time period, exceeds the value in this field, the number will be highlighted in the report. For example, if you enter 75 and a server returns 93, that value is highlighted.

- 5 **If you want to generate reports for groups of systems, select the groups from the List of Groups area.**

- 6 **To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

- 7 **If you are generating reports for specific Applications in your environment, select them from the List of Elements.**



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

- 8 **Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

### Using the Solaris Mutex Exception Report

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



**Solaris Mutex Exception**  
Date Range: 2008-04-16 00:00:00 to 2008-04-16 17:55:11 between the time range: 00:00 to 23:59  
Multi CPU systems with an Average SMTX over 150.0 are highlighted

System Name	Number of CPUs	Average SMTX
Opteron (opteron)	2	<b>158.19</b>
The Zones (vmh-t1k-a)	24	0.83

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

- Multi-CPU Usage (see page 551 for more information)
- Run Queue Length (see page 549 for more information)
- Run Queue Occupancy (see page 549 for more information)

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

## Network Bandwidth Report

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

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

- AIX
- FreeBSD
- IRIX
- MacOS
- Novell NRM

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

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

The report contains the following information:

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

## Generating a Network Bandwidth Report

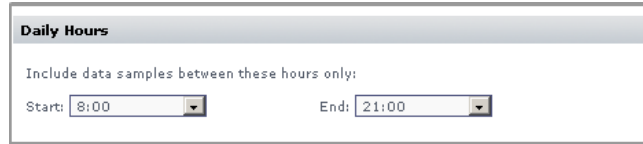
To generate a Network Bandwidth report, do the following:

- 1 In the Reports Tree panel, click Network Bandwidth.**
- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

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:**



Daily Hours

Include data samples between these hours only:

Start: 8:00 End: 21:00

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

- 4 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 111 for more information about views.

- 7 If you are generating reports for specific Applications in your environment, select them from the List of Elements.**

- 8 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

## Using the Network Bandwidth Report

The following is an example of a Network Bandwidth report:

Network Bandwidth Report			
Date Range: 2008-04-17 00:00:00 to 2008-04-17 10:59:19 between the time range: 00:00 to 23:59			
Bytes per packet: 750			
Hostname	Interface	Total MB In	Total MB Out
AIX DEV LPAR (10.1.1.57)	en0	2,186.97	2,125.59
AIX QA LPAR (10.1.1.56)	en2	417.78	336.75
AIX5 (aix51)	(unknown)	0.00	0.00
	en0	223.55	126.35
ELinux (elinux)	eth1	0.00	0.00
	sit0	0.00	0.00
	eth0	14.49	11.50
ESX4 (vmh-esx4)	vmnic0	1,311.36	5,709.02
	vmnic1	1,422.36	435.30
ESX7 (vmh-esx7)	vmnic0	2,801.19	3,546.28
	vmnic1	0.00	84.88
Exchange (uptime-exchange)	netif0	9.66	9.66
	netif1	362.76	454.91
vmh-prod	vmnic0	55,057.18	77,215.26
	eth0	14,756.37	7,039.90
WebSphere (lab-websphere51)	netif0	932.94	932.94
	netif1	24.36	124.54

In this example, the system Filter has high levels of network traffic flowing in and out of a particular network interface. Based on this information, you can generate a Network graph (see page 567 for more information) to get a better idea of why network I/O is so high on the system.

## Disk I/O Bandwidth Report

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

The report contains the following information:

- the display name of the system in [up.time](#)
- the names of each disk on the system

- where applicable, the name of the file system on the disk
- the total amount of data, measured in megabytes, that is being read from and written to the disk

### Using Regular Expressions

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

Exclude Disks: <input type="text"/>	Exceptions: <input type="text"/>
Note: you can enter regular expressions into these fields.	
Exclude File Systems: <input type="text"/>	Exceptions: <input type="text"/>
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**.

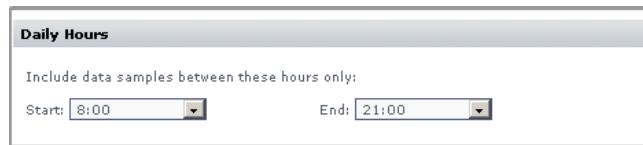


- 2 In the **Date and Time Range** area, select the dates and times on which to report.

For more information, see “Understanding Dates and Times” on page 23.

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 494 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 494 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 111 for more information about views.

- 9 If you are generating reports for specific Applications in your environment, select them from the List of Elements.**
- 10 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

### Using the Disk I/O Bandwidth Report

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

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

In this example, the systems Brightmail and Weblogic Server have high levels of disk I/O. Based on this information, you can generate a Disk Performance Statistics graph (see page 570 for more information) to get a better idea of why disk I/O is so high on the system.

## CPU Run Queue Threshold Report

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

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

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

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

This report contains the following information:

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

### Generating a CPU Run Queue Threshold Report

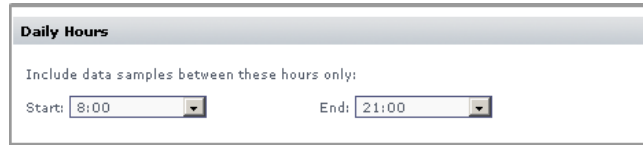
To generate a CPU Run Queue Threshold report, do the following:

- 1 In the Reports Tree panel, click CPU Run Queue Threshold.**
- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

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 the window, there is a label "Include data samples between these hours only:". Below this label, there are two dropdown menus. The first is labeled "Start:" and has "8:00" selected. The second is labeled "End:" and has "21:00" selected.

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

- 4 In the Max CPU (%) field, specify the threshold for CPU usage.**

CPU usage is considered critical when both the CPU usage and the length of the run queue exceed this threshold.

- 5 In the Threshold field, enter the number of queued up jobs that, when exceeded, is considered excessive.**

Multiple CPUs are taken into account so that the defined threshold scales up with each additional CPU present on a monitored system.

- 6 Select any of the following statistics to include in the report:**


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

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

- 7 If you want to include a list of processes that are in the run queue in the report, click Show Processes.**

- 8 **Click the Maintain Graph Scale option to keep the scale of the graphs in the reports consistent.**

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



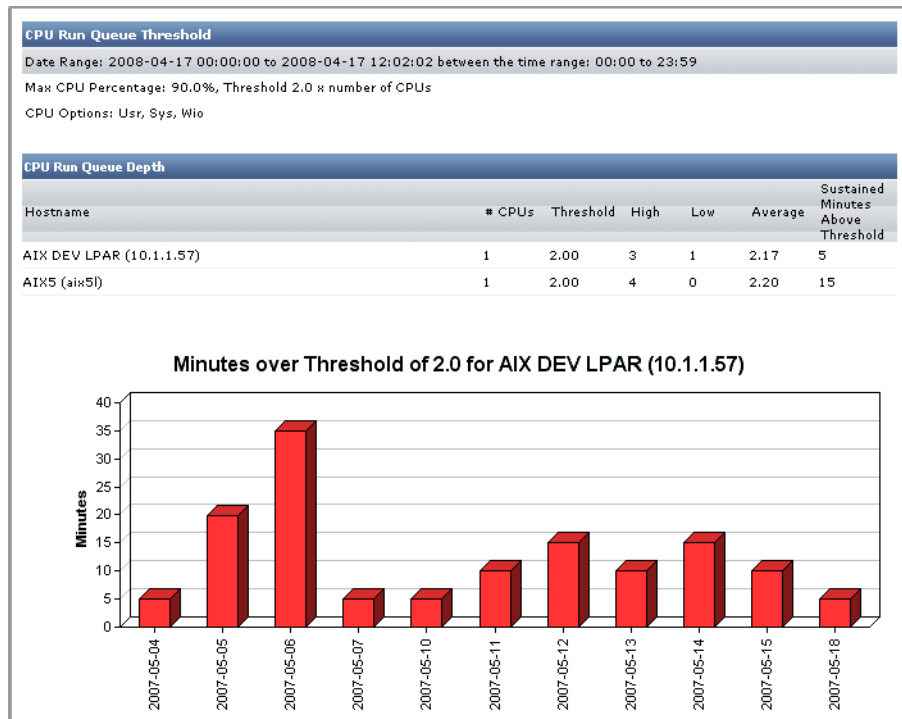
- 9 **To generate reports for groups of systems, select the groups from the List of Groups area.**
- 10 **To generate reports for one or more views, select the groups from the List of Views area.**  
See “Working with Views” on page 111 for more information about views.
- 11 **If you are generating reports for specific Applications in your environment, select them from the List of Elements.**
- 12 **Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 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 547 for more information) to get a better idea of why the system is exceeding the CPU run queue threshold.

## File System Service Time Summary Report

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

The report contains the following information:

- the name of the systems for which the report has been generated
- the names of the disks and file systems on the system
- the high, low, and average service times for each disk or file system, measured in milliseconds
- the  $n^{\text{th}}$  percentile for each disk or file system (e.g., although a file system may have had a high service time of 100ms, its 95<sup>th</sup> percentile of 40ms means 95% of the service times were 40ms or lower)

On a system with heavy disk usage, disks and file systems will be in the higher end of the percentile.

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

### Generating a File System Service Time Summary Report

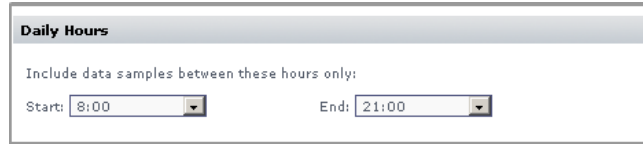
To generate a File System Service Time Summary report, do the following:

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

For more information, see “Understanding Dates and Times” on page 23.

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:**



**Daily Hours**

Include data samples between these hours only:

Start: 8:00 End: 21:00

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

- 4 Select one of the following options from the Primary Sort by dropdown list to sort the results that [up.time](#) returns:**
  - System Name
  - Disk
  - High Service Time (the default)
  - Low Service Time
  - Average Service Time
  - High Percentile
- 5 Select Ascending or Descending from the associated dropdown list.**
- 6 Optionally, do the following:**
  - Select another sort criteria from the **Secondary Sort by** dropdown list.
  - Select **Ascending** or **Descending** from the associated dropdown list.
- 7 In the Threshold field, specify the threshold for file system service time.**

Disk or file system service time is considered critical when it exceeds this threshold.



- 8 In the Percentile field, specify the percentage of time at which the service time for systems is below the threshold.**

The default is 95, which is the lowest service time that is greater than at least 95% of all of the recorded values in the time range that you specified in step 2.

- 9 If you want to include or exclude certain disks, enter the following in the Exclude Disks and Exceptions fields:**

- **The name of the disk.**
- **A regular expression. See “Using Regular Expressions” on page 494 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 494 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 111 for more information about views.

- 13 If you are generating reports for specific Applications in your environment, select them from the List of Elements.**

- 14 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

## Using the File System Service Time Summary Report

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

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

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

## Reports for Service Level Agreements

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

- SLA Summary Report
- SLA Detailed Report

### SLA Summary Report

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

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

The report answers the following questions:

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

For more information on SLA definitions, see “Working with Service Level Agreements” on page 409.

### Creating an SLA Summary Report

To create an SLA Summary Report:

- 1 **In the Reports Tree panel, click SLA Summary.**

- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 3 Select a Compliance Period to report on.**
- 4 Clear the Display Outage Tables checkbox if you want the report to display only outage graphs.**
- 5 If you want to generate reports for one or more groups that include SLAs, select the groups from the List of Groups area.**
- 6 To generate reports for one or more views that contain SLAs, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

- 7 If you are generating reports for specific Service Level Agreements, select them from the List of SLAs.**
- 8 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

## SLA Detailed Report

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

The report answers the following questions:

- Were there any outages yesterday? If so, how long were they and on which systems did they happen?
- Which business users were affected by service outages?
- What kinds of transaction volumes are we processing?

- What are the most important things we can fix in order to meet our SLA targets?

For more information on SLA definitions, see “Working with Service Level Agreements” on page 409.

## Creating an SLA Detailed Report

To create an SLA Summary Report:

- 1 In the Reports Tree panel, click SLA Detailed.**
- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 3 Select a Compliance Period to report on.**
- 4 Clear the Display Outage Tables checkbox if you want the report to display only outage graphs.**
- 5 If you want to generate reports for one or more groups that include SLAs, select the groups from the List of Groups area.**
- 6 To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

- 7 If you are generating reports for specific Service Level Agreements, select them from the List of SLAs.**
- 8 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

## Reports for Availability

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

- Application Availability Report
- Incident Priority Report
- Service Monitor Availability Report
- Service Monitor Outages Report

### Application Availability Report

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

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

For more information on Applications, see “Working with Applications” on page 103.

### 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 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 for more information.

## Incident Priority Report

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

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

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

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

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

- **Mean Time Between Failure:** The average amount of time that an Element's associated service monitors were all running (i.e., in non-critical states) over a given time period.

Elements whose associated service monitors experience no downtime are still included in the report, but will not include an MTBF count since they did not experience an incident during the time period.

- **Mean Time to Repair:** The average number of minutes any of an Element's associated service monitors were in a critical state over a given time period.



A service is considered repaired, or being repaired, when its status changes from critical to one of “MAINT”, “UNKNOWN”, “WARNING”, or “OK”.

For all report elements, a service monitor is considered to have reached a critical state—thus has caused an incident, is contributing to downtime, or is an ongoing failure—when it actually generates an alert. The period preceding the alert, during which rechecks are intermittently being performed to avoid a false positive, does not count. See “Understanding the Alert Flow” on page 431 for information on rechecks leading to a generated alert.

## Creating an Efficiency Report

To create an Efficiency report, do the following:

- 1 In the Reports Tree panel, click Efficiency.**
- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

Service monitors that, based on the selected time range, are already in a critical state will be included in calculations for downtime, incident counts, and other report elements.

- 3 In the Report Options area, select the charts you want included in the report.**
- 4 In the Report Options section, select the level of granularity at which the information will be presented (i.e., daily, weekly, or monthly).**
- 5 If you want to generate reports for groups of systems, select the groups from the List of Groups area.**
- 6 To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

- 7 If you are generating reports for specific systems in your environment, select them from the List of Elements.**

- 8 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 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 22.

### 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 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 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 23.

- 3 **Select one of the following options from the Sort by dropdown list:**

- Sample Time by Element
- Service Name by Element
- All Sample Times

- 4 From the Sort Direction dropdown list, select Ascending or Descending.**
- 5 If you want to generate reports for groups of systems, select the groups from the List of Groups area.**
- 6 To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

- 7 If you are generating reports for specific systems in your environment, select them from the List of Elements.**
- 8 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 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 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 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 261 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 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 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

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

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

- VM Sprawl Report
- vSphere Workload Report
- VMware Workload Report
- VMware Infrastructure Density Report
- LPAR Workload Report

### VM Sprawl Report

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

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

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

- 1 **monthly VM population trends for up to the last year**

For each month, stacked bars show the following VM totals:

- created during that month and are currently running
- created during that month and are currently suspended or powered off
- destroyed during that month
- created and destroyed during that same month

**2 a list of VMs created during the time period**

**3 a list of VMs destroyed during the time period**

**4 lists summarizing various VM power state scenarios:**

- existing VMs created that were powered off or suspended for 20% or more of the time, during the time period
- existing VMs that were not powered on during the time period
- existing VMs that were in a suspended state for the time period

## Creating a VM Sprawl Report

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

- 1 In the Reports Tree panel, click VM Sprawl.**
- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 3 In the Report Options section, select which types of VM population totals you would like to include in the report.**
- 4 Select the Elements whose child-object VMs you would like to include in the report totals.**

In most cases, for the greatest visibility with a high-level view, the selected Element will most likely be a VMware vCenter server or datacenter Element.

- 5 Select a report generation option. See “Report Generation Options” on page 454 for details.**

### 6 Do one of the following:

- Click the **Generate Report** button.
- Enter a name for the report in the **Save to My Portal As** field, and optionally enter text in the **Report Description** field. Then, click **Save Report**.

The report parameters are saved to the **My Portal** panel. Doing this does not generate the report.

### 7 To schedule the saved report to run at a specific time or interval, click the **Scheduled** checkbox.

See “Scheduling Reports” on page 459 for more information on configuring a scheduled report.

## vSphere Workload Report

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



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

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

## vSphere Workload Report Metrics

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

CPU Usage	<p>the following is shown for the interval:</p> <ul style="list-style-type: none"> <li>• for ESX servers, the total CPU capacity, in GHz</li> <li>• the current amount of CPU resources, in GHz, used by ESX hosts or VMs</li> <li>• the total CPU capacity of the Element</li> </ul>
Memory Usage	<p>the following is shown for the interval:</p> <ul style="list-style-type: none"> <li>• for ESX servers, total memory capacity, in GB</li> <li>• the amount of memory, in GB, used by ESX hosts or VMs</li> <li>• the total memory capacity of the Element</li> </ul>
Disk I/O	the rate, in MBps, at which ESX hosts or VMs are reading and writing data to and from disk
Network I/O	the rate, in Mbps, at which ESX hosts or VMs are receiving or transmitting data over the network
% Wait	the amount of time during the interval, as a percentage, that the VM or all VMs on an ESX host, resource pool or vApp had scheduled CPU time, but gave nothing to process
% Ready	the amount of time during the interval, as a percentage, that the VM or all VMs on an ESX host, resource pool or vApp were ready to process, but were not scheduled CPU time by the host

## Creating a vSphere Workload Report

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

- 1 In the Reports Tree panel, click vSphere Workload.**

- 2 In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 3 Select the Elements you would like to include in the vSphere Workload report from the following sections where applicable:**

- List of vCenter Servers
- List of Datacenters
- List of Clusters
- List of ESX Servers
- List of Resource Pools
- List of vApps
- List of Virtual Machines

- 4 Select which vSphere Element Summaries to Display.**

For each of the Elements you selected in the previous step, you can display resource usage summaries of their VMware vSphere child objects in the detailed sections of the report.

For example, if you select a VMware vCenter server Element to be in the report, selecting the **Datacenter** and **Cluster** summary options will include summaries for datacenters and clusters that are child objects of the vCenter server. Note that selecting a VMware vSphere child object will not automatically include its own child objects (e.g., ESX servers that are part of a cluster); a selection exclusively includes that specific object.

Only VMware vSphere child objects that are relevant to the parent will be shown:

- if you select the **Datacenter** summary option, VMware vCenter server Elements included in the report will include resource usage summaries for its child datacenter objects
- if you select the **Cluster** summary option, VMware vCenter server and datacenter Elements in the report will include resource usage summaries for their child cluster objects



- if you select the **ESX Server** summary option, datacenter and cluster Elements in the report will include resource usage summaries for their child ESX server objects
- if you select the **Resource Pool** or **vApp** summary options, any datacenter, cluster, or ESX server Elements included in this report will include resource usage summaries of its top-level resource pools and vApps
- if you select the **Resource Pool** or **vApp** summary options, any parent resource pool or vApp Elements included in this report will include resource usage summaries for their child resource pools and vApps
- if you select the **Virtual Machine** summary option, ESX servers, resource pools, and vApps included in this report will include resource usage summaries for their child VMs
- if you select the **Virtual Machine** summary option, datacenters and clusters included in this report will include percent-ready and percent-wait summaries for its child VMs

**5 In the Report Options section, select whether to Show Overview Charts in the report.**

The overview charts consist of CPU, memory, disk, and network usage overviews for the defined time range.

**6 Select a report generation option. See “Report Generation Options” on page 454 for details.**

**7 Do one of the following:**

- Click the **Generate Report** button.
- Enter a name for the report in the **Save to My Portal As** field, and optionally enter text in the **Report Description** field. Then, click **Save Report**.

The report parameters are saved to the **My Portal** panel. Doing this does not generate the report.

**8 To schedule the saved report to run at a specific time or interval, click the Scheduled checkbox.**


See “Scheduling Reports” on page 459 for more information on configuring a scheduled report.

## VMware Workload Report

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


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

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



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

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



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

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

## Creating a VMware Workload Report

To create a VMware Workload report, do the following:

- 1 **In the Reports Tree panel, click VMware Workload.**
- 2 **In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 3 **In the Report Options section, select one of the following:**

- Workload Profile - CPU

The percentage of CPU time that is being used by a VMware instance. This is a percentage of the available maximum amount of CPU time. This ensures that all of the CPU usage figures add up to the overall CPU usage of the server.

- Workload Profile - Memory

The amount of physical memory, in kilobytes, that is being used by a VMware instance.

- **Workload Profile - Disk IO**  
The amount of the disk I/O capacity, in kilobytes per second, that is being used by a VMware instance.
- **Workload Profile - Network IO**  
The amount of the network I/O capacity, in kilobits per second, that is being used by a VMware instance.
- **Workload Profile - % Ready**  
The amount of time that one or more instances running on an ESX server is ready to run, but cannot run because it cannot access the processor on the ESX server.
- **Workload Profile - % Used**  
The percentage of CPU time that an instance running on an ESX server is using.

**4 If you want to generate reports for systems in specific groups, select the groups from the List of Groups area.**

**5 To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

**6 If you are generating reports for specific systems in your environment, select them from the List of Elements.**

**7 Select a report generation option. See “Report Generation Options” on page 454 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 459 for more information on configuring a scheduled report.

## VMware Infrastructure Density Report

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

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

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

### Creating a VMware Infrastructure Density Report

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

- 1 **In the Reports Tree panel, click VMware Infrastructure Density.**
- 2 **In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 3 **In the Report Options section, indicate whether you want to Include Charts for Individual ESX Servers by selecting or clearing the check box.**

When this option is enabled, a separate chart with VM counts will be created for each ESX server that is included in the report.

- 4 In the Report Options section, select the level of granularity at which the virtual infrastructure density information will be presented (i.e., daily, weekly, or monthly).**
- 5 If you want to generate reports for groups of systems, select the groups from the List of Groups area.**
- 6 To generate reports for one or more views, select the groups from the List of Views area.**

See “Working with Views” on page 111 for more information about views.

- 7 If you are generating reports for specific systems in your environment, select them from the List of Systems and Nodes.**
- 8 Select a report generation option. See “Report Generation Options” on page 454 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 456 and “Scheduling Reports” on page 459 for more information.

## LPAR Workload Report

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

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

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

## Creating an LPAR Workload Report

To create an LPAR Workload report, do the following:

- 1 **In the Reports Tree panel, click LPAR Workload.**
- 2 **In the Date and Time Range area, select the dates and times on which to report.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 111 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 454 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 456 and “Scheduling Reports” on page 459 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 22

## Understanding Graphing

---

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


<i>Graphing in up.time</i> .....	536
<i>Using the Graph Editor</i> .....	538

## Graphing in up.time


You can graph performance information to learn about the behavior of a system in your environment. Graphs visualize information about CPU, memory, and process usage; as well as network, disk, and user activity. For more information about specific graphs, see “Using Graphs” on page 543.

up.time can generate performance data graphs in two ways:

- In Internet Explorer, the graph is generated using an ActiveX graphing control, as shown below:



- In any Java-enabled Web browser on any operating system – for example, in Firefox, on Linux – the graph is generated using a Java graphing applet, as shown below:



You can click any line in the graph or any item in either axis to zoom in on a particular time period or value. Click the R key on your keyboard to return to the original view.



You can modify ActiveX graphs after they have been generated. You cannot modify Java graphs.

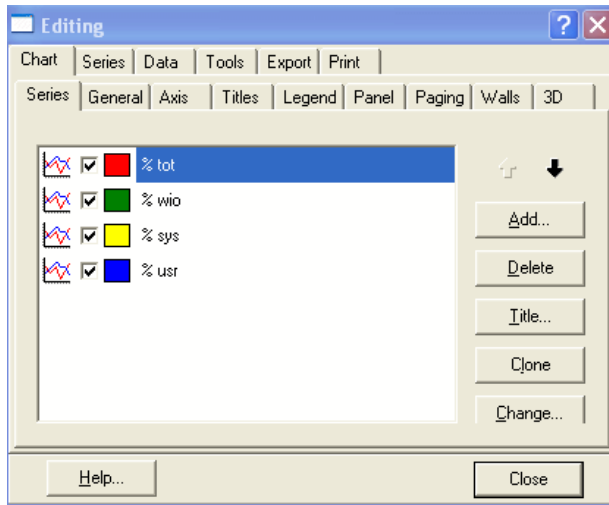
## Graphing Tool

After you generate an ActiveX graph, you can customize it using up.time's graphing tool. With the graphing tool, you can do the following:

- apply graphing line styles
- apply graphing and charting formats
- apply titles, text, and dimensioning
- manipulate a graphing axis
- apply dynamic motion to a graph

## Using the Graph Editor

The Graph Editor enables you to manipulate the presentation of your graphs, as well as apply a variety of effects to a graph to change its overall look. The following image illustrates the Graph Editor:



Use the Graph Editor to do the following:

- exclude graph lines
- change the style of the graph
- re-arrange the order of lines on your graph, or the actual data, to highlight specific Elements 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 543 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 23

## Using Graphs

---

This chapter describes each [up.time](#) graph in the following sections:

<i>Overview</i> .....	544
<i>Viewing the Status of a System</i> .....	545
<i>Monitoring CPU Performance</i> .....	547
<i>Multi-CPU Usage</i> .....	551
<i>Graphing Memory Usage</i> .....	554
<i>Graphing Processes</i> .....	557
<i>Graphing TCP Retransmits</i> .....	559
<i>Graphing User Activity</i> .....	560
<i>Workload Graphs</i> .....	561
<i>Network Graphs</i> .....	567
<i>Disk Performance Statistics Graph</i> .....	570
<i>Top 10 Disks Graph</i> .....	572
<i>File System Capacity Graph</i> .....	574
<i>VXVM Stats Graph</i> .....	575
<i>Novell NRM Graphs</i> .....	577
<i>Instance Motion Graphs</i> .....	579
<i>Displaying Detailed Process Information</i> .....	580

## Overview

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

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



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

## UNIX vs. Windows Performance Monitoring

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

## Viewing the Status of a System

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

The Quick Snapshot is typically used as a preliminary step toward root cause analysis. When you first acknowledge an issue by clicking an Element name on either **Global Scan** or the **My Alerts** section of **My Portal**, you are shown the Quick Snapshot for that Element. From here, you can scan the information provided in the charts and tables, and begin further investigation. (For example, if you notice problem while viewing the Quick Snapshot, you can generate a report to obtain more information about the problem.)

The Quick Snapshot contains the following information:

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

## Viewing a Quick Snapshot

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



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

## Monitoring CPU Performance

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

- Usage (% busy)
- Run Queue Length
- Run Queue Occupancy

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a CPU Performance Graph” on page 550.

### 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 Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click one of the following options:**
  - Usage (% busy)
  - Run Queue Length
  - Run Queue Occupancy
- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.
- 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 Infrastructure** panel, click the name of the system whose information you want to graph.
- 2 In the **Tree** panel, click the **Graphing** tab.
- 3 Click **Multi-CPU Usage**.
- 4 Select the start and end dates and times for which the graph will chart data.

For more information, see “Understanding Dates and Times” on page 23.

### 5 Click one of the following options:

- **User %**  
The percentage of CPU user processes that are in use. For Windows systems, this option is **% User Time**.
- **System %**  
The percentage of CPU kernel processes that are in use. For Windows systems, this option is **% System Time**.
- **% Privileged Time**  
On Windows systems, the percentage of time that the CPU spends executing kernel commands.
- **Wait I/O %**  
The percentage of time that a process which can be run must wait for a device to perform an I/O operation.
- **SMTX**  
The number of read or write locks that a thread was not able to acquire on the first attempt, as reported by the `mpstat` command.



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

- **XCAL**  
The number of interprocess cross-calls.  
In a multi-processor environment, one processor sends cross-calls to another processor to get that processor to do work. Cross-calls can also be used to ensure consistency in virtual memory. Heavy file system activity – such as NFS – can result in a high number of cross-calls.
- **Interrupts**  
The number of CPU interrupts. For Windows systems, this option is **% Interrupt Time**.  
Interrupts are a mechanism that a device uses to signal to the kernel that it needs attention, and that immediate processing is required on its behalf.

- Interrupts/sec  
On Windows systems, rate at which CPU handles interrupts from applications or hardware each second. If the value for Interrupts/sec is high, there could be problems with the hardware on the system.
  - Total %  
On Solaris systems, the total amount of User %, System %, and Wait I/O %.  
On Windows systems, this option is % **Total** and is the total amount of % User Time, % Privileged Time, and % Interrupt Time.
- 6 **Select the CPUs to graph from the Choose CPUs to graph list.**
  - 7 **Click Generate Graph.**

## Graphing Memory Usage

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

- Used
- Cache Hit Rate
- Paging Statistics
- Free Swap

These graphs use the same input criteria, but they return different data. For information on how to generate these graphs, see “Generating a Memory Usage Graph” on page 556.

### Used

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

### Cache Hit Rate

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

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

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

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

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

## Paging Statistics

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

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

The `handsreadpages` 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 Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 **In the Tree panel, click the Graphing tab.**
- 3 **Click one of the following options:**
  - Used
  - Cache Hit Rate
  - Paging Statistics
  - Free Swap
- 4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 558.

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 580.
- “Workload Graphs” on page 561.

### Number of Processes

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

### Process Running, Blocked, Waiting

This graph indicates whether or not there is enough CPU capacity for the processes that are being run on a system. If the size of the blocked or waiting queue is disproportionate to the running queue, then either the system does not have enough CPUs or is too I/O bound.

A blocked process signals a disk bottleneck. If the number of blocked processes approaches or exceeds the number of processes in the run queue, you should tune the disk subsystem. Whenever there are any blocked processes, all CPU idle time is treated as wait for I/O time. If database batch jobs are running on the system that is being monitored, there will always be some blocked processes. However, you can increase the throughput of batch jobs by removing disk bottlenecks.

## Process Creation Rate

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

## Generating a Process Graph

To generate a process graph, do the following:

- 1 In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click one of the following options:**
  - Number of Processes
  - Process Running, Blocked, Waiting
  - Process Creation Rate

- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 5 Click Generate Graph.**



## Graphing TCP Retransmits

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

- A header
  - Contains transmission information, such as the IP addresses of the sender and receiver, the protocol that is being used, and the packet number.
- A payload
  - Contains the data that is being sent.
- A trailer
  - Contains data that denotes the end of the packet, as well as error correction information.

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

## Generating a TCP Retransmits Graph

To generate a TCP retransmits graph, do the following:

- 1 **In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 **In the Tree panel, click the Graphing tab.**
- 3 **Click TCP Retransmits.**
- 4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.
- 5 **Click Generate Graph.**

## Graphing User Activity

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

- Login History

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

- Sessions

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

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

## Generating a User Activity Graph

To generate a user activity graph, do the following:

- 1 **In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 **In the Tree panel, click the Graphing tab.**
- 3 **Click either Login History or Sessions.**
- 4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 562.

Each workload graph captures the following metrics:

- **CPU %**  
The percentage of CPU time that is taken up by a user, group, or process.
- **Memory Size**  
The amount of the page file and virtual memory that is taken up by a user, group, or process.  
On Windows systems, Memory Size is called *Virtual Bytes*.
- **RSS**  
The Run Set Size, which is the amount of physical memory that is being used by a user, group, or process. On Windows systems, RSS is called *Working Set*.



Workload graphs that are generated for SNMP agents only chart the Memory Size metric.

## Generating a Workload Graph

To generate a workload graph, do the following:

**1 In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**

**2 In the Tree panel, click the Graphing tab.**

**3 Click one of the following options:**

- Workload - User
- Workload - Group
- Workload - Process Name

**4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

**5 Click one of the following metrics:**

- CPU %
- Memory Size or Virtual Bytes (on UNIX and Windows, respectively)
- RSS or Working Set (on UNIX and Windows, respectively)



You can only graph one metric at a time.

**6 Select one or more of the available users, groups, or processes from the list.**

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



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

- 7 Click Add.
- 8 Click Generate Graph.

## Workload Top 10 Graphs

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

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

### Generating a Workload Top 10 Graph

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

- 1 In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click one of the following options:**
  - Workload Top 10 - User
  - Workload Top 10 - Group
  - Workload Top 10 - Process Name
- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 5 Click one of the following options:**

- CPU %
- Memory Size
- RSS

Graphs generated for SNMP agents only chart the memory size metric.

- 6 Click Generate Graph.**

## LPAR Workload Graphs

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

- Workload - CPU  
The amount of CPU time that is being used by the LPAR.
- Workload - Memory  
The total amount of memory being used by an LPAR.
- Workload - Disk  
The amount of data that has been transferred to and from the disk.
- Workload - Network  
The amount of data that has been transferred over the network interface used by the LPAR.

You can also graph the CPU entitlement of individual LPARs using the CPU Utilization graph. See “LPAR CPU Utilization Graphs” for more information.

### Generating an LPAR Workload Graph

To generate an LPAR Workload graph, do the following:

- 1 In the Global Scan or My Infrastructure panel, click the name of the pSeries server which is hosting the LPARs whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click one of the following options:**
  - Workload - CPU
  - Workload - Memory
  - Workload - Disk
  - Workload - Network

- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 5 Click Generate Graph.**

## LPAR CPU Utilization Graphs

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

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

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

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

For more information, see “Understanding Dates and Times” on page 23.

- 5 Select the name of the LPAR whose information you want to graph.**

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

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



## Network Graphs

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

- I/O
- Errors
- NetFlow

The I/O and Errors graphs use the same input criteria, but return different data. NetFlow graphs are available if [up.time](#) is integrated with Scrutinizer. For information on how to generate these graphs, see “Generating a Network Graph” on page 568.

### I/O

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

The I/O graph captures the following statistics:

- In bytes  
The number of bytes received over the network interface each second.
- Out bytes  
The number of bytes sent by the network interface each second.

### Errors

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

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

- Excessive noise.

- Cabling problems.
- Problems with backbone connections.

The Errors graph captures the following statistics:

- In Errors  
A data packet was received but could not be decoded because either the header or trailer of the packet was not available.
- Out Errors  
A data packet could not be sent due to problems transmitting the packet or formatting the packet for transmission.
- Collisions  
The simultaneous presence of signals from two nodes on the network. A collision can occur when two nodes start transmitting over a network at the same time. Packets that are involved in a collision are broken into fragments and must be retransmitted.

## NetFlow

The NetFlow graphing function transfers you to your Scrutinizer instance.

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

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

See [Generating a Network Graph](#) for information on how to generate this graph.

## Generating a Network Graph

To generate network graphs, do the following:

- 1 In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**

- 3 **Click one of the following options:**
  - I/O
  - Errors
  - NetFlow (available if [up.time](#) has been integrated with Scrutinizer)
- 4 **For I/O and Errors graphs, select the start and end dates and times for which the graph will chart data. For NetFlow, select one of the set time frames.**

For more information, see “Understanding Dates and Times” on page 23.
- 5 **For I/O and Errors graphs, select one or more network interfaces from the Available Interfaces list, and then click Add.**
- 6 **Click Generate Graph.**

## Disk Performance Statistics Graph

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

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

### Generating a Disk Performance Statistics Graph

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

- 1 In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click Disk Performance Statistics.**
- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click Top 10 Disks.**
- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click File System Capacity.**
- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 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 52 for more information.

## Generating a VXVM Stats Graph

To generate a VXVM Stats graph, do the following:

- 1 **In the Global Scan or My Infrastructure panel, click the name of the system whose information you want to graph.**
- 2 **In the Tree panel, click the Graphing tab.**
- 3 **Click VXVM Stats.**
- 4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 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.

## Using Graphs *Novell NRM Graphs*


- **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 88.

## Generating a Novell NRM Graph

To generate a Novell NRM graph, do the following:

- 1 In the Global Scan or My Infrastructure panel, click the name of the Novell NRM system whose information you want to graph.**

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 23.

- 4 Click Generate Graph.**

## Instance Motion Graphs

The VMware VMotion tool enables you to move ESX instances from one server to another without any downtime or loss of data. You would use VMotion to, for example, move an instance to newer and faster hardware, or to temporarily relocate the instance while performing a hardware upgrade.

The Instance Motion graph enables you to keep track of a moving VMware instance. For a given ESX instance, the graph charts which systems it has been running on over a given time range.

### Generating an Instance Motion Graph

To generate an Instance Motion graph, do the following:

- 1 In the Global Scan or My Infrastructure panel, click the name of the ESX instance whose motion you want to graph.**
- 2 In the Tree panel, click the Graphing tab.**
- 3 Click Instance Motion.**
- 4 Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 5 Click Generate Graph.**

## Displaying Detailed Process Information

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

- **Process**  
The name of the process, which is taken from its executed path name.
- **PID**  
The number that identifies the process.
- **PPID**  
The number that identifies the parent process. The PPID can help identify possible relationships between processes.  
On Windows systems, the PPID is called the *Creating Process ID*.
- **UID**  
The ID of the user or account that has been consuming CPU time.  
On Windows systems, the UID is called the *Owner*.
- **GID**  
The ID of the group that has been consuming CPU time.  
On Windows systems, the GID is called the *Group Name*.
- **Memory Used**  
The amount of memory, expressed as a percentage of total available memory, being consumed by a process.  
On Windows systems, Memory Used is called *Virtual Bytes*.  
The **Memory Used** value can be misleading because shared memory between processes is counted multiple times. For example, if five Oracle processes are using 10% of available memory, this does not indicate that Oracle is consuming 50% of system memory.
- **RSS**  
Run Set Size – the amount of physical memory that is being used.  
On Windows systems, RSS is called the *Working Set*.

- CPU %

The percentage of the CPU time used by the process, calculated by dividing total used CPU Time by the process' running time; if applicable, the result is further divided by the number of CPUs for the Element on which the process is running.


On Windows systems, the CPU % is called *% Processor Time*.
- User Time

The amount of time (in seconds) that a particular user, group, or account has been using the CPU.


This value is not displayed for Windows systems.
- User System Time

The amount of time (in seconds) that a process has been consuming system time on the CPU.

This value is not displayed for Windows systems.

 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 Infrastructure** panel, click the name of the system whose information you want to graph.
- 2 In the **Tree** panel, click the **Graphing** tab.
- 3 Click **Detailed Process Information**.

- 4 **Select the start and end dates and times for which the graph will chart data.**

For more information, see “Understanding Dates and Times” on page 23.

- 5 **Click Display Process Information.**

A window containing a chart that lists the process information for the time period that you specified appears. The following image illustrates process information for a Solaris system:

**Detailed Process Information**

Specific Date and Time      Date Range: YYYY-MM-DD      HH:MM:SS  
 Last      From: 2008-04-16      00:00:00      24  
 Quick Date      To: 2008-04-16      23:59:59      24

**Display Process Information**

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

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

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



# CHAPTER 24

## Graphing VMware vSphere Components

---

This chapter outlines graphing options for VMware vSphere components that are monitored by [up.time](#).

This chapter describes each [up.time](#) graph in the following sections:

<i>Overview</i> .....	584
<i>Viewing the Status of a vSphere Element</i> .....	585
<i>CPU Performance Graphs</i> .....	593
<i>Memory Usage Graphs</i> .....	600
<i>Process Graphs</i> .....	609
<i>Network Graphs</i> .....	615
<i>Disk Graphs</i> .....	621
<i>VMware Management Graphs</i> .....	632

### Overview

By generating graphs, you can display performance and usage statistics for [up.time](#) Elements. Generating graphs helps you diagnose problems through root-cause analysis, as well as review the overall health and performance of monitored Elements. With vSync, you can generate graphs for VMware vSphere components that are being monitored by [up.time](#), allowing you to diagnose and review all components of your infrastructure—whether virtual assets or physical ones—from the same view.

When viewing the Graphing tab for any VMware vSphere component, you will see different graphing options depending on the type of VMware vSphere object, or in the case of VMs, the operating system that is running.

## Viewing the Status of a vSphere Element

You can view the status of your VMware vCenter servers, ESX servers, and VMs using Quick Snapshots. The Quick Snapshot summarizes both the recent and current performance of key hardware and process information for a VMware vSphere component that exists in [up.time](#) as an Element, and can help administrators identify potential issues.



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

The Quick Snapshot is typically used as a preliminary step toward root cause analysis. When you first acknowledge an issue by clicking an Element name in either **Global Scan**, or the **My Alerts** section of **My Portal**, you are shown the Quick Snapshot for that Element. From here, you can scan the information provided in the charts and tables, and begin further investigation. (For example, if you notice problem while viewing the Quick Snapshot, you can generate a report to obtain more information about the problem.)

The specific contents of the Quick Snapshot depend on the VMware vSphere component currently in focus:

- VMware vCenter server
- ESX server or cluster
- virtual machine
- virtual machine using WMI or the [up.time](#) agent to collect metrics

### vCenter Server Quick Snapshot Contents

The following information is displayed in a VMware vCenter server Quick Snapshot:

Datacenter Summary
datacenters that are being monitored are listed in alphabetical order, along with resource information for the datacenter as a whole

CPU Capacity Trend	<ul style="list-style-type: none"> <li>• the CPU usage trend of the datacenter over the last 24 hours</li> <li>• maximum and minimum total CPU usage per hour</li> <li>• clicking opens a pop-up displaying the CPU Workload graph for the datacenter showing CPU usage in MHz</li> </ul>
Memory Capacity Trend	<ul style="list-style-type: none"> <li>• the memory usage trend of the datacenter over the last 24 hours</li> <li>• maximum and minimum total CPU usage per hour</li> <li>• clicking opens a pop-up display the Memory Workload graph for the datacenter, showing the memory consumed metric</li> </ul>
Total Active CPU	the total CPU cycles, in GHz, that are available, whether they are currently being used
Total Active Memory	the total available memory, in GB, that is available, whether it is currently in use
Running VMs and Hosts Count	<ul style="list-style-type: none"> <li>• the current total number of VMs and ESX hosts running within the datacenter</li> <li>• clicking opens a custom graph in a pop-up window, displaying the number of VMs and ESX hosts running during the time period</li> </ul>
Active Outages	<ul style="list-style-type: none"> <li>• the current number of <a href="#">up.time</a> service monitors that are attached to any Element within the datacenter that are presently in a WARN or CRIT state</li> <li>• clicking displays the Service Status page for the datacenter</li> </ul>

<b>Top Clusters / Top ESX Servers</b>	
the top five clusters and ESX servers are respectively listed in order of current CPU usage	
CPU Trend 24h	<p>the CPU usage trend of the cluster or ESX server over the last 24 hours, expressed as a percentage of total available CPU cycles</p> <p>clicking opens a pop-up displaying a CPU Workload graph showing actual CPU usage over the last 24 hours</p>
Current %	the current percentage of total available CPU cycles being consumed by the cluster or ESX server
the top five clusters and ESX servers are respectively listed in order of memory used	
Memory Trend 24h	<p>the memory usage trend of the cluster or ESX server over the last 24 hours, expressed as a percentage of total available memory</p> <p>clicking opens a pop-up displaying a Memory Workload graph showing actual memory usage over the last 24 hours</p>
Current GB	the current percentage of total available memory currently being used by the cluster or ESX server
<b>Top Resource Pools</b>	
the top five resource pools are listed in order of current CPU usage	
CPU Trend 24h	<p>the CPU usage trend of the resource pool over the last 24 hours, expressed in raw GHz</p> <p>clicking opens a pop-up displaying a CPU Workload graph showing actual CPU usage over the last 24 hours</p>
Current GHz	the amount of CPU cycles, in raw GHz, currently consumed by the resource pool
the top five resource pools are listed in order of memory used	

Memory Trend 24h	<p>the memory usage trend of the resource pool over the last 24 hours, expressed in raw GBs</p> <p>clicking opens a pop-up displaying a Memory Workload graph showing actual memory usage over the last 24 hours</p>
Current GB	the amount of total memory, in raw GB, currently being used by the resource pool

## ESX Server and Cluster Quick Snapshot Contents

The following information is displayed in an ESX server or cluster Quick Snapshot:

CPU Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>total CPU capacity, in GHz, which acts as the top line of the graph</li> <li>the amount of CPU, in GHz, used by VMs</li> <li>the amount of CPU, in GHz, reserved for VMs</li> </ul>
VM state	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>total number of VMs powered on</li> <li>total number of suspended VMs</li> <li>total number of VMs powered off</li> </ul>
Memory Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>total memory capacity, in GB, which acts as the top line of the graph</li> <li>the amount of memory, in GB, used by VMs</li> <li>the amount of memory, in GB, reserved for VMs</li> </ul>

<p>Swap Usage vs. Swap Rate</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>total swap memory used by VMs, in MB or GB, whichever is more appropriate</li> <li>the rate, in MBps, at which VMs are swapping memory from disk</li> <li>the rate, in MBps, at which VMs are swapping memory to disk</li> </ul>
<p>Disk I/O Rate</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>the rate, in MBps, at which VMs are reading data from disk</li> <li>the rate, in MBps, at which VMs are writing data to disk</li> </ul>
<p>Network I/O Rate</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>the rate, in MBps, at which VMs are receiving data from the network</li> <li>the rate, in MBps, at which VMs are transmitting data over the network</li> </ul>
<p>Top VMs by CPU</p>	<ul style="list-style-type: none"> <li>top five VMs running on the ESX host, ordered by CPU usage, in MHz</li> <li>the Details link is a shortcut to a top-10 workload graph showing the 24-hour trend</li> </ul>
<p>Top VMs by Memory</p>	<ul style="list-style-type: none"> <li>top five VMs running on the ESX host, ordered by memory usage, in MB</li> <li>the Details link is a shortcut to a top-10 workload graph, showing the 24-hour trend</li> </ul>

Top VMs by Disk I/O	<ul style="list-style-type: none"> <li>• top five VMs running on the ESX host, ordered by disk usage, in KBps</li> <li>• the Details link is a shortcut to a top-10 workload graph, showing the 24-hour trend</li> </ul>
Top VMs by Network I/O	<ul style="list-style-type: none"> <li>• top five VMs running on the ESX host, ordered by network I/O throughput, in Kbps</li> <li>• the Details link is a shortcut to a top-10 workload graph, showing the 24-hour trend</li> </ul>

### Virtual Machine Quick Snapshot Contents

The following information is displayed in a virtual machine Quick Snapshot:

CPU Usage	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>• CPU usage in MHz</li> <li>• if allocated to the VM, the amount of reserved CPU cycles, in MHz</li> <li>• if allocated to the VM, the amount of entitled CPU cycles, in MHz</li> </ul>
% Wait and % Ready Time	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>• wait time: amount of time during the interval, as a percentage, that the VM had scheduled CPU time, but gave nothing to process</li> <li>• ready time: amount of time during the interval, as a percentage, that the VM was ready to process, but was not scheduled CPU time by the host</li> </ul>



<p>Memory Usage</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>• memory usage in MB or GB, whichever is more appropriate</li> <li>• if allocated to the VM, the amount of reserved memory, in MB</li> <li>• if allocated to the VM, the amount of entitled memory, in MB</li> </ul>
<p>Swap Usage vs. Swap Rate</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>• total swap memory, in KB, used during the interval</li> <li>• the rate, in KBps, at which the VM is swapping memory from disk</li> <li>• the rate, in KBps, at which the VM is swapping memory to disk</li> </ul>
<p>Disk I/O Rate</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>• the rate, in KBps, at which the VM is reading data from disk</li> <li>• the rate, in KBps, at which the VM is writing data to disk</li> </ul>
<p>Network I/O Rate</p>	<p>the following is shown for the last 24 hours, where the mouse-over segments are in 10-minute intervals:</p> <ul style="list-style-type: none"> <li>• the rate, in KBps, at which the VM is receiving data from the network</li> <li>• the rate, in KBps, at which the VM is transmitting data over the network</li> </ul>

<p>Top 10 Processes</p>	<p>note that this process list is only available if the VM Element's metrics are being reported via WMI or the <a href="#">up.time</a> Agent</p> <ul style="list-style-type: none"> <li>• top 10 processes on the VM, ranked by CPU usage</li> <li>• the Details link is a shortcut to the Detailed Process Information table for the VM Element</li> </ul>
-------------------------	---

### Viewing a Quick Snapshot for a vSphere Component

To display the Quick Snapshot page for any VMware vSphere component, do the following:

- 1 In the My Infrastructure panel, locate the VMware vCenter, ESX server, or VM Element whose Quick Snapshot you would like to view.**
- 2 Click the gear icon beside the Element.**
- 3 In the Element's Configure pop-menu, click Graph Performance.**

Note that when you are viewing an Element's profile, you can always access its Quick Snapshot by clicking the **Graphing** tab, then clicking **Quick Snapshot** in the tree panel.

## CPU Performance Graphs

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

- CPU Workload
- CPU Usage
- Multi-CPU Usage
- Wait and Ready Time

### CPU Workload

CPU Workload graphs help you gauge the demand being placed on your computing resources, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers or clusters are consuming the most CPU cycles, or for an ESX server, which VMs are creating the largest workload.

#### CPU Workload Metrics

For CPU Workload graphs for a VMware vCenter server, the amount of MHz used can be generated for the following Element types:

- Datacenters
- Clusters
- ESX servers
- virtual machines
- resource pools
- vApps

For CPU Workload graphs for an ESX server, the amount of MHz used can be generated for the following Element types:

- virtual machines
- resource pools
- vApps

### Graphing CPU Workload for a vCenter or ESX Server

To generate a CPU Workload graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click CPU Workload (for a VMware vCenter server), or Workload in the CPU section (for an ESX server).**

**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 23.

**4 Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select a VMware vSphere component whose metrics will be graphed.**

**6 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**7**

**8 Click Generate Graph.**

A pop-up window will appear, displaying the CPU Workload graph you have configured.

## CPU Usage

The CPU Usage graph shows how CPU resources are managed and used across all instances on an ESX server, or an individual VM.

## CPU Usage Metrics

The following metrics are displayed in a CPU Usage graph:

Reservation	the minimum amount of CPU resources reserved for the VM, or all VMs running on an ESX host
Entitlement	the amount of CPU resources allocated to a VM, or all VMs running on an ESX host, on the assumption that all VMs on a host are completely busy, and load is perfectly balanced across hosts
Usage MHz	the amount of CPU used, in MHz, by the VM, or all powered on VMs running on the ESX host
Usage %	the amount of CPU resources, as a percentage of total available CPU, used by the cluster, ESX host, or VM
Usage % (agent)	the amount of CPU resources, as a percentage of total available CPU, used by a VM that is reporting detailed metrics via WMI or the <a href="#">up.time</a> agent

## Graphing CPU Usage for an ESX Server or VM

To generate a CPU Usage graph, do the following:

**1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose CPU usage you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click Usage in the CPU section.**

**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 23.

**4 Select the metrics you want to include in the graph.**

**5 Click Generate Graph.**

A pop-up window will appear, displaying the CPU Usage graph you have configured.

## Multi-CPU Usage

The Multi-CPU Usage graph charts the performance statistics for ESX hosts or VMs 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.

### Multi-CPU Usage Metrics

The following metrics are displayed in a Multi-CPU Usage graph:

ESX Server	
Usage MHz	usage of total CPU resources, in MHz, by physical cores, on the ESX host
Virtual Machine	
Usage MHz	the amount of CPU usage, in MHz, by CPU, on the VM
% System Time	the amount, as a percentage, spent on system processes on each vCPU in the VM
% Wait Time	the amount, as a percentage, vCPUs on the VM spent in wait state
% Used	the amount, as a percentage, of used vCPU time
Virtual Machine with WMI- or Agent-Based Metrics	
% Total	total amount of user, privileged, and interrupt time, by vCPU
% Usr	the amount of time, as a percentage, of user processes, by vCPU

% Sys	the amount of time, as a percentage, of system process use, by vCPU
% Wait	the amount of time, as a percentage, of wait time use
Interrupts	number of CPU interrupts

## Graphing Multi-CPU Usage for an ESX Server or VM

To generate a multi-CPU usage graph, do the following:

**1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose multi-CPU workloads you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the CPU section of the Tree panel, click Multi-CPU Usage.**

**3 Select the start and end dates and times for which the graph will chart data, then click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 In the next step, select whether to graph the highest resource consumers, or specific VMs.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**5 Click Generate Graph.**

A pop-up window will appear, displaying the Multi-CPU Workload graph you have configured.

## Wait and Ready Time

CPU wait time is the amount of time a VM is given scheduled time, but there is nothing to process, resulting in an idle CPU.

CPU ready time is the amount of time that a VM is ready for processing, but could not get scheduled to run on the physical CPU.

The Wait and Ready Time graph helps you determine if a guest is waiting too often on a host, or a host is waiting too often on a guest.

## Wait and Ready Time Metrics

The following metrics are displayed in a Wait and Read Time graph:

% Wait	the amount of time during the interval, as a percentage, that the VM or all VMs on an ESX host, resource pool or vApp had scheduled CPU time, but gave nothing to process
% Ready	the amount of time during the interval, as a percentage, that the VM or all VMs on an ESX host, resource pool or vApp were ready to process, but were not scheduled CPU time by the host

## Graphing Wait and Ready Time for an ESX Server or VM

To generate a percent-wait or percent-ready graph, do the following:

**1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose wait and ready time you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the CPU section of the Tree panel, click Wait & Read Time.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select a VMware vSphere component whose metrics will be graphed.**

**6 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.



**7 Choose either the % Wait, or % Ready metric to include in the graph.**

**8 Click Generate Graph.**

A pop-up window will appear, displaying the CPU Wait and Ready Time graph you have configured.

## Graphing Wait and Ready Time for an Agent-Based VM

To generate a percent-wait or percent-ready graph for a VM that receives metrics via WMI or the [up.time](#) agent, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VM whose wait and ready time you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click Wait & Ready Time in the CPU section.**

**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 23.

**4 Select one of both of the % Ready and % Wait metrics to include in the graph.**

**5 Click Generate Graph.**

A pop-up window will appear, displaying the graph you have configured.

## Memory Usage Graphs

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

- Memory Workload
- Memory Usage
- Memory Profile
- Memory: Guest Paging
- Memory Swap

### Memory Workload

Memory Workload graphs help you gauge the demand being placed on the memory allocated to your physical and virtual, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers or clusters have been using the most memory, or for an ESX server, which VMs are the recipients of overcommitted granted memory.

### Memory Workload Metrics

The following metrics can be used to generate a Memory Workload graph:

VMware vCenter Server	
Consumed	the amount of memory, in MB, that is used by or reserved for a VM, or all VMs that are part of a resource pool or vApp, or the total physical memory used by the datacenter, cluster, or ESX host
Ballooned	the amount of guest physical memory being reclaimed through ballooning from a VM, or all VMs that are part of a resource pool or vApp, or the total memory balloon for all VMs in a datacenter, cluster, or ESX host

ESX Server	
Usage	the amount of memory used, as a percentage of total available memory, by all VMs on the ESX host, or in a resource pool or vApp
Active	the estimated amount of active memory in use, in KB, by VMs, or all VMs in a resource pool or vApp
Consumed	the amount of memory, in MB, that issued by or reserved for a VM, or all the VMs that are part of a resource pool or vApp
Granted	the amount of physical memory allocated to machine memory for a VM, or all the VMs that are part of a resource pool or vApp
Ballooned	the amount of guest physical memory being reclaimed through ballooning from a VM, or all VMs that are part of a resource pool or vApp

## Graphing Memory Workload for a vCenter or ESX Server

To generate a Memory Workload graph, do the following:

- 1 **Go to the Element's Quick Snapshot page.**  
For example, in the **My Infrastructure** panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click **Graph Performance**.
- 2 **In the Tree panel, click Memory Workload (for a VMware vCenter server), or Workload in the Memory section (for an ESX server).**
- 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 23.
- 4 **Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**
- 5 **In the first step, select a VMware vSphere component whose metrics will be graphed.**

- 6 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

- 7 Select the metric to include in the graph.**
- 8 Click Generate Graph.**

A pop-up window will appear, displaying the Memory Workload graph you have configured.

## Memory Usage

The Memory Usage graph displays the amount of memory being used on an ESX server, or a VM.

### Memory Usage Metrics

The following metrics are displayed in a Memory Usage graph:

ESX Server Metrics	
Reservation	the minimum amount of memory reserved for VMs running on an ESX host
Consumed	the amount of memory, in MB, that is used by or reserved for VMs running on an ESX host
Active	the estimated amount of active memory in use, in MB, by VMs running on an ESX host
Granted	the amount of memory allocated to VMs running on an ESX host, on the assumption that all are completely busy, and load is perfectly balanced across hosts
VM metrics	
Reservation	the minimum amount of memory reserved for the VM

Entitlement	the amount of memory, in MB, allocated to a VM, on the assumption that all VMs on a host are completely busy, and load is perfectly balanced across hosts
Consumed	the amount of memory, in MB, that is used by or reserved for a VM
Active	the estimated amount of active memory in use, in MB, by the VMs
Granted	the amount of physical memory allocated to machine memory for a VM
Used MB (Agent)	the amount of memory used by a VM that is reporting detailed metrics via WMI or the <a href="#">up.time agent</a>

## Graphing Memory Usage for an ESX server or VM

To generate a Memory Usage graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose memory usage you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click Usage in the Memory section.**

**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 23.

**4 Select the metrics you want to include in the graph.**

**5 Click Generate Graph.**

A pop-up window will appear, displaying the Memory Usage graph you have configured.

## Memory Profile

The Memory Profile graph displays trends in allocated memory for VMs on an ESX server. Use this graph to ensure VMware vSphere’s memory resource management techniques are working correctly.

### Memory Profile Metrics

The following metrics are displayed in a Memory Profile graph:

ESX Server Metrics	
Shared Common	the amount of machine memory that is shared by all powered on VMs on the ESX host
Swapped	the total amount of guest physical memory swapped out to swap files of all VMs running on the ESX host
Ballooned	the total memory balloon of all powered on VMs
Overhead	the total amount of machine memory used to run all VMs running on the ESX host
VM metrics	
Shared	the amount of guest physical memory shared with other VMs
Swapped	the amount of guest physical memory swapped out to the VMs swap file
Ballooned	the amount of guest physical memory being reclaimed through ballooning from a VM
Overhead	the amount of machine memory used to run the VM

### Graphing a Memory Profile for an ESX Server or VM

To generate a Memory Profile graph, do the following:

- 1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose Memory Profile graph you want to create, click its corresponding gear icon, then click **Graph Performance**.

- 2 **In the Tree panel, click Profile in the Memory section.**
- 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 23.

- 4 **Select the metrics you want to include in the graph.**
- 5 **Click Generate Graph.**

A pop-up window will appear, displaying the Memory Profile graph you have configured.

## Memory: Guest Paging

This graph indicates whether or not a VM guest is short of memory.

### Guest Paging Metrics

The following metrics are displayed in a Guest Paging graph:

Page In/sec	number of memory pages transferred from disk to RAM
Page Out/sec	number of memory pages transferred from RAM to disk
Page Frees/sec	number of memory pages removed from being swapped to and from RAM and secondary storage

### Graphing Guest Memory Paging Rates on an Agent-Based VM

To generate a guest memory paging rate graph for a VM that receives metrics via WMI or from the [up.time](#) agent, do the following:

- 1 **Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VM whose paging rates you want to create, click its corresponding gear icon, then click **Graph Performance**.

- 2 In the **Memory** section of the **Tree** panel, click **Guest Paging**.
- 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 23.

- 4 Select the paging metrics you want to include in the graph.
- 5 Click **Generate Graph**.

A pop-up window will appear, displaying the Memory Profile graph you have configured.

## Memory Swap

The Memory Swap graph charts the amount of memory held in swap files on ESX servers, as well as the rates at which data is swapping between memory and secondary storage.

### Memory Swap Metrics

The following metrics are displayed in a Memory Swap graph:

ESX Server Metrics	
Swapped	total guest physical memory, in MB, swapped out to swap files of VMs in resource pools, vApps, or on the ESX host
Swap Total Rate	the total swap in rate and swap out rate, in MB, for VMs in resource pools, vApps, or on the ESX host
VM Metrics	
Swap Used	the amount of guest physical memory, in MB, swapped out to the VM's swap file
Swap In Rate	the rate at which memory, in MB, is swapped into active memory



Swap Out Rate	the rate at which memory, in MB, is being swapped out to disk
Swap Used MB (Agent)	the amount of guest physical memory, in MB, swapped out to the VM's swap file, as reported via WMI or the <a href="#">up.time</a> agent

## Creating a Memory Swap Graph for an ESX server or VM

To generate a percent-wait or percent-ready graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose memory swap statistics you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Memory section of the Tree panel, click Swap.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see "Understanding Dates and Times" on page 23.

**4 Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select a VMware vSphere component whose metrics will be graphed.**

**6 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**7 Select the metric you wish to include in the graph.**

**8 Click Generate Graph.**

A pop-up window will appear, displaying the Memory Swap graph you have configured.

### Creating a Memory Swap Graph for an Agent-Based VM

To generate a Memory Swap graph for a VM that receives metrics via WMI or the [up.time](#) agent, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VM whose memory swap information you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click Swap in the Memory section.**

**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 23.

**4 Select the metrics you want to include in the graph.**

**5 Click Generate Graph.**

A pop-up window will appear, displaying the Memory Swap graph you have configured.

## Process Graphs

up.time uses the following graphs to report system process information for VMware vSphere-monitored VMs whose metrics are reported via WMI or with the up.time agent:

- Process Workload
- Detailed Process Information
- Number of Processes

### Process Workload

The Process Workload graph determines the demand that network and local services are putting on a VM whose metrics are reported via WMI or with the up.time agent. The graph charts an aggregate amount of performance information for a given user, group, or process.

#### Process Workload Metrics

The following metrics are displayed in a Process Workload graph:

% CPU	the percentage of the VM's CPU time that is taken up by a user, group, or process
Total Memory Size	<ul style="list-style-type: none"> <li>• the amount of the page file and virtual memory that is taken up by a user, group, or process</li> <li>• on Windows VMs, Memory Size is called Virtual Bytes.</li> </ul>
RSS Memory Size	<ul style="list-style-type: none"> <li>• the Run Set Size, which is the amount of physical memory that is being used by a user, group, or process</li> <li>• on Windows systems, RSS is called Working Set</li> </ul>

### Graphing Process Workloads for an Agent-Based VM

To generate a Process Workload graph for a VM that receives metrics via WMI or the [up.time](#) agent, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VM whose process workload you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Process section of the Tree panel, click Workload.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 Click one of the Quick Graphs options to display a pre-configured workload graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select an object whose workload will be graphed.**

**6 In the next step, select whether to graph the highest resource consumers, or specific processes running on the VM.**

If you select **Specific**, a process selection dialog will appear, requiring you to build a list.

**7 Select one or more process metrics to include in the graph.**

**8 Click Generate Graph.**

A pop-up window will appear, displaying the Process Workload graph you have configured.

## Detailed Process Information

Detailed process information provides insight into how various user and system processes are consuming system resources.

## Detailed Process Information Statistics

The following information is displayed in a Detailed Process Information table:

Process	the name of the process, which is taken from its executed path name
PID	the number that identifies the process
PPID	<ul style="list-style-type: none"><li>• the number that identifies the parent process. The PPID can help identify possible relationships between processes</li><li>• on Windows systems, the PPID is called the Creating Process ID</li></ul>
UID	<ul style="list-style-type: none"><li>• the ID of the user or account that has been consuming CPU time</li><li>• on Windows systems, the UID is called the Owner</li></ul>
GID	<ul style="list-style-type: none"><li>• the ID of the group that has been consuming CPU time</li><li>• on Windows systems, the GID is called the Group Name</li></ul>

Memory Used	<ul style="list-style-type: none"> <li>the amount of memory, expressed as a percentage of total available memory, being consumed by a process</li> <li>on Windows systems, Memory Used is called Virtual Bytes</li> <li>the Memory Used value can be misleading because shared memory between processes is counted multiple times (e.g., if five Oracle processes are using 10% of available memory, this does not indicate that Oracle is consuming 50% of system memory)</li> </ul>
RSS	<ul style="list-style-type: none"> <li>Run Set Size – the amount of physical memory that is being used</li> <li>on Windows systems, RSS is called the Working Set</li> </ul>
CPU %	<ul style="list-style-type: none"> <li>the percentage of the CPU time used by the process, calculated by dividing total used CPU Time by the process' running time; if applicable, the result is further divided by the number of CPUs for the Element on which the process is running.</li> <li>on Windows systems, the CPU % is called % Processor Time</li> </ul>
Mem %	the percentage of the memory used by the process
Runtime	the total runtime of the process
Children Runtime	the runtime of all child processes
Start Time	the time at which the process started, which can be used to determine its lifetime

## Listing Detailed Process Information for an Agent-Based VM

To list Detailed Process Information for a VM Element, do the following:

- 1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VM whose process information you want to display, click its corresponding gear icon, then click **Graph Performance**.

- 2 In the Tree panel, click Detailed Proc Info in the Process section.**

- 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 23.

- 4 Click Display Process Information.**

A pop-up window will appear, displaying the process information for the specified time interval.

## Number of Processes

The Number of Processes graph displays the VM's process activity with the following details:

- number of processes
- running processes
- blocked processes
- waiting processes

Analyzing processes helps you determine whether or not there is enough CPU capacity for the processes that are being run on a system. If the size of the blocked or waiting queue is disproportionate to the running queue, then either the system does not have enough CPUs or is too I/O bound.

A blocked process signals a disk bottleneck. If the number of blocked processes approaches or exceeds the number of processes in the run queue, you should tune the disk subsystem. Whenever there are any blocked processes, all CPU idle time is treated as wait for I/O time. If database batch jobs are running on the system that is being monitored, there will

always be some blocked processes. However, you can increase the throughput of batch jobs by removing disk bottlenecks.

### Process Count Statistics

The following WMI- or agent-based information is displayed in a Process Count table:

Running	the number of processes that are currently running on a system, as reported by the system kernel
Blocked	the number of processes that are unable to complete, usually due to I/O operation
Waiting	the number of processes awaiting execution
Total Processes	the total number of running, blocked, and waiting processes

### Graphing Process Counts for an Agent-Based VM

To generate a Number of Processes graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VM whose process counts you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Process section of the Tree panel, click Number of Processes.**

**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 23.

**4 Select the metrics you want to include in the graph.**

**5 Click Generate Graph.**

A pop-up window will appear, displaying the process counts graph you have configured.



## Network Graphs

up.time uses the following graphs to chart network performance and usage:

- Network Workload
- Network I/O
- Network Errors

### Network Workload

Network Workload graphs help you gauge the demand being placed on the network, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers or clusters are transmitting and receiving the most data, or for an ESX server, which VMs' virtual NICs are the busiest.

#### Network Workload Metrics

For Network Workload graphs for a VMware vCenter server, the sum of KBps transmitted and received through NICs and vNICs can be generated for the following Element types:

- Datacenters
- Clusters
- ESX servers
- virtual machines
- resource pools
- vApps

For Network Workload graphs for an ESX server, the sum of KBps transmitted and received through virtual NICs can be generated for the following Element types:

- virtual machines
- resource pools
- vApps

### Graphing Network Workload for a vCenter or ESX Server

To generate a Network Workload graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click Network Workload (for a VMware vCenter server), or Workload in the Network section (for an ESX server).**

**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 23.

**4 Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select a VMware vSphere component whose metrics will be graphed.**

**6 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**7 Click Generate Graph.**

A pop-up window will appear, displaying the Network Workload graph you have configured.

## Network I/O

The I/O graph displays rates and total sums of data that are moving in and out of physical and virtual network interfaces over a specified time period. It can be used to find where I/O bottlenecks are occurring, and whether there have been any bursts in network traffic.

## Network I/O Metrics

The following metrics are displayed in a Network I/O graph:

ESX Server Metrics	
Total Rate	the sum of data transmitted and data received, in KBps, through the physical NICs on the ESX host
Transmit Rate	the average rate, in KBps, at which data was transmitted through each physical NIC on the ESX host
Receive Rate	the average rate, in KBps, at which data was received through each physical NIC on the ESX host
VM Metrics	
Total Rate	the sum of data transmitted and data received, in KBps, through the vNICs on the VMs on the ESX host
Transmit Rate	the average rate, in KBps, at which data was transmitted through each vNIC on the VM
Receive Rate	the average rate, in KBps, at which data was received through each vNIC on the VM
Total Rate (Agent)	the sum of data transmitted and data received, in KBps, through the guest NICs on the VMs on the ESX host, as reported via WMI or the <a href="#">up.time</a> agent
Send Rate (Agent)	the average rate, in KBps, at which data was transmitted through the guest NICs on the VMs on the ESX host, as reported via WMI or the <a href="#">up.time</a> agent
Receive Rate (Agent)	the average rate, in KBps, at which data was received through the guest NICs on the VMs on the ESX host, as reported via WMI or the <a href="#">up.time</a> agent

### Graphing Network I/O Rates for an ESX Server or VM

To generate a Network I/O graph for an ESX server or VM, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose network rates you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Network section of the Tree panel, click I/O.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 If you are graphing network I/O for a VM that receives metrics via WMI or from the [up.time](#) agent, select whether to graph network rates through virtual NICs or guest NICs.**

**5 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**6 Select the network metric to include in the graph.**

**7 Click Generate Graph.**

A pop-up window will appear, displaying the network I/O rate graph you have configured.

## Network Errors

The Network Errors graph displays the rate of errors with physical NICs on ESX servers, and vNICs on VMs.

The most common types of errors for physical connections include collisions in a hubbed environment or the presence of full-duplex handshake errors between a system and a switch. Additionally, the following communication line problems can cause network errors:

- excessive noise
- cabling problems
- problems with backbone connections

## Network Error Metrics

The following metrics are displayed in a Network Errors graph

ESX Server Metrics	
Total Dropped/min	the total number of dropped packets through the physical NICs on the ESX host, per minute, during the time period
Transmit Dropped/min	the number of packets outbound through the physical NICs on the ESX host that were dropped per minute during the time period
Receive Dropped/min	the number of packets inbound through the physical NICs on the ESX host that were dropped per minute during the time period
VM metrics	
Total Errors (Agent)	the sum of in errors and out errors through a guest NIC on the VMs on the ESX host
In Errors (Agent)	the number of packets received, but unable to be decoded due to a missing header or trailer
Out Errors (Agent)	the number of packets that were not sent, due to problems transmitting the packet or formatting the packet for transmission
Collisions (Agent)	the simultaneous presence of signals from two nodes on the network (e.g., from two nodes beginning to transmit over a network at the same time)
TCP Retransmits (Agent)	the number of TCP retransmits that occurred during the time period, due to lost or broken packets

### Graphing Network Error Rates for an ESX Server or VM

To generate a network error graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose network error rate you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Network section of the Tree panel, click Errors.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**5 Select the error metric to include in the graph.**

**6 Click Generate Graph.**

A pop-up window will appear, displaying the network error rate graph you have configured.

## Disk Graphs

up.time uses the following graphs to chart disk details, performance, and usage:

- Disk Workload
- Disk I/O
- Disk Errors
- Disk Latency
- Disk Storage Capacity
- Disk Storage Profile

### Disk Workload

Disk Workload graphs help you gauge the demand being placed on your physical or guest storage, and understand specifically from where it is coming. For example, for a VMware vCenter server, you can find out which ESX servers are experiencing the heaviest disk read/write activity, or for an ESX server, which VMs are writing the most to disk, prompting further investigation into VM read and write rates.

#### Disk Workload Metrics

For Disk Workload graphs for a VMware vCenter server, the aggregated disk I/O rate, in KBps, can be generated for the following Element types:

- Datacenters
- Clusters
- ESX servers
- virtual machines
- resource pools
- vApps

For Disk Workload graphs for an ESX server, the aggregated disk I/O rate, in KBps, can be generated for the following Element types:

- virtual machines
- resource pools
- vApps

### Graphing Disk Workload for a vCenter or ESX Server

To generate a Disk Workload graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VMware vCenter server or ESX server whose workload you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Tree panel, click Disk Workload (for a VMware vCenter server), or Workload in the Disk section (for an ESX server).**

**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 23.

**4 Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select a VMware vSphere component whose metrics will be graphed.**

**6 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**7 Click Generate Graph.**

A pop-up window will appear, displaying the Disk Workload graph you have configured.

## Disk I/O

The Disk I/O charts the rates at which physical disks on ESX servers, and virtual disks on VMs are reading and writing data. VMs that are receiving



metrics via WMI or the [up.time](#) agent report deeper metrics including wait times and data transfer speeds. These metrics can help you determine which disks are busiest.

## Disk I/O Metrics

The following metrics are displayed in a Disk I/O graph:

ESX server and VM Element Metrics	
Total RW Rate	the total read rate and write rate for virtual disks on VM, or LUNs on the ESX host
Read Rate	the rate, in KBps, at which data is read from each virtual disk on the VM, or each LUN on the ESX host
Write Rate	the rate, in KBps, at which data is written to each virtual disk on the VM, or each LUN on the ESX host
Total RW/min	the total number of read requests and write request from each virtual disk on the VM, or LUN on the ESX host
Read Requests/min	the number of times data was read from each virtual disk on the VM, or LUN on the ESX host
Write Requests/min	the number of times data was written to each virtual disk on the VM, or LUN on the ESX host
VM w/ agent Metrics on Guest Disks	
% Busy (Agent)	the percentage of guest disk capacity in use
Avg Queue Requests (Agent)	the average number of processes that are waiting to access the guest disk
Transfers/sec (Agent)	the total number of disk transfer requests processed per second
Blocks/sec (Agent)	the total number of data blocks written to and from the disk

Avg Service Time (Agent)	the average time, in milliseconds, required to perform a task
Avg Wait Time (Agent)	the average time, in milliseconds, that a transaction is waiting in a queue (note that the wait time is directly proportional to the length of the queue)

## Graphing Disk I/O Rates for an ESX Server or VM

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

**1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose disk I/O rate you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Disk section of the Tree panel, click I/O.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the first step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**6 Select the disk I/O metric to include in the graph.**

**7 Click Generate Graph.**

A pop-up window will appear, displaying the disk I/O graph you have configured.

## Graphing Disk I/O Rates for an Agent-Based VM

To generate a Disk I/O graph for a VM that receives metrics via WMI or the [up.time](#) agent, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose disk I/O rates you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Disk section of the Tree panel, click I/O.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 Select whether to graph a disk I/O rate on the VM disk, or a guest disk.**

**5 In the next step, select whether to graph the highest resource consumers, or specific VMs.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**6 Select the disk I/O metric to include in the graph.**

**7 Click Generate Graph.**

A pop-up window will appear, displaying the CPU Wait and Ready Time graph you have configured.

## Disk Errors

The Disk Errors graph charts the error rates on physical disks on ESX servers, or virtual disks on VMs. High error rates indicates performance issues with the underlying hardware.

## Disk Error Metrics

The following metrics are displayed in a Disk Errors graph:

Bus Resets/min	the number of SCSI bus reset commands issued on the virtual disk on the VM, or host disk
Commands Aborted/min	the number of SCSI commands on the virtual disk on the VM, or host disk, that were aborted

## Graphing Disk Error Rates for an ESX Server or VM

To generate a disk error rate graph, do the following:

**1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or VM whose disk error rates you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Disk section of the Tree panel, click Errors.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**5 Select the disk error metric to include in the graph.**

**6 Click Generate Graph.**

A pop-up window will appear, displaying the disk error rate graph you have configured.

## Disk Latency

The Disk Latency graph indicates the health and performance of physical storage on an ESX server.

## Disk Latency Metrics

The following metrics are displayed in a Disk Latency graph:

ESX Server Metrics	
Device Latency	the average amount of time, in milliseconds, required to complete a SCSI command from the physical device on the ESX host
Device Read Latency	the average amount of time, in milliseconds, required to complete reading from the physical device on the ESX host
Device Write Latency	the average amount of time, in milliseconds, required to complete writing to the physical device on the ESX host
Kernal Latency	the average amount of time, in milliseconds, VMKernel spent processing each SCSI command on the ESX host
Kernal Read Latency	the average amount of time, in milliseconds, VMKernel spent processing each SCSI read command on the ESX host
Kernal Write Latency	the average amount of time, in milliseconds, VMKernel spent processing each SCSI write command on the ESX host
Queue Latency	the average amount of time, in milliseconds, spent in the VMKernel queue, per SCSI command
Queue Read Latency	the average amount of time, in milliseconds, taken per SCSI read command
Queue Write Latency	the average amount of time, in milliseconds, taken per SCSI write command
Command Latency	the average amount of time, in milliseconds, to process a SCSI command issued by the guest OS to the VM
Command Read Latency	the average amount of time, in milliseconds, to process a read command issued from the guest OS to the VM

Command Write Latency	the average amount of time, in milliseconds, to process a write command issued from the guest OS to the VM
-----------------------	--

### Graphing Disk Latency for an ESX Server

To generate a percent-wait or percent-ready graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server whose disk latency you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Disk section of the Tree panel, click Latency.**

**3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

**4 Click the Quick Graph option to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**

**5 In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

**6 Select the latency metric to include in the graph.**

**7 Click Generate Graph.**

A pop-up window will appear, displaying the disk latency graph you have configured.

## Disk Storage Capacity

A Disk Storage Capacity graph charts the amount of total and used space, in gigabytes, on an Element's disk. This includes VMware vCenter servers, ESX servers, and VMs. For VMs that are receiving metrics from the up.time agent, the capacity of the file systems are also available.

The current capacity and usage for each datastore is displayed, as well as the week-over-week change, and estimated time to fill. You can select a datastore to graph to see datastore usage trends over time.

## Disk Storage Capacity Details

The following metrics are displayed in a Disk Storage Capacity graph:

Total Capacity	the total amount of space available on the system or VM
Used	the amount of space current being used
Free	the amount of free space
Change Per Week	a positive or negative value indicating the amount of free space compared to one week ago
Time to Fill	based on the changes during the last week, the estimated amount of time before the disk is filled

## Graphing Storage Capacities for vSphere Components

To generate a Disk Storage graph for a VMware vCenter Server, an ESX server, or a VM, do the following:

- 1 Go to the Element's Quick Snapshot page.**  
For example, in the **My Infrastructure** panel, find the Element whose storage capacity you want to graph, click its corresponding gear icon, then click **Graph Performance**.
- 2 In the Disk section of the Tree panel (for an ESX server or VM Element), or in the Datastores section (for a VMware vCenter Element), click Storage Capacity.**
- 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 23.
- 4 Select the datastores to include in the graph.**

If the Element being graphed is a WMI- or agent-based VM, you can also select the guest file system to include in the graph.

### 5 Click **Generate Graph**.

A pop-up window will appear, displaying the storage capacity graph you have configured.

## Disk Storage Profile

The Disk Storage Profile gives a breakdown of all datastores for an ESX server or VM. The current capacity and usage for each datastore is displayed. You can select a datastore to graph to see datastore usage trends over time.

### Disk Storage Profile Details

The following information is displayed in a Storage Profile graph:

- Total Capacity
- vDisk Usage
- Snapshot Usage
- Swap Usage
- Other Usage

### Graphing a Storage Profile for an ESX Server or VM

To generate a storage profile for an ESX server or VM, do the following:

#### 1 Go to the **Element's Quick Snapshot page**.

For example, in the **My Infrastructure** panel, find the ESX server or VM whose storage profile you want to display, click its corresponding gear icon, then click **Graph Performance**.

#### 2 In the **Disk** section of the **Tree** panel, click **Storage Profile**.

#### 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 23.



- 4 **Select the datastore you want to include in the graph.**
- 5 **Click Generate Graph.**

A pop-up window will appear, displaying the process counts graph you have configured.

## VMware Management Graphs

up.time uses the following graphs to track and manage VMs:

- VM Instance Motion
- Power Consumption
- Power States

### VM Instance Motion

The VMware VMotion tool enables you to move ESX instances from one server to another without any downtime or loss of data. For example, you might use VMotion to 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.

This graph can be generated when you are viewing a VMware vCenter server, ESX server, or a VM; the topological level at which you begin to configure the graph determines which instances are available to graph.

### Generating an Instance Motion Graph

To generate an Instance Motion graph, do the following:

**1 Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the VMware vCenter server or ESX server whose VMs' instance motion you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Metrics section of the Tree panel, click Instance Motion.**

**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 23.

**4 From the drop-down menu, select a VM.**

- 5 **Click Generate Graph.**

## Power Consumption

To assist virtualization initiatives that are meant to save power costs, or to gauge the efficiency of existing virtual datacenters, power usage by watts is available as a graphing metric.

The Power Consumption graph can tell you how much power your ESX hosts are consuming by datacenter, cluster, or individual server. Additionally, you can graph the power usage of individual VMs.

You can graph power usage levels at the VMware vCenter level in order to assess current load distribution, or verify VMware’s automated distributive resource balancing is functioning.

### Graphing Power Consumption for a vCenter Server

To generate a power consumption graph for a VMware vCenter server, do the following:

- 1 **Go to the Element’s Quick Snapshot page.**  
For example, in the **My Infrastructure** panel, find the VMware vCenter server whose power consumption you want to graph, click its corresponding gear icon, then click **Graph Performance**.
- 2 **In the Metrics section of the Tree panel, click Power Consumption.**
- 3 **Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**  
For more information, see “Understanding Dates and Times” on page 23.
- 4 **Select the logical group whose power consumption you want to graph.**
- 5 **In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**  
If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.
- 6 **Click Generate Graph.**

A pop-up window will appear, displaying the power consumption graph you have configured.

## Graphing Power Consumption for an ESX Server

To generate a power consumption graph for an ESX server, do the following:

**1 Go to the Element’s Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server or cluster whose power consumption you want to graph, click its corresponding gear icon, then click **Graph Performance**.

**2 In the Metrics section of the Tree panel, click Power Consumption.**

**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 23.

**4 Click Generate Graph.**

A pop-up window will appear, displaying the power consumption of the ESX server over the specified time period.

## Power States

Power States graphs allow you to graph VM activity on an ESX host, or ESX activity on a VMware vCenter server.

The power state graphs help you manage both available computing resources within your VMware vSphere clusters and datacenters, as well as power consumption in your physical datacenters.

### VMware vCenter Element Power States

The following states are displayed in a Power States graph:

vCenter Server Metrics	
VMs Powered On	the virtual machine is powered on

ESX Powered On	the host is powered on
ESX Powered Off	the host was powered off by an administrator through the VMware vSphere Client
ESX Unknown	as is the case in the VMware vSphere Client, a host that is in an unknown state is assumed to have been powered off by an administrator
ESX In Maintenance	the host was put in maintenance mode by an administrator
ESX In Standby	the host was put in standby mode either explicitly by an administrator, or automatically by vSphere Distributed Power Management (DPM)
<b>ESX Server Metrics</b>	
VMs Powered On	the virtual machine is powered on
VMs Powered Off	the virtual machine is powered off
VMs Suspended	the virtual machine is not running, but a snapshot of its running applications and processes is retained.

## Graphing Power States for a vCenter Server

To generate a Power States graph, do the following:

- 1 Go to the Element’s Quick Snapshot page.**  
For example, in the **My Infrastructure** panel, find the VMware vCenter server whose ESX server or VM power states you want to graph, click its corresponding gear icon, then click **Graph Performance**.
- 2 In the Metrics section of the Tree panel, click Power States.**
- 3 Select the start and end dates and times for which the graph will chart data, and click Apply Date and Time.**

For more information, see “Understanding Dates and Times” on page 23.

- 4 **Click one of the Quick Graphs options to display a pre-configured graph in a pop-up window, or skip this step to manually configure a graph.**
- 5 **In the first step, select whether to graph power states by datacenter or cluster.**
- 6 **In the next step, select whether to graph the highest resource consumers, or specific VMware vSphere components.**

If you select **Specific**, an Element selection dialog will appear, requiring you to build a list.

- 7 **Select the power states you want to include in the graph.**
- 8 **Click Generate Graph.**

A pop-up window will appear, displaying the CPU Wait and Ready Time graph you have configured.

## Graphing VM Power States for an ESX Server

To generate a power status graph for an ESX server's VMs, do the following:

- 1 **Go to the Element's Quick Snapshot page.**

For example, in the **My Infrastructure** panel, find the ESX server whose power states you want to graph, click its corresponding gear icon, then click **Graph Performance**.

- 2 **In the Metrics section of the Tree panel, click Power States.**
- 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 23.

- 4 **Select the power states you want to include in the graph.**
- 5 **Click Generate Graph.**

A pop-up window will appear, displaying the power states graph you have configured.

# CHAPTER 25

## Configuring and Managing up.time

---

The configuration and management of [up.time](#), mainly through the **Config Panel** and `uptime.conf` file, is described in the following sections:

<i>Overview</i> .....	638
<i>Interfacing with up.time</i> .....	642
<i>Archiving the DataStore</i> .....	655
<i>up.time Diagnosis</i> .....	661
<i>up.time Measurement Tuning</i> .....	664
<i>Report Storage Options</i> .....	668
<i>Resource Usage Report Generation</i> .....	670
<i>Monitoring Station Interface Changes</i> .....	671
<i>License Information</i> .....	673

# Overview

**up.time** includes user-definable parameters that can control some aspects of its behavior including the following:

- Database Settings
- Mail Server Settings
- Global Scan threshold settings
- Resource Scan threshold settings
- Proxy settings
- Remote reporting settings
- RSS feed settings
- Splunk integration settings
- Web monitor settings

From a configuration perspective, there are two types of parameters:

- parameters whose modification does not require a restart of the Core service (also known as the up.time Data Collector service); these parameters can be modified in **up.time**, on the **Config** panel
- parameters whose modification requires a restart of the Core service; these parameters are found in the `uptime.conf` file



## Modifying up.time Config Panel Settings

Configuration parameters that are not directly tied to, thus do not require a restart of, the **up.time** Core service can be modified directly in the **up.time** GUI (shown below):



In general, to edit these configuration settings in the **up.time** interface, do the following:

- 1 On the **up.time** tool bar, click **Config**.
- 2 In the **Tree** panel, click **up.time Configuration**.
- 3 Enter the configuration variable and new value.
- 4 Click **Update** to save your changes.



Only the variables whose default values have been modified appear in **up.time Configuration**.

## Modifying uptime.conf File Settings

Configuration parameters that are directly tied to the **up.time** Core service are found in the `uptime.conf` file. `uptime.conf` is a text file that you can modify in any text editor, and can be found in the root **up.time** installation directory.

In addition to the [up.time](#) database, `uptime.conf` parameters affect a variety of [up.time](#) behavior.



Not all of the settings listed in this section will necessarily be found in your particular `uptime.conf` file.

## Stopping and Restarting up.time Services

In addition to the Web interface, the [up.time](#) Monitoring Station consists of the following services:

- DataStore
- Web server
- Data Collector (also called the Core)

These services run in the background and start automatically after the operating system on the server hosting [up.time](#) starts. However, system administrators may need to stop the [up.time](#) services – for example, before making configuration changes to the `uptime.conf` file, performing an upgrade, or archiving the DataStore.

### Stopping the up.time Services

To stop the [up.time](#) services in Windows, do the following:

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

To stop the [up.time](#) services on Solaris or Linux, do the following:

- 1 Log into the Monitoring Station as user root.**

- 2 **Type the following command to stop the Web server:**

```
/etc/init.d/uptime_httpd stop
```

- 3 **Type the following command to stop the Data Collector:**

```
/etc/init.d/uptime_core stop
```

- 4 **Type the following command to stop the database:**

```
/etc/init.d/uptime_datastore stop
```

## Starting the up.time Services

To restart the [up.time](#) services in Windows, do the following:

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

To restart the [up.time](#) services on Solaris or Linux, do the following:

- 1 **At the command line, log into the Monitoring Station as user root.**
- 2 **Type the following command to start the database:**

```
/etc/init.d/uptime_datastore start
```
- 3 **Type the following command to start the Data Collector:**

```
/etc/init.d/uptime_core start
```
- 4 **Type the following command to start the Web server:**

```
/etc/init.d/uptime_httpd start
```

## Interfacing with up.time

Some of the Monitoring Station's features require integration with other elements that make up your infrastructure. In some cases configuration is mandatory (e.g., an SMTP server will need to have been set at the time of installation), while in others it is required only when particular [up.time](#) features are used (e.g., using the Web Application Transaction monitor requires you to provide [up.time](#) with your proxy server settings). The following sections outline how to configure [up.time](#) to communicate with servers and databases.

### 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. The supported drivers are:

  - `com.mysql.jdbc.Driver` (for MySQL)
  - `net.sourceforge.jtds.jdbc.Driver` (for SQL Server)
  - `oracle.jdbc.OracleDriver` (for Oracle)

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. 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 3308.
- `dbName=`  
The name of the database. The default is `uptime`.
- `dbUsername=`  
The name of the default database user, which is `uptime`.
- `dbPassword=`  
The password for the default database user, which is `uptime`.
- `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](#).
- `connectionPoolMaxIdleTime=`  
(c3p0 library) Sets the amount of time a connection can be idle before it is closed. This parameter should only be modified with the assistance of uptime software Customer Support.
- `connectionPoolNumHelperThreads=`  
(c3p0 library) Sets the number of helper threads that can improve the performance of slow JDBC operations. This parameter should only be modified with the assistance of uptime software Customer Support.

## Changing the DataStore Database

The [up.time](#) DataStore is first linked to a database during the installation process, and contains important historical performance data that has since been collected. Linking the DataStore to a new database will result in lost data unless you properly migrate your data to the new database. As such, changing the DataStore's database should be done only after some consideration and planning.

In cases where you would like to migrate the database (e.g., from the default [up.time](#) MySQL implementation to Oracle) or move the DataStore to a different system from the Monitoring Station, you will modify the aforementioned database values in the `uptime.conf` file. Note that the

modification of these values is one of a series of steps. Refer to the Knowledge Base for more information on migrating your DataStore.

### Monitoring Station Web Server

Monitoring Stations include a Web server component that drives the user interface. Any Monitoring Station that is accessed by users or administrators requires a URL. The Web address used to access the Monitoring Station is configured through the following `uptime.conf` parameter:

```
httpContext = http://<hostname>:<port>
```

- `<hostname>` is the host name of the server on which **up.time** is running (e.g., `localhost`)
- `<port>` is the port on which the **up.time** Web server is listening for requests (e.g., `9999`); you can optionally omit the port number

If the **up.time** interface is being accessed via SSL, the value for this parameter should be stated as `https` instead of `http`.

### SMTP Server

**up.time** uses a mail server to send alerts and reports to its users. After installing **up.time** for the first time, the administrator was asked to enter SMTP server information. These initial values can be modified in the **Mail Servers** configuration panel.

#### Modifying the SMTP Server Used by up.time

To configure **up.time**'s mail server, do the following:

- 1 **On the **up.time** tool bar, click **Config**.**
- 2 **In the **Tree** panel, click **Mail Servers**.**
- 3 **In the sub panel, click **Edit Configuration**.**
- 4 **Type the name of the mail server in the **SMTP Server** field.**

This value was set the first time the **up.time** administrator logged in after installation; the default value is the name of the host on which the Monitoring Station was installed at that time.

The name of the server could follow the “smtp.<domain\_name>” convention, or could be its host name or IP address.

- 5** **Optionally, enter the port used by the mail server in the SMTP Port field.**

- 6** **In the SMTP Sender field, enter the email address that up.time uses to send alert notifications and reports.**

This value was set the first time the **up.time** administrator logged in after installation, and should be set to your domain (e.g., admin@mail.uptimesoftware.com).

A sender’s name can be encapsulated with double quotes, in which case, the email address is encapsulated with angled brackets:

```
"uptime administrator" <admin@uptimesoftware.com>
```

- 7** **In the SMTP Helo String field, enter the string that identifies the domain from which a message is being sent.**

For example, uptimesoftware.com.

- 8** **In the SMTP User field, enter the user name that is used to authenticate connections with the SMTP server.**

- 9** **In the SMTP Password field, enter the password that is used to authenticate connections.**

- 10** **Click Save.**

The edit window closes, and you are returned to the **Mail Server Configuration** panel.

- 11** **To test the mail server configuration, click the Test Configuration button.**

The 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.

## Configuring Global Data Collection Methods

A Windows-based Element can retrieve metric data either through the up.time Agent, or via WMI. Initially set when the Element is added to up.time, the data collection method can be switched from an agent-based to agentless method, or vice versa. This change can be made on a per-Element basis, or multiple Elements can be switched in a single batch. (See “Agentless WMI Systems” on page 83 for more information.) In order to use the latter option, you must configure up.time so that it is aware of a data collection source that will be used for bulk conversions.

For configuration, you can provide information for either the up.time Agent, or your organization’s WMI credentials, or both. Note that multiple Windows-based Elements can only be converted to a particular data collection source when it has been configured in the Global Element Settings panel.

### Configuring Global WMI Credentials

To provide WMI credentials that can be used to switch Windows Elements from agent-based data collection:

- 1 On the up.time tool bar, click **Config**.
- 2 In the **Tree** panel, click **Global Element Settings**.
- 3 In the **WMI Agentless Global Credentials** sub panel, click **Edit Configuration**.
- 4 In the **Edit Global Element Settings** pop-up window, enter the **Windows Domain** in which WMI has been implemented.
- 5 In the **Username** field, enter the user ID that has administrative access to WMI on the Windows domain.
- 6 In the **Password** field, enter the password for the WMI account.
- 7 Click **Save** to retain your changes and close the pop-up window.
- 8 Click **Test Configuration** to ensure the credentials provided are correct.



## Configuring a Global up.time Agent Configuration

To provide up.time Agent information that can be used to switch Windows Elements from agentless, WMI-based data collection, do the following:

- 1 On the up.time tool bar, click **Config**.
- 2 In the Tree panel, click **Global Element Settings**.
- 3 In the up.time Agent Global Configuration sub panel, click **Edit Configuration**.
- 4 In the **Edit Global Element Settings** pop-up window, enter the port through which the up.time Agents communicate with the up.time Monitoring Station.



The port number entered reflects what the up.time Agents are configured to use; this setting does not modify the agent-side configuration.

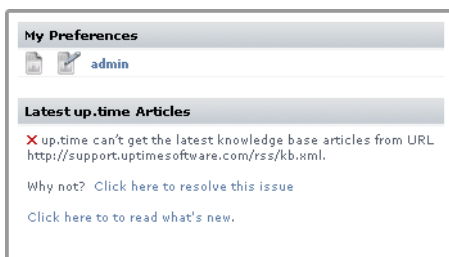
- 5 Select the **Use SSL** check box if the agents securely communicate with the Monitoring Station using SSL.
- 6 Click **Save** to retain your changes and close the pop-up window.
- 7 Click **Test Configuration** to ensure the credentials provided are correct.

## RSS Feed Settings

up.time displays a list of recent knowledge base articles in the **My Portal** panel. This list is fed to the **My Portal** panel via RSS (Really Simple Syndication, a method for delivering summaries of and links to Web content). Clicking the title of an article opens it in your Web browser.

By default, RSS feeds are drawn directly from the uptime software Support Portal without the use of proxy server information. If your Monitoring

Station accesses the Internet through one, feeds will most likely not be available, and the following message will appear in the **My Portal** panel:



You can change the RSS feed settings to point to the proxy server rather than directly to the uptime software Web site by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 639.

### Changing Proxy Server Information for RSS Feeds

You can manually configure the settings for RSS feeds through the following parameters (default values, if applicable, are shown):

- `rssFeedUrl=http://support.uptimesoftware.com/rss/kb.xml`  
The URL of the RSS feed.
- `httpProxyHost`  
The host name of the proxy server that the Monitoring Station uses to access the Internet.
- `httpProxyPort`  
The port through which the Monitoring Station communicates with the proxy server.
- `httpProxyUsername`  
The user name required to use the proxy server.
- `httpProxyPassword`  
The password required to use the proxy server.

## VMware vCenter Orchestrator Integration

Administrators can configure Action Profiles to automatically carry out tasks in the event of an up.time alert. One such task is the initiation of contact with VMware vCenter Orchestrator, and the execution of a workflow. To have access to this functionality, up.time needs to know how to communicate with Orchestrator.

For information about Action Profiles and VMware vCenter Orchestrator, see “Action Profiles” on page 441.

### Integrating up.time with VMware vCenter Orchestrator

To configure up.time integration with Orchestrator to execute workflows, do the following:

- 1 On the up.time tool bar, click **Config**.
- 2 In the **Tree** panel, click **VMware vCenter Orchestrator**.
- 3 In the sub panel, click **Edit Configuration**.
- 4 Ensure the **VMware Orchestrator Enabled** check box is selected.
- 5 In the **VMware Orchestrator Server** field, enter the host name of, or IP address assigned to the Orchestrator server when it was configured.
- 6 In the **VMware Orchestrator Port** field, enter the port the Orchestrator server was configured to use in order to communicate with other systems.
- 7 Optionally select the **Use SSL** check box if Orchestrator was configured to use an SSL certificate.
- 8 Enter the **Username and Password** of an appropriate user account on the Orchestrator server.

For proper integration, an Orchestrator account with View and Execute permissions is required.

- 9 Click **Save**.

The configuration window closes, and you are returned to the **VMware vCenter Orchestrator Configuration** panel.

- 10 To ensure the settings you provided are correct, click the Test Configuration button.**

The Monitoring Station will try to communicate with the VMware vCenter Orchestrator server. If an error message appears in the subpanel, edit and then re-test the configuration.

## Web Application Monitor Proxy Settings

When the Web Application Transaction monitor is recording a user session on an external site, it is intercepting URLs by acting as your browser's proxy. For the monitor to do this, you must replace your organization's proxy server information with the Web Application Transaction monitor in your browser settings. In order for the monitor to access the Internet, you must provide your proxy settings in [up.time](#).

This monitor-specific proxy information is used during transaction recording; during session playback, the proxy normally used by [up.time](#) (defined by the `httpProxy*` settings) is used.

For more information about the Web Application Transaction monitor, see “Web Application Transactions” on page 275.

You can change [up.time](#)'s proxy server configuration by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 639

### Changing Proxy Server Information for up.time

You can configure the proxy server settings used by [up.time](#) when running the Web Application Transaction monitor with the following parameters:

- `webmonitor.proxyHost`  
The host name of the proxy server that the Web Application Transaction monitor uses to access the Internet during transaction recording.
- `webmonitor.proxyPort`  
The port through which the Web Application Transaction monitor communicates with the proxy server during transaction recording.

## Remote Reporting Settings

If you are using a reporting instance (an [up.time](#) instance that only generates and serves reports), the remote reporting settings enable you to specify the location of the reporting instance, and the port on which it is listening.

### Modifying the Remote Reporting Server Settings

To configure the remote reporting instance used by [up.time](#), do the following:

- 1 On the [up.time](#) tool bar, click **Config**.
- 2 In the **Tree** panel, click **Remote Reporting**.
- 3 In the sub panel, click **Edit Configuration**.
- 4 Ensure the **Reporting Instance Enabled** check box has been selected.
- 5 In the **Remote Reporting Server** field, enter the host name or IP address of the server on which the remote reporting instance is found.
- 6 Enter the port used to communicate with the server.
- 7 Click **Save**.

The edit window closes, and you are returned to the **Remote Reporting Instance Configuration** panel.

- 8 To test the remote reporting server configuration, click **Test Configuration**.

A pop-up window appears, indicating whether [up.time](#) was able to connect to the remote reporting instance. If an error message is displayed, correct your configuration and re-test it.

Note that the modification of these values is one of a series of steps performed to correctly set up a remote reporting instance. Refer to the Knowledge Base article entitled “*Setting up a reporting instance*” for more information.

## User Interface Instance Settings

A UI instance is an **up.time** installation that does not perform any data collection tasks, and is primarily used for real-time monitoring and report generation. UI instances can divert traffic from a standard Monitoring Station implementation, and are helpful when there are many **up.time** users who do not need to perform full administrative tasks.

You can manually configure UI instance settings with the following `uptime.conf` parameters:

- `uiOnlyInstance = true`  
Determines whether the Monitoring Station functions only as a user interface instance.
- `uiOnlyInstance.monitoringStationHost = HOSTNAME`  
The host name or IP address of the **up.time** Monitoring Station that is performing data collection, and to which this UI instance will connect.
- `uiOnlyInstance.monitoringStationCommandPort = 9996`  
The port through which the UI instance can communicate with the data-collecting Monitoring Station.

A Monitoring Station that is acting as a UI instance must have the same database settings as the data-collecting Monitoring Station. See “Database Settings” on page 642 for more information.

## Scrutinizer Settings

Scrutinizer is a NetFlow analyzer that can be installed to monitor network traffic managed by compatible switches and routers. Scrutinizer can be integrated with **Global Scan**, as well as **up.time**’s graph generation for node-type Elements, and other hosts that are also monitored with Scrutinizer.

In order to access Scrutinizer, **up.time** needs to be pointed to your installation.

## Modifying the Scrutinizer Settings

You can configure Scrutinizer's integration with [up.time](#) through the following parameters:

- `netflow.enabled`  
Determines whether Scrutinizer is integrated with the Monitoring Station.
- `netflow.hostname`  
The host name or IP address of your Scrutinizer installation.
- `netflow.port`  
The HTTP port through which Scrutinizer sends and receives communication.
- `netflow.username`  
The user name required to log in to Scrutinizer.
- `netflow.password`  
The password required to log in to Scrutinizer.

## Splunk Settings

Splunk is a third-party search engine that indexes log files and data from the devices, servers, and applications in your network. Using Splunk, you can quickly analyze your logs to pinpoint problems on a server or in a network, or ensure that you are in compliance with a regulatory mandate or Service Level Agreements. You install Splunk on a server in your datacenter.

When values are provided for the Splunk settings listed below, the Splunk icon ( `splunk>` ) will appear in the **My Portal** panel beside the names of services that are in WARN or CRIT states. When you click the Splunk icon, you will be automatically logged in to your Splunk search page.

You can change your [up.time](#)-Splunk integration by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 639.

### Changing Splunk Server Information for up.time

You can enable automatic login to the Splunk search page, or modify an existing configuration through the following parameters:

- `splunk.url`  
The URL of the server on which your Splunk search page is hosted (e.g., `http://webportal:8000`).
- `splunk.username`  
The user name required to log in to your Splunk search page.
- `splunk.password`  
The password required to log in to your Splunk search page.
- `splunk.soapurl`  
The URL that points to the SOAP management port that Splunk uses to communicate with the splunk daemon (e.g., `https://webportal:8089`).  
In the URL, you must include the port on which the Splunk server listens for requests. See the Splunk Admin Manual for more information.



## Archiving the DataStore

Depending on the amount of disk space available for the continuously growing DataStore, administrators can set an archive policy that determines how many month's worth of data is retained. Old performance data is automatically archived and removed from the DataStore. This archiving procedure works with all databases that are compatible with [up.time](#).

The existing archive policy can be viewed and modified on the **Archive Policy** subpanel, which is accessed from the main **Config** panel. Here, the main archive categories are shown, along with the number of months for which collected data is retained in the DataStore.

Every month, [up.time](#) checks the DataStore's entries; data that is older than the limit set in the archive policy are written to XML files. The XML archives use the following format:

```
<table_name>_<date>.xml.gz
```

The archives created reflect the database table structure used to store performance data, as well as the date that the stored data represents:

```
performance_cpu_2006-09-13.xml.gz
```

The DataStore is trimmed and the XML files are compressed and stored in the `/archives` directory.

For example, if you installed [up.time](#) in the default location, the path to the archived data will be:

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



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

Once backed up, archives can be stored offline. If required, they can be temporarily imported into the DataStore.

## Archive Categories

The following table lists the statistical categories whose archiving can be configured, along with the corresponding DataStore database table:

Archive Policy Category	Database Table
Overall CPU/Memory	performance_cpu
Multi-CPU	performance_aggregate
Detailed Process	performance_psinfo
Disk Performance	performance_disk
File System Capacity	performance_fscap
Network	performance_network
User Information	performance_who
Volume Manager	performance_vxvol
Retained Data	erdc_int_data erdc_decimal_data erdc_string_data

## Configuring an Archive Policy

To set an archive policy, do the following:

- 1 **On the up.time tool bar, click Config.**
- 2 **In the Tree panel, click Archive Policy.**
- 3 **For the following categories, specify the number of months worth of data that will be retained in the DataStore before being removed and archived:**
  - Overall CPU/Memory Statistics
  - Multi-CPU Statistics
  - Detailed Process Statistics

- Disk Performance Statistics
  - File System Capacity Statistics
  - Network Statistics
  - User Information Statistics
  - Volume Manager Statistics
  - Retained Data
- 4 **Ensure the Enable Archiving checkbox is selected.**
  - 5 **Click Set Archive Policy.**
  - 6 **Optionally, you can click the Archive Now button to immediately create archives of the data in your DataStore.**

[up.time](#) will check the DataStore entries and archiving anything that is older than the limits you have configured.

## Restoring Archived Data

If you need to generate graphs or reports on older data that has already been archived, and is no longer in the DataStore, you can import specific archives using the `restorearchive` command line utility. The command's parameters allow you to import archives in the following manner:

- a single archive that represents a specific archive category and date; the collected data for each archive category and 24-hour period is exported to individual XML files
- all archives for a specific date (i.e., 24-hour period)

## Importing Archived Data into the DataStore

To import archived data into the DataStore, do the following:

- 1 **At the command line, navigate to the following directory:**
  - Linux: `/usr/local/uptime/scripts/`
  - Solaris: `/opt/uptime/scripts/`

- Windows: C:\Program Files\uptime software\uptime\archives

### 2 Run the `restorearchive` command with one or more of the following options:

- `-f <filename>`  
Imports a single file (i.e., an archive category's data for a single date). You must specify the full path to the file name.
- `-d <date>`  
Imports all files with the specified date (in YYYY-MM-DD format).
- `-D <directory>`  
The directory containing the archived files. Note that you must specify this option when using the `-d` option.
- `-c <directory>`  
The full directory path to the file `uptime.conf`.

For example, enter the following command to import all of the data archived on September 18, 2006 which are located in the default directory for archived data:

```
restorearchive -d 2006-09-18 -D /usr/local/uptime/archives/ -c /usr/local/uptime
```

## Exporting and Importing the DataStore

In cases where you need to perform a wholesale backup of the existing DataStore (e.g., migrating your DataStore to another database), `up.time` includes two command line utilities:

- `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 **Ensure the database hosting the DataStore is running.**
- 2 **Stop the `up.time` Data Collector service.**
- 3 **Navigate to the scripts folder under the directory where `up.time` is installed.**
- 4 **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.



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

## Restoring the DataStore

To restore your DataStore, do the following:

- 1 **Ensure the database hosting the DataStore is running.**
- 2 **Use the `resetdb` utility with the `really` and `nodata` options to delete, then recreate the database structure that is used by `up.time` by running one of the following commands:**
  - Linux: `/usr/local/uptime/resetdb --nodata really`
  - Solaris: `/opt/uptime/resetdb --nodata really`
  - Windows: `C:\Program Files\uptime software\uptime\resetdb --nodata 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 an archive that is located in **up.time**'s root installation directory, enter the following:

```
fulldatabaseimport uptimedump_2007-01-02.xml.gz
```



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

## up.time Diagnosis

The following options assist you with diagnostic steps that you may need to perform should you encounter problems with **up.time**. You have access to two types of logs: system logs and audit logs that track user actions. Additionally, you can generate a problem report for **up.time** Customer Support if further analysis is required.

System and audit logs are written to the `/logs` directory, and problem reports are found in the `/GUI` directory, both of which are found in the **up.time** installation directory:

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



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

## System Event Logging

**up.time** automatically logs system events to the `/logs` directory. These weekly logs follow the `uptime.log.<year>--<week>.log` naming format. You can determine the type of system information **up.time** writes to the log by using one of the following values:

- DEBUG
- INFO
- WARN
- ERROR
- FATAL
- ALL

- OFF

The default setting, `DEBUG`, essentially logs all system event types. To reduce the number of log entries, you can limit logging to events with a higher level of severity, from `INFO` to `FATAL`. Note that each severity level is a subset of higher levels (e.g., setting `loggingLevel` to `WARN` means any `WARN`-, `ERROR`- or `FATAL`-level events are written to the log).

Logging is configured through the following `uptime.conf` parameter:

```
loggingLevel = DEBUG
```

## Audit Logs

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

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

The following is an example of an audit log entry:

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

Audit Logging is enabled or disabled, with “yes” or “no” values, respectively, through the following `uptime.conf` parameter:

```
auditEnabled = yes
```

## Problem Reporting

When you encounter a problem with `up.time`, Client Care needs specific information to diagnose and fix the problem. `up.time` can automatically collect this information and compress it in an archive which you can send to Client Care.



The archive contains the following: **up.time** configuration files; system information; log files; database information and error files; and a listing of the DataStore directory. Optionally, the archive will also contain a copy of the configuration data from your DataStore.

The archive is saved to the `GUI/problemreports` directory on the Monitoring Station and has a file name with the following format:

`prYYYYMMDD-HHMMSS.zip`

- YYYYMMDD is the date on which the report was generated (e.g., 20061212).
- HHMMSS is the time at which the report was generated (e.g., 142306).

## Generating a Problem Report

To generate a problem report, do the following:

- 1 **On the **up.time** tool bar, click **Config**.**
- 2 **In the **Tree** panel, click **Problem Reporting**.**

If you have generated problem reports in the past, they appear in the subpanel.

- 3 **If you do not want to include a copy the configuration data from your DataStore, click the **Include config database dump** option.**
- 4 **Click the **Generate Report** button.**

A message such as the following appears in the subpanel:

Problem report created : `pr20061017-094927.zip`

Click the name of the problem report to download it to your local file system, then send the archive to uptime software Client Care.

## up.time Measurement Tuning

In some cases, you can make measurement adjustments to **up.time**'s default values. Changes can be made to the following:

- the number of threads allocated to service monitors
- status thresholds in the **Resource Scan** and **Global Scan** panels
- how often performance and status are checked for monitored hosts

### Service Monitor Thread Counts

By default, the number of Java threads allocated to service and performance monitors is 100. This can be modified with the following `uptime.conf` parameter:

```
serviceThreads = 100
```

### Status Thresholds

The **Global Scan** threshold settings determine when a cell in the **Global Scan** panel changes state to reflect a host's status change: green represents normal status, yellow represents Warning status, and red represents Critical.

The Resource Scan threshold settings determine the size of the gauge ranges on the **Resource Scan** view: green represents normal status, yellow represents Warning status, and red represents Critical status.

You can change the thresholds used to determine status by manually inputting settings in the **up.time Configuration** panel, as outlined in "Modifying up.time Config Panel Settings" on page 639.



Changes to Global Scan thresholds are not retroactively applied to all Elements; only Elements added after threshold changes will reflect those changes.

## Changing Global Scan Threshold Settings

You can modify the **Global Scan** threshold settings through the following parameters (default values are shown):

- `globalscan.cpu.warn=70`  
A Warning-level status is reported when CPU usage is at 70% or greater.
- `globalscan.cpu.crit=90`  
A Critical-level status is reported when CPU usage is at 90% or greater.
- `globalscan.diskbusy.warn=70`  
A Warning-level status is reported when a disk on the host is busy for 70% or more of a five-minute time frame.
- `globalscan.diskbusy.crit=90`  
A Critical-level status is reported when a disk on the host is busy for 90% or more of a five-minute time frame.
- `globalscan.diskfull.warn=70`  
A Warning-level status is reported when 70% or more of the disk space on the host is used.
- `globalscan.diskfull.crit=90`  
A Critical-level status is reported when 90% or more of the disk space on the host is used.
- `globalscan.swap.warn=70`  
A Warning-level status is reported when 70% or more of the swap space on a disk is in use.
- `globalscan.swap.crit=90`  
A Critical-level status is reported when 90% or more of the swap space on a disk is in use.

### Resource Scan Threshold Settings

You can modify the **Resource Scan** threshold settings through the following parameters (default values are shown):

- `resourcescan.cpu.warn=70`  
The Warning-level range in the **CPU Usage** gauge begins at this value (70%), and ends at the Critical-level range.
- `resourcescan.cpu.crit=90`  
The Critical-level range in the **CPU Usage** gauge is between this value (90%) and 100%.
- `resourcescan.memory.warn=70`  
The Warning-level range in the **Memory Usage** gauge begins at this value (70%), and ends at the Critical-level range.
- `resourcescan.memory.crit=90`  
The Critical-level range in the **Memory Usage** gauge is between this value (70%) and 100%.
- `resourcescan.diskbusy.warn=70`  
The Warning-level range in the **Disk Busy** gauge begins at this value (70%), and ends at the Critical-level range.
- `resourcescan.diskbusy.crit=90`  
The Critical-level range in the **Disk Busy** gauge is between this value (70%) and 100%.
- `resourcescan.diskcapacity.warn=70`  
The Warning-level range in the **Disk Capacity** gauge begins at this value (70%), and ends at the Critical-level range.
- `resourcescan.diskcapacity.crit=90`  
The Critical-level range in the **Disk Capacity** gauge is between this value (70%) and 100%.

## Platform Performance Gatherer Check Intervals

The Platform Performance Gatherer is a core performance monitor that resides on all agent-based Elements. (See “The Platform Performance Gatherer” on page 162 for more information.)

By default, the Platform Performance Gatherer checks the host Elements’ performance levels every 300 seconds. You can change the interval by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 639

### Changing the Performance Monitor Check Interval

You can modify the Platform Performance Gatherer check interval through the following parameter (the default value is shown):

```
performanceCheckInterval = 300
```



A change to the Platform Performance Gatherer check interval is not retroactively applied to all Elements; only Elements added after an interval change will reflect that change.

## Report Storage Options

When an **up.time** user generates a report, that report is stored in the /GUI/reportcache directory; when a scheduled report is automatically generated and published, it is stored in the /GUI/published directory. Both of these directory paths are found in the **up.time** installation directory:

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



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

By default, generated reports are cached on the Monitoring Station for 30 days; additionally, the location for published reports is also on the local Monitoring Station file system. Both options can be modified. In the latter case, automatically publishing reports to a publicly accessed directory on the network is an ideal way for non-IT staff to view them. See “Saving Reports to the File System” on page 456 for more information.

## Changing the Number of Days Reports Are Cached

You can change a report’s expiry time limit by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 639.

Change the expiry limit through the following parameter (the default value is shown):

```
reportCacheExpiryDays=30
```

## Changing the Published Report Location

This can be modified with the following `uptime.conf` parameter:

```
publishedReportRoot=<location>
```

If the intended published report directory is on a system other than the Monitoring Station, the provided location should be a full network path to the system in addition to the directory path on that system.

## Resource Usage Report Generation

Due to the large number of options available for the Resource Usage report, generating an extensive report for a large group of Elements can take several minutes. If exhaustive report generation is necessary, but taking too long, you can increase the number of report images (the default being “6”) that **up.time** concurrently generates for this type of report.

Note that the default number is optimal in most cases; increasing the amount may improve performance, but the law of diminishing returns applies, as too many concurrent threads can tax the PDF generation process overall.

Logging is configured through the following `uptime.conf` parameter:

```
reporting.prefetch.images.threads = 6
```



## Monitoring Station Interface Changes

Some configuration options affect the Monitoring Station interface. These can be modified by manually inputting settings in the **up.time Configuration** panel, as outlined in “Modifying up.time Config Panel Settings” on page 639.

### Status Alert Acknowledgement

When services reach a warning or critical state, administrators can flag an alert as “acknowledged,” which prevents subsequent alerts from being broadcasted, giving them time to investigate the issue. See “Acknowledging Alerts” on page 115 for more information.

Service status alert acknowledgements can be reported in the status tables on the **Global Scan** panel. By default, status alert acknowledgement counts are not shown; if enabled a new column (labelled ACK) appears in the **Service Status** section of **Global Scan**. When the current status of a monitor is acknowledged, it appears in the ACK column instead of in the WARN or CRIT column.

You can enable or disable status acknowledgement (i.e., add or remove the ACK column from the status tables) through the following parameter (the default value is shown):

```
acknowledgedSeparate=false
```

### 3D Graphs

When performance and availability graphs are generated, the Graph Editor is used to manipulate the appearance of graphed data (see “Using the Graph Editor” on page 538. Transformations from a three-dimensional perspective are possible if the user account permits it (see “Adding Users” on page 389), and the user is connecting to the Monitoring Station using Internet Explorer.

This 3D presentation option can be disabled outright. You can determine whether ActiveX graphs are displayed in 3D for users with Internet Explorer through the following parameter (the default value is shown):

```
default3DGraphs=true
```

## Custom Dashboard Tabs

Custom dashboards can be added to **My Portal** to display custom content that is relevant to the particular user who is currently logged in. Up to 50 dashboards can be added, each of which is accessed through, and viewed in, its own tab at the top of **My Portal**.

A custom dashboard tab is configured by pointing [up.time](#) to a custom Web page, and indicating which User Group will be able to view it. You can enable and configure the first dashboard through the following parameters:

```
myportal.custom.tab1.enabled=true  
myportal.custom.tab1.name=<DashboardNameOnTab>  
myportal.custom.tab1.URL=<URLtoCustomPage>  
myportal.custom.tab1.usergroups=<UserGroupName>
```

Values for the first three parameters are required. If no name is specified for the User Group parameter (or, if no User Groups have been defined), the custom dashboard will be visible to all [up.time](#) users. Thus, a User Group parameter is only required if you want to restrict or refine user access to a particular custom dashboard.

To create additional tabs, add the same set of parameters, but increment the tab count:

```
myportal.custom.tab2.enabled=true  
myportal.custom.tab2.name=<DashboardNameOnTab>  
myportal.custom.tab2.URL=<URLtoCustomPage>
```

## License Information

If your [up.time](#) package did not come with a license key, then either contact your sales representative to request a key or send an email to [support@uptimesoftware.com](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 Information** 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.

### Installing or Updating a License

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 License.**

### Setting a Notification Group for vSync-Related Licensing Errors

In the **License Notification** section of the **License Information** page, you can select the Notification Group that receives alerts should there be any licensing errors related to syncing with VMware vSphere.

For more information, see “Managing vSync” on page 178.



# APPENDIX A

## Reference

---

This appendix contains the following sections:

<i>Frequency Definitions</i> .....	676
<i>Time Period Definitions</i> .....	677

## Frequency Definitions

To define synchronization frequencies in [up.time](#), you input a string that represents the amount of time between actions. These units of time can be days, hours, minutes, seconds, or a combination. Frequency definitions are used when configuring user detail synchronization, when configuring up.time to use an Active Directory or LDAP listing for user authentication and management. (See “Changing How Users Are Authenticated” on page 401 for more information.)

All time units are represented by a one-letter abbreviation:

- days: d
- hours: h
- minutes: m
- seconds: s

Frequency definitions can be a combination of any of these time units and their values, in descending order, without spaces:

- 1d
- 1d12h
- 1h30m
- 30s

## Time Period Definitions

When defining new, or editing existing, Maintenance Profiles and Monitoring Periods, you need to use precise definitions that **up.time** can correctly interpret. Time period definitions use a controlled vocabulary that allow you to precisely define, combine, and exclude time periods.



Although all examples listed in the following sections are written in mixed case (e.g., “Every Oct 28”), none of the terms used in time period definitions is case sensitive.

### Building Blocks

The following tables outline the basic components of all time period definitions.

#### Time Units

- Units of time that act as building blocks in definitions include times of day, days of the week, months, years, and exact dates.

Times		
Required	<ul style="list-style-type: none"> <li>• hour of day</li> <li>• 12-hour clock suffix, inputted as “AM” or “PM”</li> </ul>	correct: 8:00 PM
Optional	<ul style="list-style-type: none"> <li>• minutes of the hour</li> <li>• spaces</li> </ul>	correct: 8 PM, 8:00PM, 8PM
Not Accepted	<ul style="list-style-type: none"> <li>• missing 12-hour clock suffix</li> <li>• 24-hour clock convention</li> </ul>	incorrect: 8:00 20:00, 20:00 PM

## Time Period Definitions

<b>Days</b>		
Required	three-letter abbreviation	correct: Sun, Mon, Tue Wed, Thu Fri, Sat
Not Accepted	<ul style="list-style-type: none"> <li>• full spellings</li> <li>• other abbreviation styles</li> </ul>	incorrect: S, M, T We, Th Friday, Saturday
<b>Dates</b>		
Required	single- or two-digit number	correct: 8, 09, 10
Not Accepted	<ul style="list-style-type: none"> <li>• ordinal suffixes</li> <li>• full spellings</li> </ul>	incorrect: 8th, 9th, tenth
<b>Months</b>		
Required	three-letter abbreviation	correct: Jan, Feb, Mar, Apr May, Jun, Jul, Aug Sep, Oct, Nov, Dec
Not Accepted	other abbreviation styles	incorrect: J, F, M, A June, July, August Se, Oc, No, De
<b>Years</b>		
Required	full year	correct: 2008
Not Accepted	any abbreviation of the year	incorrect: 08, '08, Y2K+8



## Lists and Ranges

Days can be inputted as a list:

- each day is separated by a comma (e.g., “mon, tue, wed”)
- spaces are optional (e.g., “mon,tue,wed”)

Times and days can be inputted as ranges:

- Elements in the range must be separated by hyphens
- spaces are optional; the following examples are correct:
  - 8AM-8PM
  - 8:00 AM - 8:00 PM
  - Fri - Mon
  - Fri-Mon
- ranges wrap around day and week boundaries:
  - “10PM - 2AM” is interpreted as 10:00 p.m. to 11:59 p.m. on one day, and 12:00 a.m. to 2:00 a.m. the following calendar day
  - “Fri-Mon” is interpreted as Friday through Saturday on one week, then Sunday through Monday the following week
- **up.time** converts day ranges to lists (e.g., “Fri-Mon” becomes “Fri, Sat, Sun, Mon”)
- day ranges and lists can be mixed; the following examples are correct:
  - Fri - Sun, Mon
  - Fri-Sun,Mon

## Basic Expressions

Using the building blocks outlined in the previous section, use the following templates to create basic expressions that are used to define time periods in **up.time**. Note that shaded components of a template are optional.

## Time Period Definitions

### Fixed Dates

<month>	<date>	,	<year>	<time range>
---------	--------	---	--------	--------------

**Basic example:**

Oct 28, 2008

**Spaces are optional:**

Oct28,2008

**Time ranges are optional:**

Oct 28, 2008 7 PM - 11 PM

Oct28,20087PM-11PM

**Note:** Fixed dates that do not include a time range are interpreted to include the entire day (i.e., 12:00 a.m. through 11:59 p.m.), although this will not automatically appear in the defined time period.

### Fixed Date Ranges

from	<month>	<date>	<year>	<time range>
to	<month>	<date>	<year>	<time range>

**Basic example:**

From Oct 28, 2008 to Oct 29, 2008

**Spaces are optional:**

FromOct28,2008toOct29,2008

**Time ranges are optional:**

From Oct 28, 2008 7 PM to Oct 29, 2008 2 AM

**Note:** A fixed date without a time that is at the end of a date range is interpreted to include the first minute of the next day (e.g., up.time converts “From Oct 28, 2008 to Oct 29, 2008” into “From Oct 28, 2008 12:00AM to Oct 30, 2008 12:00AM”).

**Note:** The time range in a fixed date range merely acts as a more precise start point and end point; a fixed date range is a contiguous block of time that has no gaps.

## Weekly Recurrence

every	<day> / <day range / list>	<time range>
-------	----------------------------	--------------

**Basic example:**

Sun  
Sun - Tue  
Every Sun, Mon, Tue

**Spaces are optional:**

Sun-Tue  
EverySun, Mon, Tue

**Time ranges are optional:**

Sun 9 AM - 5 PM  
Sun - Tue 9AM - 5PM  
EverySun, Mon, Tue9AM-5PM

**Note:** Recurring days that do not include a time range are interpreted to include the entire day (i.e., 12:00 a.m. through 11:59 p.m.), although this will not automatically appear in the defined time period.

## Yearly Recurrence

every	<month>	<date>	<time range>
-------	---------	--------	--------------

**Basic example:**

Every Oct 28

**Ordinal suffixes are optional:**

Every Oct 28th

**Time ranges are optional:**

Every Oct 28 7PM - 11PM

**Note:** You cannot define a date range within a yearly recurrence; instead, combine a separate yearly recurrences for each date in the date range.

## Monthly Recurrence

every month on the	<date>	<time range>
--------------------	--------	--------------

### Basic example:

Every month on the 28

### Ordinal suffixes are optional:

Every month on the 28th

### Time ranges are optional:

Every month on the 28 6 PM - 11 PM

Every month on the 28th 6PM-11PM

## Monthly Ordinal Recurrence

every month on the	<ordinal_as_word>	<day>	<time range>
-----------------------	-------------------	-------	--------------

### Basic example:

Every month on the last Fri

### Time ranges are optional:

Every month on the last Fri 6 PM - 11 PM

Every month on the last Fri 6PM-11PM

**Note:** The ordinal must be stated as a word: first, second, third, fourth, and last.

## Combining Expressions and Excluding Time Periods

Elaborate time period definitions are built from a combination of the basic expressions defined in the previous section:

- fixed dates
- fixed date ranges

- weekly recurrences
- monthly recurrences
- monthly ordinal recurrences
- yearly recurrences

## Combinations

Combine basic expressions by writing each one on a new line in the **Definition** box when defining a Maintenance Profile or Monitoring Period. The following examples demonstrate combinations of different basic expressions used to define a maintenance window:

Combining fixed dates:

```
Dec 25, 2008 12AM - 12PM  
Jan 1, 2009 12AM - 12PM
```

Combining a fixed date and a fixed date range:

```
Dec 25, 2008 12AM - 12PM  
From Dec 31, 2008 11PM to Jan 1, 2009 12PM
```

Combining weekly recurrences:

```
Mon-Fri 1AM-3AM  
Sat 1AM-5:30AM  
Sun
```

Combining yearly recurrences:

```
Every Dec 25 12AM-12PM  
Every Dec 31 11PM-11:59PM  
Every Jan 1st 12AM-12PM
```

Combining monthly recurrences:

```
Every month on the 2  
Every month on the 16th
```

Combining monthly ordinal recurrences:

```
Every month on the first Fri  
Every month on the third Fri  
Every month on the last Fri
```

## *Time Period Definitions*

Note that when a time period consists of more than one component time period expression, a condition met within *any* of those component time periods applies to the entire time period. For example, if a Monitoring Period named “Open Hours” is defined as:

```
Mon-Fri 9AM-5PM
Sat 10AM-5PM
Sun 12PM-5PM
```

An alert-worthy event that occurs on Sunday at 1:00 p.m. means the entire time period definition has been fulfilled.

## **Exclusions**

Time periods can be excluded from greater time period definitions by using the term “exclude” as a prefix to the exclusionary expression. The following examples demonstrate the use of exclusions in time periods:

Excluding a monthly recurrence from a weekly recurrence:

```
Sun 3PM-5PM
Exclude every month on the last Sunday
```

Defining two yearly recurrences to exclude from a weekly recurrence:

```
Mon-Fri 2AM-3AM
Exclude every Jan 1
Exclude every Jan 2
```

# APPENDIX B

## End User License Agreement

---

Before downloading [up.time](#), obtaining a license key, or using [up.time](#), please read the following End User License Agreement for [up.time](#). The [up.time](#) End User License Agreement defines the rights, permissions, and limitations that you agree to by choosing [up.time](#).

The [up.time](#) End User License Agreement is detailed in the following sections:

<i>NOTICE TO USER..</i>	<i>686</i>
<i>License.....</i>	<i>686</i>
<i>Intellectual Property and Confidentiality.....</i>	<i>688</i>
<i>License Fees .....</i>	<i>689</i>
<i>Term and Termination.....</i>	<i>690</i>
<i>Remedies and Indemnification .....</i>	<i>690</i>
<i>Disclaimer.....</i>	<i>691</i>
<i>Limitation of Liability .....</i>	<i>691</i>
<i>General Terms .....</i>	<i>692</i>

## NOTICE TO USER

This End User License Agreement (the “Agreement”) is a legal contract between you, as either an individual or a business entity, and Uptime Software Inc. (“Uptime”).

PLEASE READ THIS CONTRACT CAREFULLY BEFORE DOWNLOADING UPTIME’S PROPRIETARY SOFTWARE (the “SOFTWARE”) OR OBTAINING A LICENSE KEY TO THE SOFTWARE OR USING THE SOFTWARE. BY CLICKING ON THE “I ACCEPT” BUTTON AND BY DOWNLOADING THE SOFTWARE OR OBTAINING A LICENSE KEY TO THE SOFTWARE YOU REPRESENT AND WARRANT THAT YOU ARE EITHER THE REPRESENTATIVE OF THE COMPANY WITH THE AUTHORITY TO ENTER INTO THIS AGREEMENT AND TO BIND THE COMPANY OR YOU ARE AN INDIVIDUAL OVER THE AGE OF 18 AND THAT YOU HAVE READ THIS AGREEMENT, THAT YOU UNDERSTAND IT, AND THAT YOU ACCEPT AND AGREE TO BE BOUND BY ITS TERMS. IF YOU ARE UNWILLING TO BE BOUND BY THE TERMS OF THIS AGREEMENT YOU SHOULD CLICK THE “I DO NOT ACCEPT” BUTTON BELOW, TERMINATE THE DOWNLOAD PROCESS AND REFRAIN FROM ACCESSING OR USING THE SOFTWARE. THIS AGREEMENT REPRESENTS THE ENTIRE AGREEMENT BETWEEN YOU AND UPTIME CONCERNING THE SOFTWARE AND THIS AGREEMENT SUPERSEDES AND REPLACES ANY PRIOR PROPOSAL, REPRESENTATION, COMMUNICATION, ADVERTISEMENT OR UNDERSTANDING YOU MAY HAVE HAD WITH UPTIME RELATING TO THE SOFTWARE.

### 1. License

#### 1.1 Grant of License.

Uptime hereby grants to you and you accept, a limited, non-exclusive license to use the Software in machine-readable, object code form only and the user manuals accompanying the Software (the “Documentation”), only as authorized in this Agreement. For purposes of this Agreement, the “Software” includes any updates, enhancements, modifications, revisions or additions to the Software made by Uptime and made available to end



users through Uptime’s web site. Notwithstanding the foregoing, Uptime shall be under no obligation to provide any updates, enhancements, modifications, revisions or additions to the Software.

## 1.2 Scope of Use

You may use the Software activated by a license key on a single server designated by you as the monitoring station. If you have multiple license keys for the Software each key will be activated on a designated server. For purposes of this Agreement, “use” of the Software means loading the Software into the temporary or permanent memory of a computer. The Software may not be used on or distributed to a greater number of servers than you have license keys. There is no restriction on the number of users who may access the designated servers and use the Software.

## 1.3 Copies and Modifications

You may not reverse engineer, decompile, disassemble or otherwise translate the Software or attempt to derive the source code of the Software or any license keys you have obtained. You may not modify or adapt the Software or any license keys that you have obtained in any way. You may make one (1) copy of the Software, the Documentation and any license keys that you have obtained, solely for backup or archival purposes. Any such copies of the Software, Documentation or license keys shall include any copyright or other proprietary notices that were included on such materials when you first received them. Except as authorized in this Section 1.3, no copies of the Software, Documentation or license keys, or any part thereof, may be made by you or any person under your authority or control.

## 1.4 Assignment of Rights

The license granted under this Agreement is personal to you. You are not permitted to grant access to, distribute, sell, transfer, publish, disclose, display, sublicense, lease, rent or lend your rights in the Software, Documentation or license keys as granted by this Agreement for any purpose or in any manner.

## 1.5 Licenses Required for Third-party Software

The Software enables you to monitor multiple instances of third-party operating systems and application programs. You are responsible for obtaining and complying with any licenses necessary to operate any such third-party software, including Operating Systems and/or application programs.

## 2. Intellectual Property and Confidentiality

### 2.1 Use Reporting, License Violations and Remedies

Uptime reserves the right to gather data on key usage including license key numbers, server IP addresses, domain counts and other information deemed relevant to ensure that its products are being used in accordance with the terms of this Agreement. Uptime expressly prohibits simultaneous, multiple installations of its licensed products and domain count overrides without its prior written approval. Any unauthorized use shall be considered by Uptime to be a violation of this Agreement. Uptime reserves the right to remedy violations immediately upon discovery, by charging the then-current list price of unauthorized keys to the end user or by any other means necessary. You agree not to block, electronically or otherwise, the transmission of data required for compliance with this Agreement. Any blocking of data required for compliance under this Agreement is considered to be violation of this Agreement and will result in immediate termination of this Agreement pursuant to Section 4.

### 2.2 License Automatic Update and Expiration

Your license may include an expiration date that can result in the termination of the license. For permanent license keys, the license updates will be available to you upon payment of the appropriate, then-current Uptime license fees. You must contact Uptime to take the appropriate steps to obtain the permanent key. If your license key is stolen or if you suspect any improper or illegal usage of your license outside of your control you should promptly notify Uptime of such occurrence. A replacement license will be issued to you and the suspect license will be allowed to expire. For your convenience Uptime provides license expiration warnings in the product interface should there be any issues that would cause the license to

expire. It is your responsibility to contact Uptime regarding any potential expiration. Uptime is not liable for any damages or costs incurred in connection with an expiring license.

## 2.3 Proprietary Rights to Software and Trade Marks

You acknowledge that the Software and the Documentation are proprietary to Uptime and the Software and Documentation are protected under Canadian copyright law and international treaties. You further acknowledge and agree that, as between you and Uptime, Uptime owns and shall continue to own all right, title and interest in and to the Software and Documentation including associated intellectual property rights under copyright, trade secret, patent or trade mark laws. This Agreement does not grant you any ownership interest in or to the Software or the Documentation but only a limited right of use that is revocable in accordance with the terms of this Agreement. Any and all trade marks or service marks that Uptime uses in connection with the Software or with services rendered by Uptime are marks owned by Uptime. This Agreement does not grant you any right, license or interest in such marks and you shall not assert any right, license or interest in such marks or any words or designs that are confusingly similar to such marks.

## 2.4 Confidentiality

You shall permit only authorized users who possess rightfully obtained license keys to use the Software or to view the Documentation. Except as expressly authorized by this Agreement you shall not make the Software, Documentation or any license key available to any third party. You will use your best efforts to co-operate with and assist Uptime in identifying and preventing any unauthorized use, copying or disclosure of the Software, Documentation or any part thereof.

## 3. License Fees

The Software will be available to you for use upon your receipt of one or more license keys. Upon acceptance of this Agreement you may obtain one or more temporary license keys and permanent license keys using the procedure set forth on Uptime's web site including, but not limited to, payment of Uptime's license fees. The license fees paid by you are paid in

## *NOTICE TO USER*

consideration of the license granted under this Agreement. Uptime does not refund license fees. By accepting this Agreement you fully understand that once license fee payment is made to Uptime you will have no recourse for receiving a refund of any part of the fees.

### **4. Term and Termination**

This Agreement is effective upon your acceptance of the Agreement or upon your downloading, accessing and using the Software, even if you have not expressly accepted this Agreement. This Agreement shall continue in effect until terminated. Without prejudice to any other rights, this Agreement will terminate automatically if you fail to comply with any of the limitations or other requirements described herein. If you have a temporary key and fail to pay the applicable license fees for continuation of use the key will expire. You may terminate this License Agreement at any time by: (i) providing written notice of your decision to terminate the Agreement to Uptime; and, (ii) either returning the Software, Documentation, all copies thereof and all license keys that you have obtained to Uptime or destroying all such materials and providing written verification of such destruction to Uptime. Uptime reserves the right to physically verify that the Software has been removed. Uptime may terminate this License Agreement if you breach any term of the Agreement by giving you written notice of your breach and Uptime's decision to terminate the Agreement. Upon termination by Uptime you agree to either return the Software, Documentation and all copies thereof and all license keys that you have obtained to Uptime or to destroy all such materials and provide written verification of such destruction to Uptime.

### **5. Remedies and Indemnification**

#### **5.1**

If you learn of any actual or threatened infringement or piracy of the Software or, if any infringement or piracy claim is made against you by a third party in connection with your use of the Software, you shall notify Uptime in writing of the infringement, piracy or claim as soon as is reasonably possible. Uptime shall, in its sole discretion, determine what action, if any, to take with respect to the foregoing and shall assume the

defense or bear the expenses of any such action (except to the extent, if any, to which such dispute or costs arise from your negligence, willful misconduct or modification of the Software). In the event that the use of the Software in accordance with the provisions of this Agreement is declared by a court of competent jurisdiction to infringe the rights of any third party, as your sole remedy, Uptime, at its option may: (i) procure for you the right to use the Software; or, (ii) modify the Software to render it non-infringing.

## 5.2

You will, at your expense, indemnify and hold Uptime and all its officers, directors and employees, harmless from and against any and all claims, actions, liabilities, losses, damages, judgments, grants, costs and expenses, including reasonable lawyer fees (collectively “Claims”) arising out of any use of the Software by you, any party related to you or any party acting upon your authorization in a manner that is not expressly authorized by this Agreement.

## 6. Disclaimer

THE SOFTWARE, DOCUMENTATION AND ANY (IF ANY) SUPPORT SERVICES ARE LICENSED “AS IS” AND UPTIME AND ITS SUPPLIERS DISCLAIM ANY AND ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. WITHOUT LIMITING THE GENERALITY OF THE FOREGOING, UPTIME EXPRESSLY DOES NOT WARRANT THAT THE SOFTWARE WILL MEET YOUR REQUIREMENTS OR THAT OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR FREE. YOU ASSUME RESPONSIBILITY FOR SELECTING THE SOFTWARE TO ACHIEVE YOUR INTENDED RESULTS AND FOR THE RESULTS OBTAINED FROM YOUR USE OF THE SOFTWARE. YOU SHALL BEAR THE ENTIRE RISK AS TO THE QUALITY AND THE PERFORMANCE OF THE SOFTWARE.

## 7. Limitation of Liability

UPTIME’S CUMULATIVE LIABILITY TO YOU OR ANY PARTY RELATED TO YOU FOR ANY LOSS OR DAMAGES RESULTING

## *NOTICE TO USER*

FROM ANY CLAIMS, DEMANDS OR ACTIONS ARISING OUT OF OR RELATING TO THIS AGREEMENT INCLUDING, WITHOUT LIMITATION, UPTIME'S INTELLECTUAL PROPERTY INDEMNIFICATION OBLIGATIONS SHALL BE LIMITED TO THE AMOUNT OF LICENSE FEES PAID TO UPTIME BY YOU UNDER THIS AGREEMENT. BUT, IN NO EVENT SHALL SUCH LIABILITY EXCEED CDN. \$2,000.00 IN THE AGGREGATE FOR ALL OCCURRENCES. THIS LIMITATION APPLIES TO ALL CAUSES OF ACTION OR CLAIMS IN THE AGGREGATE, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, INDEMNITY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS. IN NO EVENT SHALL UPTIME BE LIABLE TO YOU OR ANY PARTY RELATED TO YOU FOR ANY INDIRECT, INCIDENTAL, CONSEQUENTIAL, SPECIAL, EXEMPLARY OR PUNITIVE DAMAGES OR LOST PROFITS EVEN IF UPTIME HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE FOREGOING LIMITATIONS, EXCLUSIONS AND DISCLAIMERS SHALL APPLY TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW EVEN IF ANY REMEDY FAILS ITS ESSENTIAL PURPOSE.

## **8. General Terms**

### **8.1 Governing Law and Choice of Forum**

This Agreement shall be governed by and interpreted in accordance with the laws of the Province of Ontario, Canada, without regard to the conflicts of law rules thereof. Any claim or dispute arising in connection with this Agreement shall be resolved in the federal or provincial courts situated with the City of Toronto, Ontario. To the maximum extent permitted by law, you hereby consent to the jurisdiction and venue of such courts and waive any objections to the jurisdiction or venue of such courts. To the extent any terms and conditions on a purchase order or other ordering document submitted to Uptime by you conflicts with the terms of this Agreement, the terms of this Agreement shall control and notwithstanding any term of your order which states to the contrary.

## 8.2 Severability

If any term or provision of this Agreement is declared void or unenforceable in a particular situation by any judicial or administrative authority this declaration shall not affect the validity or the enforceability of the remaining terms and provisions hereof or the validity or enforceability of the offending term or provision in any other situation.

## 8.3 Survival

Sections 2, 5, 6, 7 and 8 of this Agreement and all subsections thereof shall survive the termination of this Agreement regardless of the cause for termination and shall remain valid and binding indefinitely.

## 8.4 Headings

The Article and Section headings contained in this Agreement are incorporated for reference purposes only and shall not affect the meaning or interpretation of this Agreement.

## 8.5 No Waiver

The failure of either party to enforce any rights granted hereunder or to take action against the other party in the event of any breach hereunder shall not be deemed a waiver by that party as to subsequent enforcement of rights or subsequent actions in the event of future breaches.

## 8.6 Amendment

Uptime reserves the right, in its sole discretion, to amend this Agreement from time to time. If there is a conflict between this Agreement and the most-current version of this Agreement posted at [www.uptimesoftware.com](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 115
- Action Profiles 441
- action profiles
  - creating 443
  - editing 447
  - viewing 447
- Active Directory
  - authentication 401
  - monitor 245
- adding
  - alert settings 152
  - Applications 103
  - Distribution Lists 396
  - groups 108
  - monitor information 145
  - multiple systems 94
  - nested groups 109
  - nested views 112
  - Notification Groups 399
  - Novell NRM system 88
  - OIDs 366
  - service groups 156
  - systems 69
  - user groups 394
  - views 111
  - VMware instance 81
- advanced monitors 141, 373
  - choosing 143
  - custom 142, 376
  - external check 142, 380
  - guidelines 375
  - overview 374
  - types 142
  - with retained data 142, 378
- agent monitors 139
  - File System Capacity 201
  - overview 200, 212
  - Performance Check 204
  - Process Count Check 208
- agentless monitors 140
- agents

- installing 42
  - Linux 44
  - pSeries 45
    - with HMC 45
    - without HMC 47
  - Solaris 43
  - UNIX 44
  - Windows 42
- overview 14
- Alert Profiles 433
  - editing 436
- alert profiles
  - custom formats 437
- alerts
  - acknowledging 115
  - applying to Applications 436
  - creating profiles 434
  - custom formats 437
  - editing profiles 436
  - monitor settings 151
  - overview 430
  - profiles 433
- Application Availability report 508
- application monitors
  - ESX Workload 225
  - Exchange 252
  - IIS 258
  - Live Splunk Listener 290
  - Splunk Query 288
  - up.time Agent 250
  - WebLogic 261
  - WebSphere 269
- Application, Web, monitoring 275
- Applications
  - adding 103
  - applying Action Profiles 441
  - applying Alert Profiles 436
  - deleting 114
  - editing 106
  - maintenance 126
  - offline 126
  - status in Global Scan 126
  - viewing details 105
  - viewing in Global Scan 126
- Archive Policy 655
- auto discovery 76
  - ESX 79
  - pSeries with HMC 80

# Index

using 77

## C

cloning service monitors 154  
Config Panel 9, 637  
    Archive Policy 655  
    Global Scan thresholds 665  
    License Information 673  
    Mail Servers 644  
    Resource Scan thresholds 666  
Config panel  
    Problem Reporting 662  
configuring response time 148  
configuring user roles 386  
configuring users 385  
CPU Run Queue Threshold report 497  
CPU Usage graph 547  
    Linux, UNIX, Novell 548  
    Windows 547  
CPU Utilization Ratio report 474  
CPU Utilization Summary report 471  
creating Alert Profiles 434  
critical threshold 147  
custom alert formats 437

## D

database monitors  
    MySQL (Advanced Metrics) 296  
    MySQL (Basic Checks) 303  
    Oracle (Advanced) 305  
    Oracle (Basic Checks) 308  
    Oracle Tablespace Check 311  
    SQL Server (Advanced Metrics) 318  
    SQL Server (Basic Checks) 314  
    SQL Server Tablespace Check 322  
    Sybase 327  
DataStore  
    archiving 659  
    restoring 659  
dates and times 23  
deleting  
    Applications 114  
    systems 114  
    views 114  
Disk I/O Bandwidth report 493  
Disk Performance Statistics graph 570  
Distribution Lists 396

adding 396  
editing 397  
viewing 397  
DNS monitor 332

## E

editing  
    Action Profiles 447  
    Alert Profiles 436  
    Distribution Lists 397  
    host check 161  
    Notification Groups 400  
    service groups 158  
    system profile 100  
    user groups 394  
    user roles 388  
editing views 113  
email delivery time monitor 282  
enabling Windows Messaging Service 433  
end-user monitoring 275, 282  
Enterprise CPU Utilization report 480  
ESX Workload monitor 225  
Exchange monitor 252  
exiting up.time 51  
External Check monitor 380

## F

File System Capacity 201  
File System Capacity graph 574  
File System Capacity Growth report 483  
File System Capacity monitor 201  
File System Service Time Summary report 501  
filtering 59, 60  
frequency definitions 676  
FTP monitor 335

## G

generating reports 453  
Global Element Settings 646  
Global Scan 117  
    groups 120  
    overview 118  
    Resource Scan 132  
        chart 133  
        gauges 132  
        graph 133

- view all Applications 126
- view all Elements 129
- view all services 131
- Global Scan panel 7
- Graph Editor 538
- Graphing Tool 537
- graphs
  - ActiveX 391, 536
  - appearance 542
  - CPU performance 547
    - generating 550
    - Run Queue Length 549
    - Run Queue Occupancy 549
    - Usage 547
  - disk performance statistics 570
  - displaying process information 580
  - exporting 542
  - file system capacity 574
  - formatting Elements 541
  - Graph Editor 538
  - Graphing Tool 537
  - instance motion 579
  - Java 537
  - LPAR entitlement 566
  - memory usage 554
    - Cache Hit Rate 554
    - Free Swap 555
    - generating 556
    - Paging Statistics 555
    - Used 554
  - Multi-CPU Usage 551
  - network 567
    - errors 567
    - generating 568
    - I/O 567
    - NetFlow 568
  - Novell NRM 577
  - overview 13, 536, 544, 584
  - process 557
    - creation rate 558
    - generating 558
    - number of processes 557
    - running, blocked, waiting 557
  - Quick Snapshot 545, 585
  - setting date and time ranges 23
  - TCP retransmits 559
  - top 10 disks 572

- trend lines 540
- user activity 560
- viewing quick snapshot 546, 592
- viewing system status 545, 585
- VXVM stats 575
- workload 561
- workload top 10 564
- groups
  - adding 108
  - adding nested 109
  - Global Scan 120

## H

- host check 4, 57
  - changing 161
  - ping 161
  - up.time agent 161
- HTTP (Web services) monitor 337

## I

- icons 11
  - Clone 11, 154
  - critical 126
  - Delete 11, 395
  - Edit 11, 161, 436, 447
  - View 11
- IIS monitor 258
- IMAP (Email Retrieval) monitor 341
- installation
  - agents 42
    - Linux 44
    - pSeries 45
      - with HMC 45
      - without HMC 47
    - Solaris 43
    - UNIX 44
    - Windows 42
  - guidelines 28
  - Monitoring Station 31
    - UNIX/Linux 34
    - VMware 37
    - Windows 32
  - post-installation tasks 39
  - requirements 29
    - browsers 30
    - hardware 30

# Index

- Monitoring Station 29
    - up.time agents 30
  - upgrading 41
  - Instance Motion graphs 579
  - interface
    - overview 6
    - up.time tool bar 6
      - Config 9
      - Global Scan 7
      - My Infrastructure 7
      - My Portal 7
      - Reports 9
      - Services 8
      - Users 8
- L**
- LDAP
    - authentication 401
    - monitor 343
  - license information 673
  - Live Splunk Listener monitor 290
  - LPAR
    - adding 83
    - entitlement graphs 566
    - workload graphs 565
  - LPAR workload graphs 565
- M**
- mail servers 644
  - maintenance on demand 169
  - monitor template 144
    - alert settings 151
    - configuring response time 148
    - configuring settings 145
    - identification 144
    - Monitoring Period 153
    - response time 148
    - thresholds 147
    - timing setting options 150
    - timing settings 149
  - Monitoring Periods 449
  - monitors 137
    - adding alert settings 152
    - adding information 145
    - adding timing settings 151
    - advanced 141, 373
    - external check 380
    - with retained data 378
  - agent
    - File System Capacity 201
    - overview 200, 212
    - Performance Check 204
    - Process Count Check 208
  - agentless
    - NIS/YP 349
  - alert settings 151
  - application
    - Exchange 252
    - IIS 258
    - Live Splunk Listener 290
    - Splunk query 288
    - up.time Agent 250
    - WebLogic 261
    - WebSphere 269
  - cloning 154
  - comparisons 146
  - database
    - MySQL (Advanced Metrics) 296
    - MySQL (Basic Checks) 303
    - Oracle (Advanced) 305
    - Oracle (Basic Checks) 308
    - Oracle Tablespace 311
    - SQL Server (Advanced Metrics) 318
    - SQL server (Basic Checks) 314
    - SQL Server Tablespace 322
    - Sybase 327
  - editing performance 162
  - getting help 153
  - identification 144
  - monitor settings 145
  - Monitoring Period 153
  - Monitoring Periods 449
  - network
    - DNS 332
    - FTP 335
    - HTTP 337
    - IMAP 341
    - LDAP 343
    - NFS 347
    - NIS/YP 349
    - NNTP 351

- Ping 355
  - POP 357
  - SMTP 361
  - SNMP 363
  - SSH 359
  - TCP 370
  - overview 18, 138
  - template 144
  - testing 155
  - timing setting options 150
  - timing settings 149
  - types 139
  - Windows
    - Active Directory 245
    - Event Log Scanner 236
    - Service Check 240
    - SMB 243
  - Multi-System CPU report 470
  - My Infrastructure 67
    - acknowledge alerts 115
    - adding Applications 103
    - adding groups 108
    - adding nested groups 109
    - adding nested views 112
    - adding systems 69
    - Application details 105
    - editing Applications 106
    - editing system profile 100
    - overview 68
    - views 111
  - My Infrastructure panel 7
  - My Portal panel 7, 63, 64
  - MySQL (Advanced Metrics) monitor 296
  - MySQL (Basic Checks) monitor 303
- ## N
- Net-SNMP 363
  - Network Bandwidth Report 490
  - Network graphs 567
  - network monitors
    - DNS 332
    - FTP 335
    - HTTP 337
    - IMAP 341
    - LDAP 343
    - NFS 347
    - NIS/YP 349
    - NNTP 351
    - Ping 355
    - POP 357
    - SMTP 361
    - SNMP 363
    - SSH 359
    - TCP 370
  - NFS monitor 347
  - NIS/YP monitor 349
  - NNTP (Network News) monitor 351
    - command implementation 351
    - response category 352
    - response codes 352
  - Notification Groups 399
  - notification groups
    - adding 399
    - editing 400
    - overview 399
    - viewing 400
  - Novell NRM 88
    - Abended Thread Count 91
    - abended thread count 91
    - adding 88
    - Allocated Server Processes 90
    - allocated server processes 90
    - Available Disk Space 93
    - available disk space 93
    - Available Event Control Blocks (ECBs) 93
    - Available Memory 92
    - available memory 92
    - Available Server Processes 91
    - available server processes 91
    - Connection Usage 92
    - connection usage 92
    - CPU Utilization 91
    - CPU utilization 91
    - Disk Throughput 93
    - disk throughput 93
    - DS Thread Usage 92
    - DS thread usage 92
    - ECBs 93
    - LAN Traffic 93
    - LAN traffic 93
    - Packet Receive Buffers 92
    - packet received buffers 92
    - Statistics Available 89
    - statistics captured 89

# Index

Work To Do Response Time 90  
work to do response time 90

## O

OIDs

deleting 366  
manually adding 366

Oracle

(Advanced Metrics) monitor 305  
(Basic Checks) monitor 308  
Tablespace Check monitor 311  
using as the up.time database 39

Orchestrator 441, 443

## P

Performance Check monitor 204

Ping monitor 355

POP (Email Retrieval) monitor 357

Problem Reporting 662

Process Count Check 208

Configuring 208  
Settings 208

Process Count Check monitor 208

pSeries

adding LPARs 83

## R

Report Log 462

completed reports 463  
deleting 464  
pending reports 462  
running reports 462  
viewing 463

reports 465

Application Availability 508  
background 453  
CPU Run Queue Threshold 497  
CPU Utilization Ratio 474  
CPU Utilization Summary 471  
Disk I/O Bandwidth 493  
dynamic 453  
Enterprise CPU Utilization 480  
File System Capacity Growth 483  
File System Service Time Summary 501  
generating 453  
generation options 454

email 454  
to screen 454  
XML 454

incidents 509

mean time between failure 509

mean time to repair 509

Multi-System CPU 470

Network Bandwidth 490

overview 13

Report Log 462

Resource Usage 466

saving 456

saving to file system 456

scheduling 459

searching saved 458

Server Virtualization 484

Service Monitor Availability 512

Service Monitor Metrics 477

Service Monitor Outages 513

setting date and time ranges 23

Solaris Mutex Exception 488

viewing saved 457

VMware Workload 528

Wait I/O 475

WebSphere 515

Reports panel 9

Resource Scan 132

Resource Usage report 466

Response Time 148

## S

scheduled maintenance 167

assigning to a host 168

assigning to a service 169

profiles 167

viewing profiles 168

Scrutinizer 74, 135, 568, 652

search box 59

searching 59

Server Virtualization report 484

using 487

service groups 156

creating 156

editing 158

overview 21

service level agreements 171, 409

adding and editing 423

- creating 418
- end-user performance 280
- objectives 411, 417, 425
- reports for 505
- status 415
- viewing 121, 412
- service monitor
  - File System Capacity 201
  - Process Count Check 208
  - SQL Server (Advanced) 318
  - up.time Agent 250
- Service Monitor Availability report 512
- Service Monitor Metrics report 477
- Service Monitor Outages report 513
- service monitors 137
  - adding alert settings 152
  - adding information 145
  - adding timing settings 151
  - advanced 141, 373
    - external check 380
    - with retained data 378
  - agent
    - File System Capacity 201
    - overview 200, 212
    - Performance Check 204
    - Process Count Check 208
- alert settings 151
- application
  - ESX Workload 225
  - Exchange 252
  - IIS 258
  - Live Splunk Listener 290
  - Splunk query 288
  - up.time Agent 250
  - WebLogic 261
  - WebSphere 269
- cloning 154
- comparisons 146
- database
  - MySQL (Advanced Metrics) 296
  - MySQL (Basic Checks) 303
  - Oracle 308
  - Oracle (Advanced) 305
  - Oracle tablespace 311
  - SQL Server (Advanced Metrics) 318
  - SQL server (Basic Checks) 314
  - SQL Server Tablespace 322
    - Sybase 327
- getting help 153
- identification 144
- monitor settings 145
- Monitoring Period 153
- network
  - DNS 332
  - FTP 335
  - HTTP 337
  - IMAP 341
  - LDAP 343
  - NFS 347
  - NIS/YP 349
  - NNTP 351
  - Ping 355
  - POP 357
  - SMTP 361
  - SNMP 363
  - SSH 359
  - TCP 370
- overview 18, 138
- template 144
- testing 155
- timing settings 149
- types 139
- Windows
  - Active Directory 245
  - Event Log Scanner 236
  - Service Check 240
  - SMB 243
- services
  - filtering 60
  - service groups 21
  - starting 641
  - stopping 640
  - viewing 131
- Services panel 8
- SMB Check monitor 243
- SMTP monitor 361
- SNMP 363
  - MIB Browser 364
  - supported versions 364
- SNMP monitor 363
  - adding OIDs 366
  - deleting OIDs 366

# Index

- loading MIBs 365
- manually adding OIDs 366
- MIB browser 364
- Solaris Mutex Exception report 488
- Splunk
  - Action Profile 446
  - Live Splunk Listener monitor 290
  - query monitor 288
- SQL Server (Advanced Metrics) monitor 318
- SQL Server (Advanced) 318
- SQL Server (Advanced) monitor
  - multiple 318
- SQL Server (Basic Checks) monitor 314
- SQL Server Tablespace Check monitor 322
- SSH monitor 359
- starting up.time 51
- supported Web browsers 30
- Sybase monitor 327
- System List (Syslist) 10
- systems
  - adding 71
  - deleting 114

## T

- temporary maintenance 169
- testing monitors 155
- time period definitions 676, 677
- Top 10 Disks graph 572
- topological dependencies 164
  - adding 165
  - viewing 166

## U

- UNIX vs. Windows 544
- up.time
  - administrator account 50
  - exiting 51
  - installing 27
  - interface 6
  - monitoring concepts 4
  - overview 2
  - service information 54
  - services
    - starting 641
    - stopping 640
    - stopping and starting 640

- starting 51
- starting and exiting 50
- system information 52
- tool bar 6
  - Config 9
  - Global Scan 7
  - My Infrastructure 7
  - My Portal 7
  - Reports 9
  - Services 8
  - Users 8
- viewing information 52
- up.time Agent Monitor 646
- up.time Agent monitor 250
- upgrading to up.time 4 41
- uptime.conf
  - database 642
  - NetFlow 652
  - remote reporting 651
  - RSS feed 647
  - Splunk 653
  - UI instance 652
  - Web monitor proxy 650
- user groups
  - adding 394
  - deleting 395
  - editing 394
  - overview 393
  - viewing 394
- user roles
  - adding 386
  - editing 388
  - overview 386
  - viewing 387
- users
  - adding 389
  - configuring 385
  - Distribution Lists 396
  - editing 392
  - Notification Groups 399
  - overview 389
  - roles 386
  - viewing 392
- Users panel 8

## V

- viewing



- Action Profiles 447
  - all Elements 129
  - detailed process information 580
- Distribution Lists 397
- Notification Groups 400
- Quick Snapshot 546, 592
- report logs 463
- Resource Scan 132
- scheduled Maintenance Profiles 168
- service information 54
- services 131
- system information 52
- system status 545, 585
- user groups 394
- user roles 387
- users 392
- views
  - adding 111
  - adding nested 112
  - deleting 114
- virtual infrastructure
  - density 531
- VM
  - density 531
- VMware 81, 484
  - Instance Motion graph 579
- VMware instance 81
- VMware vCenter Orchestrator 441, 443, 649
- VMware Workload report 528
- VXVM Stats graph 575

## W

- Wait I/O report 475
- warning threshold 147
- Web application transaction monitoring 275
- WebLogic monitor
  - configuring 265, 267
  - metrics 262
- Weblogic monitor 261
- WebLogic report
  - using 520
- WebSphere monitor 269
  - configuring 273
  - counters 269
- WebSphere report 515
  - using 517
- Windows Event Log Scanner monitor 236

- Windows File Shares monitor 243
- Windows monitors
  - Active Directory 245
  - Event Log Scanner 236
  - Service Check 240
  - SMB 243
- Windows Service Check monitor 240
- WMI 83, 646
  - agentless system, adding 84
- workflows 441, 443
- Workload graphs 561
- Workload Top 10 graphs 564