

**ProTune™**  
*Console User's Guide*  
Version 1.5



MERCURY INTERACTIVE

ProTune Console User's Guide, Version 1.5

This manual, and the accompanying software and other documentation, is protected by U.S. and international copyright laws, and may be used only in accordance with the accompanying license agreement. Features of the software, and of other products and services of Mercury Interactive Corporation, may be covered by one or more of the following patents: U.S. Patent Nos. 5,701,139; 5,657,438; 5,511,185; 5,870,559; 5,958,008; 5,974,572; 6,138,157; 6,144,962; 6,205,122; 6,237,006; 6,341,310; and 6,360,332. Other patents are pending in the U.S. and other countries.

WinRunner, XRunner, and LoadRunner are registered trademarks of Mercury Interactive Corporation. WinRunner, XRunner, and LoadRunner are registered trademarks of Mercury Interactive Corporation in the United States and/or other countries.

Astra QuickTest, Astra LoadTest, Astra FastTrack, RapidTest, QuickTest, Visual Testing, Action Tracker, Link Doctor, Change Viewer, Dynamic Scan, Fast Scan, Visual Web Display, ActiveTest, ActiveTest SecureCheck, ActiveWatch, POPs on Demand, Topaz, Topaz ActiveAgent, Topaz Observer, Topaz Prism, Topaz Delta, Topaz Rent-a-POP, Topaz Open DataSource, Topaz AIMS, Topaz Console, Topaz Diagnostics, Topaz WeatherMap, Twinlook, TurboLoad, LoadRunner TestCenter, SiteReliance and Global SiteReliance are trademarks of Mercury Interactive Corporation in the United States and/or other countries.

All other company, brand and product names are registered trademarks or trademarks of their respective holders. Mercury Interactive Corporation disclaims any responsibility for specifying which marks are owned by which companies or which organizations.

Mercury Interactive Corporation  
1325 Borregas Avenue  
Sunnyvale, CA 94089  
Tel. (408)822-5200 (800) TEST-911  
Fax. (408)822-5300

© 2003 Mercury Interactive Corporation, All rights reserved

If you have any comments or suggestions regarding this document, please send them via e-mail to [documentation@merc-int.com](mailto:documentation@merc-int.com).

---

# Table of Contents

|   |      |
|---|------|
| Welcome to ProTune.....                   | xiii |
| Online Resources .....                    | xiii |
| ProTune Documentation Set .....           | xiv  |
| Using the ProTune Documentation Set ..... | xv   |
| Documentation Updates .....               | xvi  |
| Typographical Conventions.....            | xvii |

## **PART I: UNDERSTANDING PROTUNE**

|                                      |   |
|--------------------------------------|---|
| Chapter 1: Introducing ProTune ..... | 3 |
| ProTune's Features .....             | 4 |
| Tuning Workflow .....                | 5 |

## **PART II: MAPPING A BUSINESS PROCESS**

|  |    |
|--|----|
| Chapter 2: Creating a Topology .....                 | 9  |
| About Creating a Topology .....                      | 10 |
| Building a Topology Diagram .....                    | 11 |
| Setting the Component Properties.....                | 16 |
| Selecting Monitors.....                              | 18 |
| Automatically Assigning Monitors and Alerts .....    | 23 |
| Accessing the SiteScope Administration Console ..... | 26 |

## **PART III: DESIGNING A TUNING SESSION**

|   |     |
|---|-----|
| <b>Chapter 3: Adding Session Steps</b> .....      | 29  |
| About Adding Session Steps .....                  | 29  |
| Getting Started with Designing a Session .....    | 30  |
| Understanding Session Steps.....                  | 31  |
| Adding Session Steps .....                        | 40  |
| Managing Steps .....                              | 44  |
| Specifying Step Execution Order.....              | 45  |
| Managing Scripts .....                            | 46  |
| Setting an Initial Load (Manual Profiles) .....   | 50  |
| Configuring Script Details .....                  | 51  |
| Using Relative Paths for Scripts.....             | 54  |
| <b>Chapter 4: Managing Load Generators</b> .....  | 57  |
| About Managing Load Generators .....              | 57  |
| Configuring Load Generators .....                 | 57  |
| Configuring Load Generator Settings .....         | 61  |
| <b>Chapter 5: Configuring Session Steps</b> ..... | 69  |
| About Configuring a Session.....                  | 69  |
| Configuring Session Run-Time Settings.....        | 70  |
| Setting Timeout Intervals .....                   | 71  |
| Setting the Run-Time File Location .....          | 74  |
| Specifying Path Translation .....                 | 76  |
| <b>Chapter 6: Defining Alerts</b> .....           | 77  |
| About Defining Alert Schemes .....                | 77  |
| Types of Alerts .....                             | 78  |
| Specifying Alert Conditions .....                 | 80  |
| Specifying Alert Actions .....                    | 83  |
| Viewing Alert Descriptions.....                   | 85  |
| Creating, Configuring, and Deleting Alerts.....   | 85  |
| Enabling and Disabling the Alert Mechanism .....  | 91  |
| Viewing Alerts in the Output Window.....          | 92  |
| <b>Chapter 7: Scheduling Session Steps</b> .....  | 95  |
| About Scheduling Session Steps .....              | 95  |
| Specifying Execution Time.....                    | 96  |
| Creating and Selecting a Profile .....            | 98  |
| Creating a Manual Profile .....                   | 101 |
| Creating a Goal-Oriented Profile .....            | 107 |

|   |     |
|---|-----|
| <b>Chapter 8: Preparing to Run a Session Step</b> ..... | 113 |
| About Preparing to Run a Session Step .....             | 113 |
| Specifying a Results Location .....                     | 114 |
| Results Directory File Structure .....                  | 116 |
| Collating Results .....                                 | 117 |

## **PART IV: EXECUTING A TUNING SESSION**

|  |     |
|--|-----|
| <b>Chapter 9: Running a Session</b> .....                            | 121 |
| About Running a Session Step .....                                   | 121 |
| Running an Entire Session .....                                      | 123 |
| Controlling a Specific Number of Vusers.....                         | 124 |
| Continuing With Subsequent Steps.....                                | 125 |
| Adding Vusers to a Running Session.....                              | 126 |
| Viewing and Controlling Vusers.....                                  | 129 |
| Invoking the System Topology Window .....                            | 133 |
| <b>Chapter 10: Viewing Vusers During Execution</b> .....             | 135 |
| About Viewing Vusers During Execution.....                           | 135 |
| Monitoring Vuser Status .....  | 136 |
| Viewing the Output Window .....                                      | 140 |
| Viewing the Script Log .....   | 143 |
| Viewing the Agent Summary .....                                      | 145 |
| <b>Chapter 11: Viewing the Session Summary</b> .....                 | 147 |
| About Viewing the Session Summary .....                              | 147 |
| The Session Summary Window.....                                      | 148 |
| Viewing Step Run Information .....                                   | 148 |
| Viewing Tuning Information .....                                     | 150 |
| Adding Session Notes .....   | 151 |
| Deleting Session Summary Entries.....                                | 151 |
| <b>Chapter 12: Generating a Session Report</b> .....                 | 153 |
| About Generating a Session Report.....                               | 153 |
| Session Summary Section .....  | 153 |
| Step Information .....   | 154 |
| Generating the Report.....   | 154 |
| <b>Chapter 13: Working with Firewalls</b> .....                      | 157 |
| About Using Firewalls in ProTune.....                                | 157 |
| Overview of Running or Monitoring over the Firewall .....            | 159 |
| Configuring the ProTune Agents in LAN1.....                          | 160 |
| Configuring the Firewall to Allow Agent Access.....                  | 168 |
| Installing and Configuring the MI Listener in LAN2 .....             | 169 |
| Configuring Console to Run or Monitor Vusers over the Firewall...170 |     |

## **PART V: MONITORING A SESSION**

|  |     |
|--|-----|
| <b>Chapter 14: Online Monitoring</b> .....                   | 175 |
| About Online Monitoring .....                                | 176 |
| Choosing Monitors and Measurements.....                      | 178 |
| Viewing the Monitors .....                                   | 180 |
| Opening Online Monitor Graphs .....                          | 182 |
| Customizing the Graph Display View .....                     | 184 |
| Configuring Online Monitors .....                            | 184 |
| Setting Monitor Options .....                                | 185 |
| Configuring Online Graphs .....                              | 187 |
| Merging Graphs .....   | 192 |
| Understanding Online Monitor Graphs .....                    | 193 |
| Configuring Online Measurements .....                        | 194 |
| Exporting Online Monitor Graphs .....                        | 198 |
| Viewing Data Offline .....                                   | 198 |
| <b>Chapter 15: Monitoring over a Firewall</b> .....          | 199 |
| About Monitoring over the Firewall .....                     | 200 |
| Installing Monitors over Firewall .....                      | 200 |
| Installing MI Listener .....                                 | 205 |
| Preparing for Data Collection .....                          | 205 |
| Configuring Server Monitor Properties .....                  | 205 |
| Adding and Removing Measurements .....                       | 208 |
| Configuring Measurement Frequency .....                      | 209 |
| Configuring the Network Delay Monitor over a Firewall.....   | 209 |
| <b>Chapter 16: Run-Time and Transaction Monitoring</b> ..... | 211 |
| About Run-Time and Transaction Graphs .....                  | 211 |
| Run-Time Graphs .....  | 212 |
| User-Defined Data Points Graph .....                         | 213 |
| Transaction Monitor Graphs .....                             | 214 |
| Enabling the Transaction Monitor .....                       | 215 |
| Adding Transactions to a Script .....                        | 216 |
| Enabling Web Page Breakdown .....                            | 217 |

|   |     |
|---|-----|
| <b>Chapter 17: Web Resource Monitoring</b> .....                | 219 |
| About Web Resource Monitoring.....                              | 219 |
| Hits per Second Graph .....                                     | 220 |
| Throughput Graph .....  | 220 |
| HTTP Responses per Second Graph .....                           | 221 |
| Pages Downloaded per Second Graph .....                         | 223 |
| Retries per Second Graph .....                                  | 224 |
| Connections Graph .....   | 225 |
| Connections per Second Graph .....                              | 225 |
| SSL Connections per Second Graph .....                          | 225 |
| <b>Chapter 18: System Resource Monitoring</b> .....             | 227 |
| About System Resource Monitoring.....                           | 227 |
| Configuring the Windows Resources Monitor .....                 | 229 |
| Configuring the UNIX Resources Monitor .....                    | 234 |
| Configuring an rstatd Daemon on UNIX .....                      | 236 |
| Configuring the SNMP Resources Monitor .....                    | 238 |
| Configuring the TUXEDO Monitor .....                            | 241 |
| Configuring the Antara FlameThrower Monitor .....               | 247 |
| Configuring the SiteScope Monitor .....                         | 259 |
| <b>Chapter 19: Network Monitoring</b> .....                     | 263 |
| About Network Monitoring.....                                   | 263 |
| Network Monitoring from a UNIX Source Machine .....             | 265 |
| Configuring the Network Monitor .....                           | 268 |
| Viewing the Network Delay Time Graph .....                      | 272 |
| <b>Chapter 20: Firewall Server Performance Monitoring</b> ..... | 275 |
| About the Firewall Server Monitor.....                          | 275 |
| Configuring the CheckPoint FireWall-1 Server Monitor .....      | 275 |
| <b>Chapter 21: Web Server Resource Monitoring</b> .....         | 279 |
| About Web Server Resource Monitors.....                         | 279 |
| Configuring the Apache Monitor .....                            | 280 |
| Configuring the Microsoft IIS Monitor .....                     | 283 |
| Configuring the iPlanet/Netscape Monitor .....                  | 285 |
| Configuring the iPlanet (SNMP) Monitor .....                    | 289 |
| Monitoring Using a Proxy Server .....                           | 297 |

|   |     |
|---|-----|
| <b>Chapter 22: Web Application Server Resource Monitoring</b> ..... | 299 |
| About Web Application Server Resource Monitors .....                | 300 |
| Configuring the Ariba Monitor .....                                 | 300 |
| Configuring the ATG Dynamo Monitor .....                            | 304 |
| Configuring the BroadVision Monitor .....                           | 310 |
| Configuring the ColdFusion Monitor .....                            | 320 |
| Configuring the Fujitsu INTERSTAGE Monitor .....                    | 323 |
| Configuring the iPlanet (NAS) Monitor .....                         | 325 |
| Configuring the Microsoft Active Server Pages Monitor .....         | 342 |
| Configuring the Oracle9iAS HTTP Monitor .....                       | 344 |
| Configuring the SilverStream Monitor .....                          | 349 |
| Configuring the WebLogic (SNMP) Monitor .....                       | 353 |
| Configuring the WebLogic (JMX) Monitor .....                        | 357 |
| Configuring the WebSphere Monitor .....                             | 363 |
| Configuring the WebSphere (EPM) Monitor .....                       | 377 |
| <b>Chapter 23: Database Resource Monitoring</b> .....               | 387 |
| About Database Resource Monitoring.....                             | 387 |
| Configuring the DB2 Monitor .....                                   | 388 |
| Configuring the Oracle Monitor .....                                | 405 |
| Configuring the SQL Server Monitor .....                            | 412 |
| Configuring the Sybase Monitor .....                                | 415 |
| <b>Chapter 24: Streaming Media Monitoring</b> .....                 | 423 |
| About Streaming Media Monitoring.....                               | 423 |
| Configuring the Windows Media Server Monitor .....                  | 424 |
| Configuring the RealPlayer Server Monitor .....                     | 426 |
| Viewing the RealPlayer Client Online Graph.....                     | 428 |
| Viewing the Media Player Client Online Graph.....                   | 429 |
| <b>Chapter 25: ERP/CRM Server Resource Monitoring</b> .....         | 431 |
| About ERP/CRM Server Resource Monitoring .....                      | 431 |
| Setting up the Monitoring Environment.....                          | 432 |
| Setting Up the SAP Monitor .....                                    | 433 |
| Configuring the SAP Monitor .....                                   | 433 |
| Configuring the SAP Portal Monitor.....                             | 438 |
| Configuring the Siebel Monitor .....                                | 441 |
| Configuring the Siebel Server Manager Monitor .....                 | 445 |
| <b>Chapter 26: Java Performance Monitoring</b> .....                | 451 |
| About Java Performance Monitoring .....                             | 451 |
| EJB Performance Monitoring .....                                    | 452 |
| JProbe Performance Monitoring .....                                 | 468 |
| Sitraka JMonitor Performance Monitoring.....                        | 470 |



|   |     |
|---|-----|
| <b>Chapter 27: J2EE Performance Monitoring</b> .....                | 491 |
| About J2EE Performance Monitoring.....                              | 492 |
| Installing the J2EE Monitor on the Application Server .....         | 493 |
| Initial J2EE Monitor Configuration Settings.....                    | 495 |
| Activating the J2EE Monitor on the Client Machine .....             | 497 |
| Examples of Modifying Application Server Configurations .....       | 500 |
| Troubleshooting the J2EE Monitor .....                              | 507 |
| <b>Chapter 28: Application Deployment Solution Monitoring</b> ..... | 509 |
| About Application Deployment Solution Monitoring .....              | 509 |
| Monitoring Citrix MetaFrame Servers.....                            | 509 |
| Configuring the Citrix MetaFrame Server Monitor .....               | 510 |
| <b>Chapter 29: Middleware Performance Monitoring</b> .....          | 519 |
| About Middleware Performance Monitoring.....                        | 519 |
| Configuring the IBM WebSphere MQ Monitor .....                      | 520 |
| Configuring the TUXEDO Monitor .....                                | 530 |
| <b>Chapter 30: Application Traffic Management</b> .....             | 535 |
| About Application Traffic Management Monitoring.....                | 535 |
| Configuring the F5 BIG-IP Monitor .....                             | 536 |
| <b>Chapter 31: Troubleshooting Online Monitors</b> .....            | 541 |
| Troubleshooting Server Resource Monitors .....                      | 541 |
| Troubleshooting the Network Delay Monitor .....                     | 544 |
| Network Considerations.....   | 546 |

## **PART VI: TUNING YOUR SYSTEM**

|  |     |
|--|-----|
| <b>Chapter 32: Tuning Your System from the Console</b> ..... | 551 |
| About Tuning Your System from the Console.....               | 552 |
| Supported Operating Systems .....                            | 552 |
| Applications That ProTune Can Tune.....                      | 553 |
| Tuning Flow.....   | 554 |
| Host Requirements .....                                      | 554 |
| Configuring Host Connection Parameters.....                  | 557 |
| Connecting to the Host Computer .....                        | 562 |
| Viewing the Host Information .....                           | 567 |
| Viewing Windows Services.....                                | 570 |
| Resynchronizing the Information Tab.....                     | 571 |
| Using Expert Mode .....                                      | 572 |
| Changing Tuning Parameter Values .....                       | 572 |
| Updating the Host or Service with Changes .....              | 574 |
| Configuring Special Tuner Agent Settings .....               | 575 |

|   |     |
|---|-----|
| <b>Chapter 33: Exporting and Importing Configuration Settings</b> ..... | 583 |
| About Exporting and Importing Configuration Settings.....               | 583 |
| Exporting a Host or Service's Configuration Settings.....               | 584 |
| Importing Configuration Settings for a Host or Service .....            | 585 |
| Saving and Loading Profiles .....                                       | 587 |
| Creating a New Profile .....  | 589 |
| <b>Chapter 34: Configuring Tuning Agents</b> .....                      | 591 |
| About Configuring Tuning Agents .....                                   | 591 |
| Changing Tuning Agent Passwords .....                                   | 592 |
| Changing the Tuning Agent's Port .....                                  | 593 |
| Using the Performance Tuner Registry .....                              | 596 |
| Automatically Starting the Tuning Agent when Booting.....               | 600 |
| Starting and Stopping the Apache and IBM HTTP Servers .....             | 601 |
| <b>Chapter 35: Tune Tab Functions</b> .....                             | 607 |
| About Tune Tab Functions.....   | 607 |
| Start a Service.....  | 608 |
| Stop a Service.....   | 608 |
| Reboot Host Machines .....  | 609 |
| Reconnect the Console to a Server.....                                  | 610 |
| Stop the Tuning Agent .....   | 610 |
| Print Host Configurations.....  | 611 |
| Reload Host Configuration .....   | 611 |
| Remove Host from Server Configurations Tree .....                       | 611 |
| <b>Chapter 36: Tuning UNIX Hosts</b> .....                              | 613 |
| Using Telnet .....  | 613 |
| Redirecting Script Output.....  | 613 |
| Solaris Requirements .....  | 614 |
| IBM AIX Requirements.....   | 615 |
| HP-UX Requirements .....  | 616 |
| Linux Requirements .....  | 617 |

**PART VII: APPENDIXES**

|   |     |
|---|-----|
| <b>Appendix A: Troubleshooting the Console</b> .....          | 621 |
| About Troubleshooting .....                                   | 621 |
| ProTune Communications .....                                  | 622 |
| Failure to Communicate with a Load Generator .....            | 623 |
| Failure to Connect to the AUT Database .....                  | 629 |
| Failure to Access Files .....                                 | 629 |
| Failed Vusers or Transactions .....                           | 631 |
| Increasing the Number of Vusers on a Windows Machine .....    | 634 |
| Troubleshooting Firewalls .....                               | 635 |
| Troubleshooting Remote Tuning .....                           | 643 |
| <b>Appendix B: Working in Expert Mode</b> .....               | 645 |
| Entering Expert Mode .....                                    | 645 |
| Options - Agent Settings .....                                | 646 |
| Options - General Settings .....                              | 647 |
| Options - Debug Information Settings .....                    | 648 |
| Options - Output Settings .....                               | 650 |
| Options - Monitor Settings .....                              | 651 |
| Load Generator Information - UNIX Environment Settings .....  | 652 |
| Load Generator Information - Connection Log Settings .....    | 653 |
| <b>Appendix C: Performing Path Translation</b> .....          | 655 |
| Understanding Path Translation .....                          | 655 |
| Adding Entries to the Path Translation Table.....             | 657 |
| Editing the Path Translation Table .....                      | 659 |
| Path Translation Examples.....                                | 660 |
| <b>Appendix D: Working with Server Monitor Counters</b> ..... | 661 |
| Changing a Monitor's Default Counters .....                   | 661 |
| Useful Counters for Stress Testing .....                      | 662 |
| <b>Appendix E: Configuring Multiple IP Addresses</b> .....    | 665 |
| About Multiple IP Addresses .....                             | 666 |
| Adding IP Addresses to a Load Generator .....                 | 667 |
| Using the IP Wizard.....                                      | 667 |
| Configuring Multiple IP Addresses on UNIX.....                | 672 |
| Updating the Routing Table.....                               | 673 |
| Enabling Multiple IP Addressing from the Console .....        | 674 |
| <b>Appendix F: Working with Digital Certificates</b> .....    | 675 |
| Using Digital Certificates with Firewalls .....               | 675 |
| Creating and Using Digital Certificates .....                 | 676 |
| <b>Index</b> .....  | 683 |



---

# Welcome to ProTune

Welcome to ProTune, the proactive solution for optimizing production systems. ProTune's system-wide approach to optimization is a product of Mercury Interactive's expert knowledge and tuning methodologies.

ProTune's chief purpose is to enable the user to explore the network, detect bottlenecks, and assist during the tuning phase to enhance performance. ProTune combines the network topology, predefined tuning sessions and goals into a set of tests that pinpoint problematic components in the client network. In addition, ProTune helps the more advanced user to explore and tune the business processes by providing a simple and organized methodology.

## Online Resources



ProTune includes the following online tools:

**Read Me First** provides last-minute news and information about ProTune.

**Books Online** displays the complete documentation set in PDF format. Online books can be read and printed using Adobe Acrobat Reader, which is included in the installation package. Check Mercury Interactive's Customer Support Web site for updates to ProTune online books.

**ProTune Function Reference** gives you online access to all of ProTune's functions that you can use when creating Vuser scripts, including examples of how to use the functions. Check Mercury Interactive's Customer Support Web site for updates to the online *ProTune Function Reference*.

**ProTune Context Sensitive Help** provides immediate answers to questions that arise as you work with ProTune. It describes dialog boxes, and shows you how to perform ProTune tasks. To activate this help, click in a window and press F1. Check Mercury Interactive's Customer Support Web site for updates to ProTune help files.

**Technical Support Online** uses your default Web browser to open Mercury Interactive's Customer Support Web site. This site enables you to browse the knowledge base and add your own articles, post to and search user discussion forums, submit support requests, download patches and updated documentation, and more. The URL for this Web site is <http://support.mercuryinteractive.com>.

**Support Information** presents the locations of Mercury Interactive's Customer Support Web site and home page, the e-mail address for sending information requests, and a list of Mercury Interactive's offices around the world.

**Mercury Interactive on the Web** uses your default Web browser to open Mercury Interactive's home page (<http://www.mercuryinteractive.com>). This site enables you to browse the knowledge base and add your own articles, post to and search user discussion forums, submit support requests, download patches and updated documentation, and more.

## ProTune Documentation Set

ProTune is supplied with a set of documentation that describes how to:

- install ProTune
- create scripts
- use the ProTune Console
- use the ProTune Analysis

## Using the ProTune Documentation Set

The ProTune documentation set consists of one installation guide, a Console user's guide, an Analysis user's guide, guides for creating Virtual User scripts, and a tutorial.

### Installation Guide

For instructions on installing ProTune, refer to the *ProTune Installation Guide*. The installation guide explains how to install:

- ▶ the ProTune Console—on a Windows-based machine
- ▶ Virtual User components—for Windows and UNIX platforms

### Console User's Guide

The ProTune documentation pack includes one Console user's guide:

The *ProTune Console User's Guide* describes how to create and run ProTune sessions using the ProTune Console in a Windows environment. The Console user's guide presents an overview of the ProTune testing process.

### Analysis User's Guide

The ProTune documentation pack includes one Analysis user's guide:

The *ProTune Analysis User's Guide* describes how to use the ProTune Analysis graphs and reports after running a session in order to analyze system performance.

### Guide for Creating Scripts

The ProTune documentation pack includes one guide for creating scripts.

The *ProTune Creating Vuser Scripts User's Guide* describes how to create scripts using VuGen. When necessary, supplement this document with the online *ProTune Function Reference* and the *WinRunner User's Guide* for creating GUI scripts.

## ProTune Tutorial

The ProTune documentation pack includes one tutorial.

The *ProTune Tutorial* is a self-paced guide that teaches you how to use ProTune. It instructs you on how to create tests and run them on your production system, and use the test results to tune your system.

| <b>For information on</b>     | <b>Look here...</b>                                |
|-------------------------------|--|
| Installing ProTune            | <i>ProTune Installation Guide</i>                  |
| The ProTune testing process   | <i>ProTune Console User's Guide</i>                |
| Creating scripts              | <i>ProTune Virtual User Generator User's Guide</i> |
| Creating and running sessions | <i>ProTune Console User's Guide</i>                |
| Analyzing test results        | <i>ProTune Analysis User's Guide</i>               |

## Documentation Updates

Mercury Interactive is continuously updating its product documentation with new information. You can download the latest version of this document from Mercury Interactive's Customer Support Web site (<http://support.mercuryinteractive.com>).

### To download updated documentation:

- 1** In the Customer Support Web site, click the **Documentation** link.
- 2** Select the product name.

Note that if ProTune does not appear in the list, you must add it to your customer profile. Click "My Account" to update your profile.



- 3 Click **Retrieve**. The Documentation page opens and lists all the documentation available for the current release and for previous releases. If a document was recently updated, **Updated** appears next to the document name.
- 4 Click a document link to download the documentation.

## Typographical Conventions

This book uses the following typographical conventions:

|                   |  |
|-------------------|--|
| <b>1, 2, 3</b>    | Bold numbers indicate steps in a procedure.  |
| ►                 | Bullets indicate options and features.   |
| >                 | The greater than sign separates menu levels (for example, <b>File &gt; Open</b> ).   |
| <b>Stone Sans</b> | The <b>Stone Sans</b> font indicates names of interface elements on which you perform actions (for example, “Click the <b>Run</b> button.”).           |
| <b>Bold</b>       | <b>Bold</b> text indicates method or function names  |
| <i>Italics</i>    | <i>Italic</i> text indicates method or function arguments, file names or paths, and book titles.   |
| Arial             | The Arial font is used for examples and text that is to be typed literally.  |
| <>                | Angle brackets enclose a part of a file path or URL address that may vary from user to user (for example, < <i>Product installation folder</i> >\bin). |
| [ ]               | Square brackets enclose optional arguments.  |
| { }               | Curly brackets indicate that one of the enclosed values must be assigned to the current argument.  |
| ...               | In a line of syntax, an ellipsis indicates that more items of the same format may be included.   |



# Part I

---

## Understanding ProTune



# 1

---

## Introducing ProTune

ProTune offers a proactive solution for validating and optimizing the capacity of an application and its underlying infrastructure to process business transactions. It combines tuning processes, key technologies, and integrated tuning knowledge into a flexible and easy-to-use software package for tuning deployment and production systems.

Mercury Interactive's Safe Deployment System™ (SDS) is key to delivering the advanced capabilities in ProTune. SDS uses a combination of technology and a knowledge base that has evolved over years of experience in successful customer deployments. The SDS methodology provides a systematic approach to deployment tuning, which includes:

- ▶ **Systematic Identification:** ProTune automates the process of infrastructure and application optimization by examining the system as a whole and finding problems in a logical step-by-step process.
- ▶ **Problem Isolation:** ProTune pinpoints the precise areas where bottlenecks may occur by using expert-designed component tests.
- ▶ **Expert Recommendations:** ProTune recommends the correct action using a built-in knowledge base and tuners that reflect the learning gathered through years of customer engagements and working with all of Mercury Interactive's key infrastructure vendors.
- ▶ **Automated Improvements and Validation:** ProTune automates the process of making configuration changes directly to the devices on the system and automatically validates these changes to ensure they have changed the system in a positive way.

This section describes:

- ▶ ProTune's Features
- ▶ Tuning Workflow

## ProTune's Features

ProTune includes the following unique features for tuning deployment and production systems.

### **System Topology Mapping**

ProTune's System Topology window lets you draw a map of your system, defining the individual components and their properties. The System Topology window provides the system-wide perspective necessary for tuning multiple tier systems, as well as the capability to perform drill-down tuning exercises on any component.

### **Automatic Assignment of Monitors**

Complex systems can contain hundreds of counters and statistics. A user can find it difficult to know which counters are critical for performance. To eliminate this confusion, you can use ProTune to automatically select the monitors most critical for each component's performance.

### **Configurable Alerts**

*Configurable alerts* allow for fail-safe control of your tuning exercise, so that your deployment or production system is protected during the exercise. ProTune lets you establish safe performance thresholds. If a danger point is reached during a test, ProTune can halt or scale down the test, or stop increasing the load on the component.

In addition, by using alerts to inform you when errors occur, you can identify bottlenecks on complex multiple tier systems. Instead of running 4 or 5 different tools, which require you to look at several different screens at the same time to determine where a problem exists, ProTune's alerts locate and display the cause of the problem.

### **Canned Scripts**

To save time and effort, ProTune includes a large number of pre-configured tuning exercises that are automatically configured for your topology. These are known as *canned scripts*.

## Tuners

Once bottlenecks have been identified, *tuners* display your components' current settings, recommend which settings to change, and let you tune your components remotely from one location.

Tuners also help you manage your configurations, allowing you to roll back any changes made in ProTune, so you can restore your system to its original state.

## Performance Tuning Network

Mercury Interactive's Performance Tuning Network is a knowledge base relating specifically to performance tuning. ProTune uses the knowledge base to give you tips on locating bottlenecks and tuning your components, and provides access to it via the Help menu.

## Flexibility and Scalability

To provide the broadest flexibility and greatest scalability possible, ProTune includes the entire suite of Mercury Interactive's monitoring technologies. These include multi-layer tuning, capacity planning, security validation under load, capacity limit and reliability testing.

# Tuning Workflow

ProTune guides you through the tuning process in a few easy steps: topology, design, execute, analyze, and tune.

To tune a system, you do the following:

- 1 Identify the goals for tuning. This includes:
  - Desired number of concurrent users
  - Expected transaction rate
  - Expected response time
  - Maximum acceptable error rate

- 2** Create a session. This includes:
  - define the production system's topology, using ProTune's topology mapping tool.
  - assign monitors
- 3** Create session steps to test the suspected components. This includes:
  - choose scripts to run tests against components. ProTune's canned scripts save you time and provide structure and expert guidance on how to begin your tuning sessions. To expand the functionality of ProTune's recommended scripts, you can also create custom scripts for your application.
  - define alerts for each session step. These measure the state of your systems and establish safe performance thresholds.
  - choose profiles for each session step
- 4** Execute session steps. This is where you test the components of your system. ProTune's monitors and analysis display your performance in real time. During the session execution, ProTune records your application's performance under different loads.
- 5** View analysis data resulting from the session step to analyze the application's performance.
- 6** After ProTune helps to identify the problem areas, tune the system, in line with the session step results. You can use ProTune's integrated tuners to make changes to system component properties from the Console machine.
- 7** Run the session step again to validate the improvement in performance.
- 8** Continue to tune the system and run session steps based on the session step results, until you achieve your performance goal.



# Part II

---

## Mapping a Business Process



# 2

---

## Creating a Topology

The first step in mapping the business process is defining a topology of the servers used in the business process. This chapter describes how to create a topology.

This chapter discusses:

- ▶ Building a Topology Diagram
- ▶ Setting the Component Properties
- ▶ Selecting Monitors
- ▶ Automatically Assigning Monitors and Alerts
- ▶ Accessing the SiteScope Administration Console

## About Creating a Topology

The first step in creating a test for tuning your session is defining a topology. The topology refers to the architecture of the system. You indicate the servers and their types, such as database, Web, or application. You also define how the servers are connected. The topology also includes other hardware such as routers, firewalls, and load generator machines.

ProTune provides several topology templates for typical network configurations. Each template contains a number of different components.

Following is a list of ProTune's topology templates:

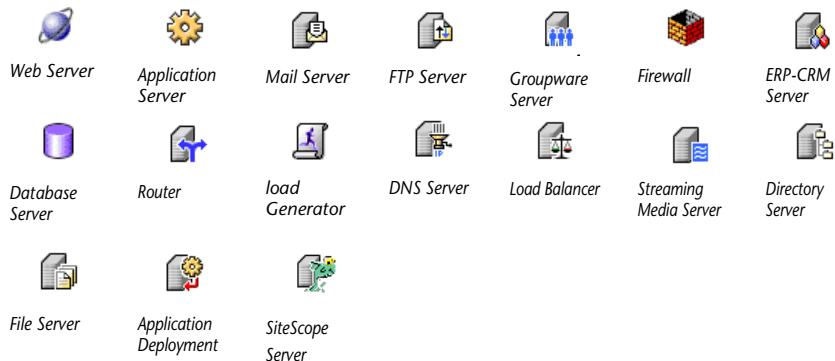
- ▶ Single Application Server
- ▶ Single Database Server
- ▶ Single Web Server
- ▶ 2-Tier WebAppServer
- ▶ 2-Tier WebDB
- ▶ 3-Tier FWLdbWebAppDb
- ▶ 3-Tier LdbWebAppDb
- ▶ 3-Tier Two System Arch
- ▶ Three Tier WebAppDb
- ▶ Siebel 5.x Three Tier
- ▶ Siebel 7.x Three Tier
- ▶ PeopleSoft Applications Three Tier
- ▶ SAP Enterprise Three Tier
- ▶ SAP GUI Three Tier
- ▶ SAP Enterprise Portal Three Tier
- ▶ ERP-CRM General Three Tier

You can use one of these templates as it is, or modify it. You can also create your own topology from an empty template. When you add items to the topology, ProTune arranges the elements in the most common layout. You can rearrange this layout or delete any unnecessary components.

## Building a Topology Diagram

Before building a topology diagram, make sure you have a clear understanding of your system architecture. An incorrect representation of your system will result in inaccurate results.

Using ProTune's design window, you can drag topology elements from the design palette into the topology diagram. The design palette contains the following elements:

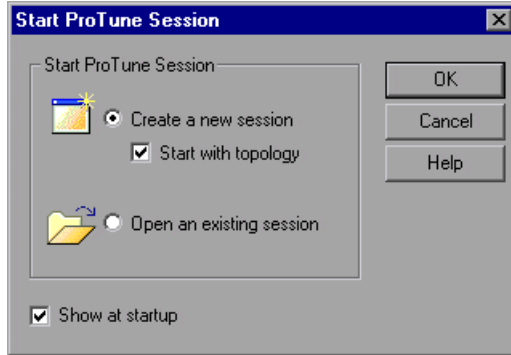


You can map a physical component to more than one logical topology element. For example, if a computer hosts both a Web server and a database server, you can define two elements—a Web server element and a database server element—and map them to the same computer.

You can use an existing topology, base your topology on a template, or create a new topology from scratch. ProTune allows you to modify templates and existing topologies to suit your specific needs.

**To open the System Topology window:**

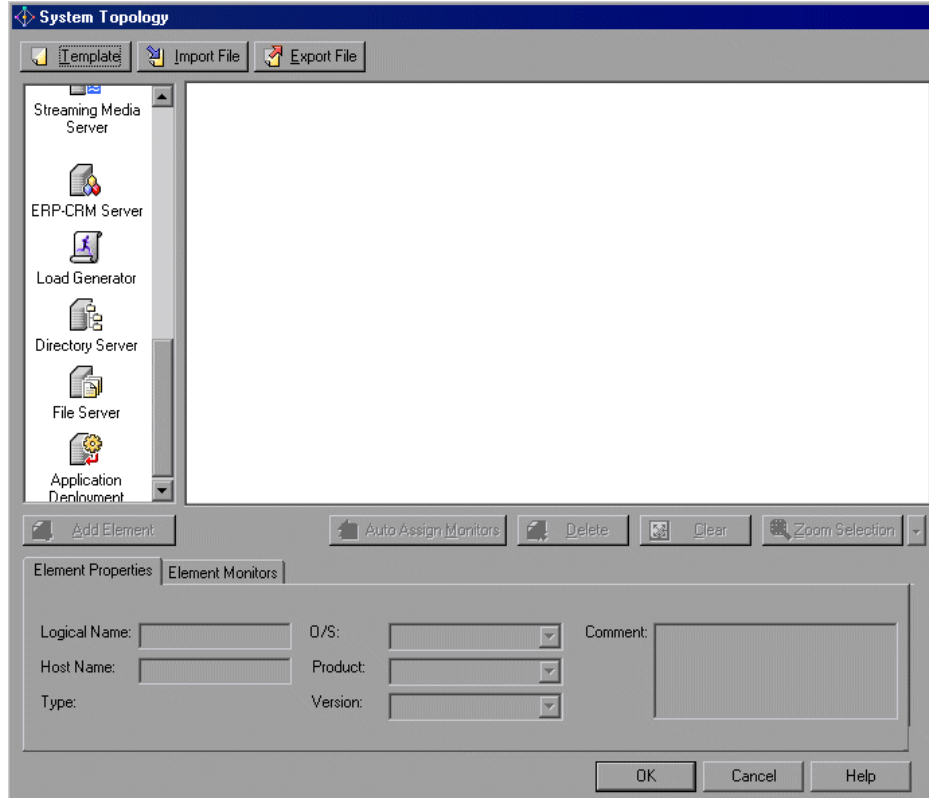
- 1** Invoke ProTune Console by choosing **Start > Programs > ProTune > Console**. ProTune displays the Start ProTune Session window:



You can create a new session or open an existing one. If you choose to create a new session, you can also decide whether you want the System Topology window to be displayed immediately.

- 2** Click **OK**.
- 3** If the Console is running but the System Topology window is not displayed, click the **Topology** button or choose **Tools > System Topology** to open it.

The System Topology window is displayed:




---

**Note:** You can also open the System Topology window by double-clicking the white space in the **Execute** tab's topology pane. See “Running a Session,” on page 121.

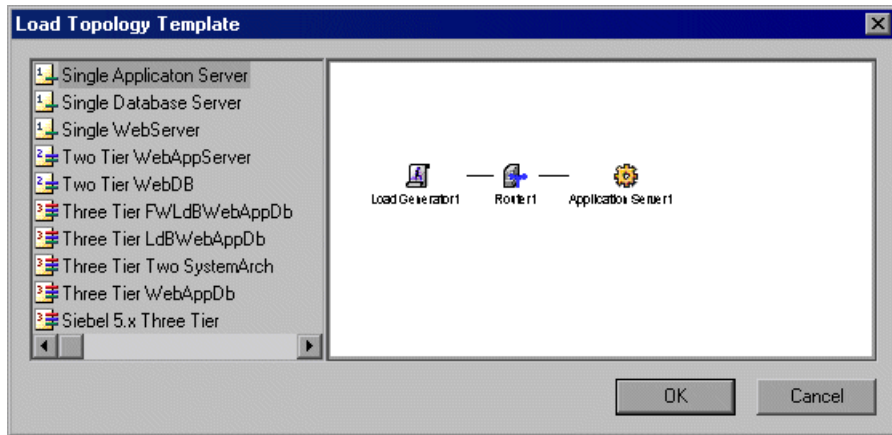
---

### To import an existing topology diagram:

- 1** Click **Import File** and browse to the desired topology (a file with a *.tpl* extension).
- 2** Click **Open**. ProTune displays your topology.

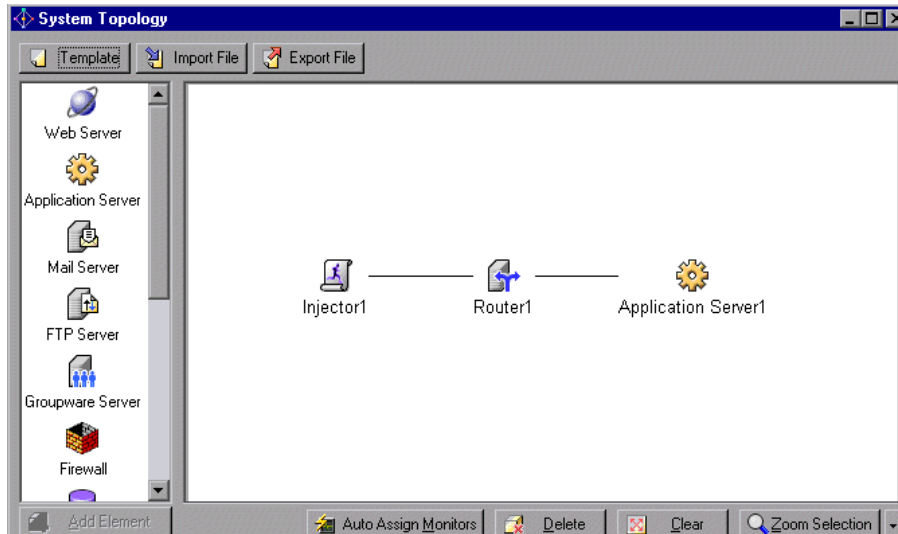
**To create a topology diagram from a template:**

- 1 Click **Template** to open the Load Topology Template dialog box.



- 2 Select a template and click **OK**.

The topology diagram you selected appears in the right pane, as in the following example:






**To create or modify a topology:**

- 1** To add a component to a topology, select the component in the left pane and click **Add Element**, or drag the component from the left pane into the right pane.
- 2** To change a component's position, click the component and drag it to its new position in the topology.
- 3** To connect components in the topology, drag a line from one component to the other.
- 4** To delete a component, select the component by clicking it in the diagram, and click the **Delete** button or press your keyboard's Delete key.
- 5** To delete all the components from the topology diagram, click **Clear**.

**To zoom in on a component:**

- 1** Select the component and click **Zoom Selection**.
- 2** Click the arrow to the right of the **Zoom Selection** button to view all zoom options:



- 3** Click **Zoom In** to increase the size of the topology diagram.
- 4** Click **Zoom Out** to reduce the size of the topology diagram.
- 5** Click **Zoom to Original Size** to restore the original size of the diagram.

**To save the topology diagram for use in another session:**

- 1** Click **Export File**.
- 2** Browse to the desired location, specify a filename with a *.tpl* extension, and click **Save**.

## Setting the Component Properties

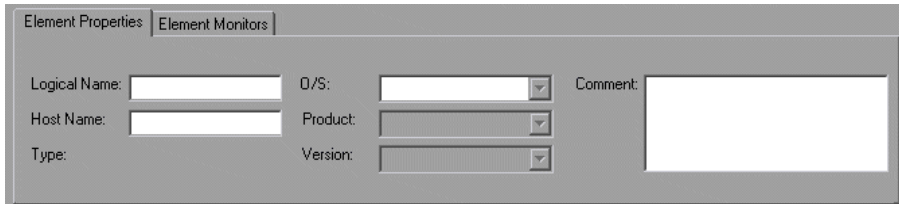
After you create a topology diagram, you specify the properties for each component. ProTune builds a tuning session based on these settings.

You can specify the following properties for each component:

- logical name
- host name
- operating system
- product
- version

**To specify a component's properties:**

- 1** Select the component whose properties you want to specify.
- 2** Click the **Element Properties** tab.




- 3** In the Logical Name box, enter the name that you want ProTune to use for the component in the topology diagram, graphs, reports and other places where the component appears. You can use up to 255 characters. You can use the logical name to indicate the task performed by the component (for example, "Backup\_DB\_Server"). (Optional)
- 4** Specify the component's host name in the **Host Name** box, using up to 50 characters.

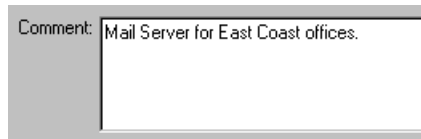
---

**Note:**

- If you do not specify the host name, ProTune copies the Logical Name to the Host Name field. If you specify neither host name nor logical name, ProTune will not allow you to select monitors for the element.
- If you delete a component's host name and then change the focus to another element or exit the topology window, ProTune displays the following warning:



- 
- 5** Select your component's operating system from the **O/S** list.
  - 6** Select the product or vendor for the selected component from the **Product** list.
  - 7** Select the application's version from the **Version** list.
  - 8** In the Comment field, add any useful information about the component, as in the following example, using up to 255 characters.

A screenshot of a text input field labeled "Comment:". The field contains the text "Mail Server for East Coast offices." and is surrounded by a light gray border.

## Selecting Monitors

After you define each component's properties, you select the measurements that you want to monitor for the component.

---

**Note:** You need to define the component properties before selecting the monitors, since ProTune uses the properties to locate the available monitors in its database.

---

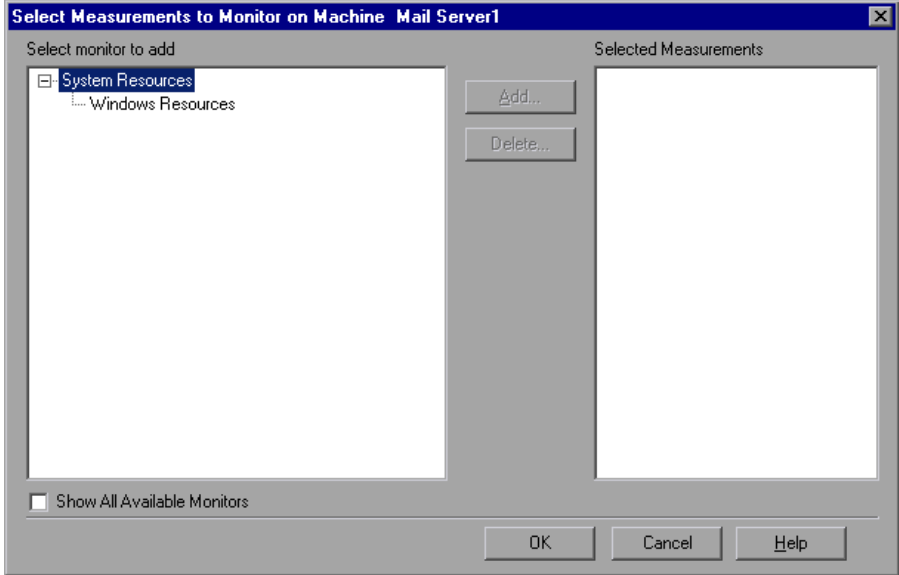
### To select a measurement to monitor:

- 1** In the topology diagram, select the component whose properties you want to monitor.
  - 2** Click the **Element Monitors** tab and then click **Add Monitor**. The **Select measurements to monitor** dialog box opens.
- 

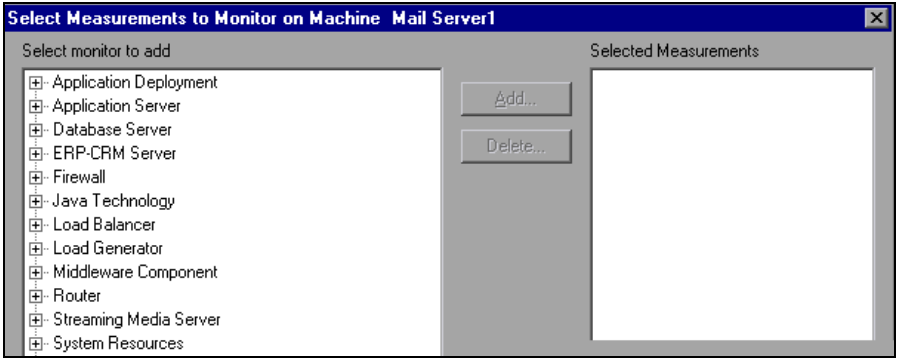
**Note:** To select monitors and measurements when you are not in the System Topology window, click the **Monitors** button on the main toolbar. See "Choosing Monitors and Measurements," on page 178.

---

If the ProTune knowledge base contains a list of monitors for your element, ProTune displays only those monitors available for the component in its specified operating system, product and version, as in the following example:

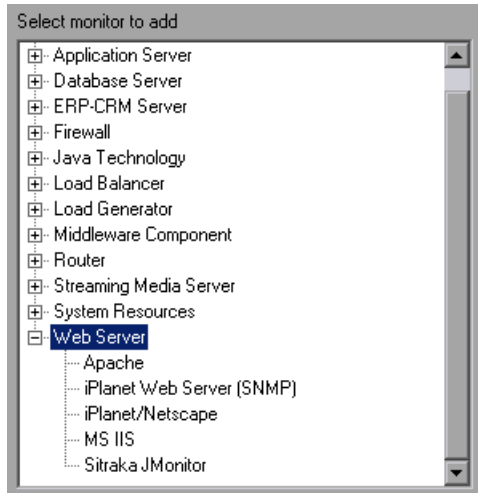


If the knowledge base does not contain a list of monitors for the specified element (possibly because you haven't specified all the element's properties), ProTune automatically displays all its monitors:



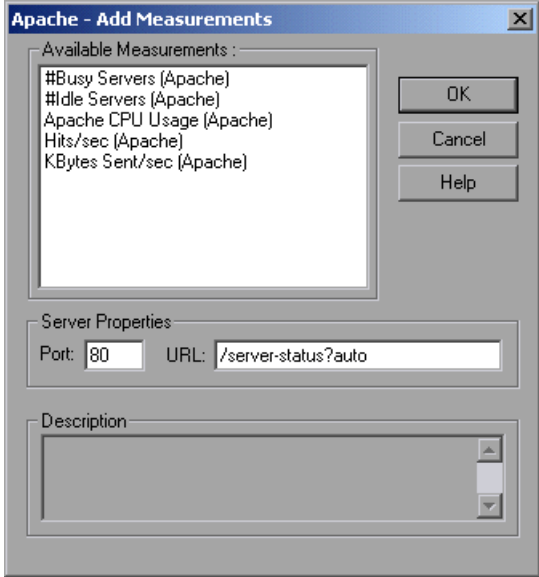
You can also display all the monitors in the knowledge base by checking the **Show All Available Monitors** box.

The full list of available monitors is displayed as a list of component types. If the monitor you need is not visible, expand the appropriate component type to display the list of components in your category. For example, clicking Web Server shows you the list of Web Server applications whose measurements you can monitor:



Click a component to select it and then click **Add**, or double-click the component. For most component types, ProTune opens a dialog box with a list of the component's available measurements.

For example, if you choose Apache, ProTune displays the following list of Apache measurements that you can monitor:

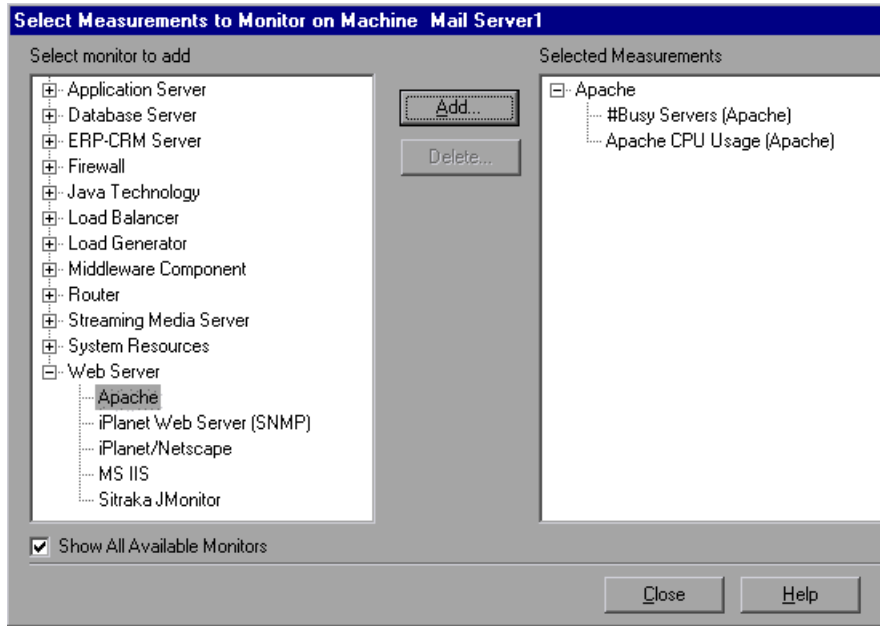


---


**Tip:** For information on how to specify the measurements specific to each component, see the appropriate chapter in the section “Monitoring a Session,” on page 173.

---

- 3 Select the measurements that you want to monitor, and click **OK**. ProTune adds the measurements to the Selected Measurements pane.



- 4 To delete a measurement from the Selected Measurements pane, select the measurement and click the **Delete** button.
- 5 Repeat steps 1 through 3 for all the components that you want to monitor, and then click **Close**. ProTune displays the newly added monitors in the Element Monitors section.



For some monitors, a **Test Monitor** button now appears in the Element Monitors section. Click this button to test whether you can access the monitor.



- 6 To delete a monitor that was previously assigned to the component, select the monitor in the **Element Monitors** tab and click **Delete Monitor**.
- 7 To add a measurement to the list of measurements for monitoring, select the monitor in the **Element Monitors** tab and click **Add Measurement**. ProTune opens the Add Measurements dialog box for the selected monitor.
- 8 To delete a measurement from the list of measurements for monitoring, select the monitor in the **Element Monitors** tab, select the measurement in the right pane, and then click **Delete Measurement**.

## Automatically Assigning Monitors and Alerts

ProTune's Auto Assign feature assigns monitors automatically to a selected component, based on the component's type. This allows you to skip the task of choosing monitors for your components.

---

**Note:** The Auto Assign feature is currently enabled only if the Console machine is running an English version of Windows.

---

For example, if your component is an Apache Web Server, the Auto Assign feature assigns to your component all the monitors for all the Apache measurements.

---


**Tip:** For details of ProTune's monitors, see "Monitoring a Session," on page 173.

---

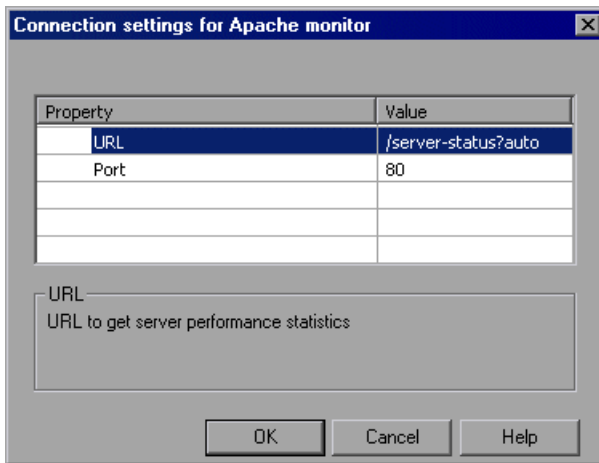
In addition to automatically assigning monitors, you can also automatically assign alerts. This assigns a default set of alerts for various measurements, based on definitions in ProTune's knowledge base. For information on alerts, see "Defining Alerts," on page 77.

**To automatically assign monitors and alerts to a component:**

- 1 In the topology diagram, select the component whose properties you want to monitor, and click **Auto Assign Monitors**. Alternatively, you can right-click the component and click **Auto Assign Monitors**. ProTune displays the Auto Assign Monitors dialog box:



- 2 To automatically assign alerts, check the **Auto Assign Alerts** box.
- 3 Click **OK**. If a monitor that is being assigned has configurable parameters or requires additional data, ProTune displays a Connection Settings dialog box for the monitor, as in the following example:



The dialog box displays one or more of the following parameters, depending on the monitor:

- Username
- Password
- Server Name
- URL
- Port

The dialog box typically allows you to insert values and change one or more of the displayed default values. After making your changes, click **OK**.

ProTune assigns the monitors to the selected components, and displays them in the Element Monitors section of the Topology window. If you chose to assign alerts, you can view the assigned alerts by clicking the **Alerts Definition** button.

## Accessing the SiteScope Administration Console

If you define a SiteScope Server element, ProTune allows you to access the SiteScope Administration Console via the Topology window.

**To access the SiteScope Administration Console:**

- Right-click the SiteScope Server element in the Topology window, and choose Open SiteScope. The SiteScope Administration Console opens in your browser, as in the following example.



# Part III

---

## Designing a Tuning Session



# 3

---

## Adding Session Steps

After creating a topology, you add steps to your session and specify the order in which they should be executed. Each step reflects a specific business process, or part of a whole business process. This chapter describes how to create and manage session steps.

This chapter discusses:

- Getting Started with Designing a Session
- Understanding Session Steps
- Adding Session Steps
- Managing Steps
- Specifying Step Execution Order
- Managing Scripts
- Setting an Initial Load (Manual Profiles)
- Configuring Script Details
- Using Relative Paths for Scripts

### About Adding Session Steps

For your tuning session, you create test steps. Each step performs a specific business process. ProTune provides a set of *canned* or prepared scripts. These scripts include tests for the Web Infrastructure.

You can use the canned scripts or write your own and incorporate them in the session steps.

You select one or more scripts for each session step. You can also indicate upon which servers to execute these steps. When selecting multiple scripts, you can specify whether the scripts should all be part of one session step or each script should be contained in a separate step. Since ProTune tests single steps, combining scripts in a step causes them to be tested together.

After setting up the steps for your session, you can indicate the number of Vusers you want to emulate. Vusers are virtual users that emulate real users. You can also choose a goal for each session step. For more information, see Chapter 7, “Scheduling Session Steps.”

## Getting Started with Designing a Session

Designing your tuning session consists of the following steps:

**Add Test Steps:** Add steps that emulate a business process to the session, for each server. For more information, see “Adding Session Steps,” on page 40.

**Manage Test Steps:** Manage the steps that you created by arranging their order, adding scripts, disabling them temporarily, or deleting them. For more information, see “Managing Steps,” on page 44.

**Specify Step Execution Order:** Specify which step should be executed after the current step completes execution, depending on whether the step succeeds or fails. For more information, see “Specifying Step Execution Order,” on page 45.

**Define a Schedule Profile:** Specify how and when the step should run, and whether it should be goal-oriented. For more information, see “Scheduling Session Steps,” on page 95.

**Assign Scripts:** Once you have defined a goal, you can use the built-in script or assign additional scripts to emulate other aspects of the business process. For more information, see “Managing Scripts,” on page 46.

**Set Alerts:** You set alerts for your session. ProTune issues an alert for the element that reaches the specified threshold. For more information, see Chapter 6, “Defining Alerts.”



**Set the Initial Load:** Indicate the number of Vusers you want to emulate in your session. For more information, see “Setting an Initial Load (Manual Profiles),” on page 50.

**Configure the Script:** Configure the values for your script. For more information, see Chapter 3, “Configuring Script Details.”

**Schedule the Step:** Set up a schedule for your session steps in order to automate them and simulate the true environment. For more information, see Chapter 7, “Scheduling Session Steps.”

**Set the Run-Time Settings:** Set the run-time settings for each script. These relate primarily to the pacing of the script, the think time, and the logging options. For information on configuring the run-time settings, refer to the *ProTune Creating Virtual User Scripts* guide.

## Understanding Session Steps

ProTune’s built-in scripts help you test the connection capacity or rate goals. You can create session steps in the following areas:

- Infrastructure
- Winsock
- Web Server
- FTP Server
- Mail Server
- Streaming Server
- Database Server
- PeopleSoft
- Siebel
- Security

## Infrastructure

The Infrastructure steps relate to the infrastructure of your network architecture. The available sub-step is:

**DNS Request Rate:** Checks the rate at which host names are resolved on the Domain Name Server. The script title as it appears in the script list is *dns\_rqst\_rate*. The supported server is the DNS server. **Note:** The DNS protocol is not supported by UNIX.

## Winsock

The Winsock steps relate to the TCP/IP and SSL connections of your servers. The available sub-steps are:

**TCP Connection Capacity:** Sustains simultaneous TCP connections with the target device. The script title as it appears in the script list is *tcp\_conn\_cpty*. The supported servers are Web servers, load balancers and database servers.

**SSL Connection Capacity:** Sustains simultaneous SSL connections with the target device. The script title as it appears in the script list is *tcp\_ssl\_conn\_cpty*. The supported servers are Web servers, load balancers and database servers.

## Web Server

The Web Server tests are used for testing Web server performance. The available sub-steps are:

**HTTP Connection Rate:** Generates HTTP requests against a target URL, without keep-alive enabled. Therefore, each new request must establish a new connection with the Web server. The script title as it appears in the script list is *http\_conn\_rate\_nokeepalive*. The supported servers are Web servers and load balancers.

**HTTP Request Rate:** Generates HTTP requests against a target URL, with keep-alive enabled. Therefore, each new request reuses the same connection with the Web server. Note: the Web server must be configured to support keep-alive connections in order for this script to work properly. The script title as it appears in the script list is *http\_rqst\_rate\_keepalive*. The supported servers are Web servers and load balancers.

**HTTP Downstream Bandwidth:** Generates downstream data transmission by a continuous download of a large file from the Web server via HTTP protocol. The HTTP requests are set with keep-alive enabled. Therefore each new request re-uses the same connection with the web server. The script title as it appears in the script list is *http\_downstream\_bandwidth*. The supported servers are Web servers and load balancers.

### **FTP Server**

The FTP Server tests are used for testing FTP server performance. The available sub-steps are:

**FTP Connection Capacity:** Sustains simultaneous FTP connections with the file server. The script title as it appears in the script list is *ftp\_conn\_cpty*. The supported servers are FTP servers, Web servers and load balancers.

**FTP Get File Rate:** Generates FTP GET requests to download a specific file from an FTP server. The script title as it appears in the script list is *ftp\_get\_rate*. The supported servers are FTP servers, Web servers and load balancers.

**FTP Put File Rate:** Generates FTP PUT requests to upload a specific file to an FTP server. The script title as it appears in the script list is *ftp\_put\_rate*. The supported servers are FTP servers, Web servers and load balancers.

### **Mail Server**

The Mail Server tests are used for testing the following mail protocols: SMTP, MAPI, POP3 and IMAP. The available sub-steps are:

**SMTP Connection Capacity:** Sustains simultaneous SMTP connections with the mail server. The script title as it appears in the script list is *smtp\_conn\_cpty*. If a connection is closed by the server, it is re-established in order to sustain the same number of sessions. The supported servers are Mail Servers and Application Servers.

**SMTP Send Mail:** Submits e-mail messages to a mail server via the SMTP protocol. The script title as it appears in the script list is *smtp\_send\_mail*. The supported servers are Mail Servers and Application Servers.

**MAPI Connection Capacity:** Sustains simultaneous MAPI sessions with MS Exchange Server. These sessions are created only once. If the session terminates, it is not re-established. The script title as it appears in the script list is *mapi\_conn\_cpty*. The supported servers are Mail Servers and Application Servers.

**MAPI Send Mail:** Sends an e-mail with a file attachment to the specified recipient, using the given MS Exchange profile. The script title as it appears in the script list is *mapi\_send\_mail*. The supported servers are Mail Servers and Application Servers.

**POP3 Connection Capacity:** Creates simultaneous POP3 session connections with the mail server. If a mail server does not allow enough connections, it will reject new attempts to connect, and performance will suffer. The purpose of this script is to test the maximum number of connections that the specified mail server allows. The script title as it appears in the script list is *pop3\_conn\_cpty*. The supported servers are Mail Servers and Application Servers.

**POP3 Retrieve Mail:** Puts stress on the infrastructure involved in retrieving mail messages on a POP3 server. The script title as it appears in the script list is *pop3\_retrieve\_mail*. The supported servers are Mail Servers and Application Servers.

**IMAP Connection Capacity:** Sustains simultaneous IMAP sessions with the mail server. The script title as it appears in the script list is *imap\_conn\_cpty*. The supported servers are Mail Servers and Application Servers.

**IMAP Search Mail:** Searches for mail in the specified folder using the given search criteria. The script title as it appears in the script list is *imap\_search\_mail*. The supported servers are Mail Servers and Application Servers.

**IMAP Store Mail:** Stores a mail message in the specified mail folder using the IMAP protocol. The *imap\_mail.dat* file is an IMAP message supplied with ProTune. To use a different message, modify or replace this file. The script title as it appears in the script list is *imap\_store\_mail\_file*. The supported servers are Mail Servers and Application Servers.

## Streaming Server

The Streaming Server tests are used for testing streaming servers that support the Real and Microsoft Media Stream protocols. The available sub-steps are:

**Real Connection Capacity:** Sustains simultaneous connections through the RTSP protocol with the Real server. The script title as it appears in the script list is *rtsp\_conn\_cpty*. The supported servers are Web Servers and Application Servers.

**RealPlayer Play Media:** Plays a media stream through the RTSP protocol. The script title as it appears in the script list is *rtsp\_play\_media*. The supported servers are Web Servers and Application Servers.

**MMS Play Media:** Plays a media stream via the Microsoft Media Stream (MMS) protocol. The script title as it appears in the script list is *mms\_play\_media*. The supported servers are Web Servers and Application Servers.

---

**Note:** Before running this script in the context of a load test, make sure that your Streaming Media Server public directory includes the file *wmload.asf* (this is a Microsoft requirement), or, if you are not running in the context of a load test, add the *mms\_disable\_host\_check()* function to your script.

---

## Database Server

The Database Server tests are used for testing database servers that support the ODBC driver and ADO-DB interface. The available sub-steps are:

**ADO-DB Connection Rate:** Opens a new connection to the SQL server, executes the specified query on the server and closes the connection. The script title as it appears in the script list is *adodb\_open\_sql\_close\_rate*. The supported servers are Database Servers.

**ADO-DB SQL Query Rate:** Executes the specified SQL statement on the server. The connection is created only once at the script initialization stage. The script title as it appears in the script list is *adodb\_sql\_query\_rate*. The supported servers are Database Servers.

**ODBC SQL Connection Rate:** Executes the specified SQL statement on the database server. A new connection is opened in each iteration. The script title as it appears in the script list is *odbc\_sql\_connection\_rate*. The supported servers are Database Servers.

**ODBC SQL Query Rate:** Executes the specified SQL statement on the database server. The connection to the database server is created only once at the script initialization stage. The script reuses the same connection for all of its iterations, but creates a new database cursor for each iteration. The script title as it appears in the script list is *odbc\_sql\_query\_rate*. The supported servers are Database Servers.

**ODBC SQL Query Rate (reuse cursor):** Executes the specified SQL statement on the database server. The connection to the database server is created only once at the script initialization stage. The script reuses the same connection and database cursor for all of its iterations. The script title as it appears in the script list is *odbc\_sql\_query\_rate\_reuse\_cursor*. The supported servers are Database Servers.

## **PeopleSoft**

The PeopleSoft tests are generalized business processes used to create load on most of the application's components. Steps are available for PeopleSoft 8.1x and PeopleSoft 8.4x.

The available sub-steps for PeopleSoft 8.1x are:

**Financials Process Financial Information:** Browses the Process Financial Information tree in the PeopleSoft8.1x Financials application, using the following schema:

**Process financial Information > Review Financial Information > Inquire > Journal**

The PeopleSoft login and logout are not included in the iteration loop.

**Financials Browse Admin Procurement:** Browses the Administer Procurement tree in the PeopleSoft8.1x Financials application, using the following schema:

**Administer Procurement > Create Payments > Inquire > Voucher Inquiry**

The PeopleSoft login and logout are not included in the iteration loop.

**Financials Browse Manage Assets:** Browses the Manage Assets tree in the PeopleSoft8.1x Financials application, using the following schema:

Manage Assets > Manage Assets > Use > Asset ExpressAdd

The PeopleSoft login and logout are not included in the iteration loop.

**Financials Browse Manage Sales Activities:** Browses the Manage Sales Activities tree in the PeopleSoft8.1x Financials application, using the following schema:

Manage Sales Activities > Maintain Customers > Use > General Information

The PeopleSoft login and logout are not included in the iteration loop.

The available sub-steps for PeopleSoft 8.4x are:

**CRM Browse FieldService Service Orders:** Browses the FieldService tree in the PeopleSoft8.4x CRM application, using the following schema:

FieldService > Service Orders

The PeopleSoft login and logout are not included in the iteration loop.

**CRM Browse FieldService Reports:** Browses the FieldService tree in the PeopleSoft8.4x CRM application, using the following schema:

FieldService > Reports > Expense Report

The PeopleSoft login and logout are not included in the iteration loop.

**CRM Browse Sales Leads:** Browses the Sales tree in the PeopleSoft8.4x CRM application, using the following schema:

Sales > Leads > Lead Details

The PeopleSoft login and logout are not included in the iteration loop.

**CRM Browse Sales Opportunities:** Browses the Sales tree in the PeopleSoft8.4x CRM application, using the following schema:

Sales > Opportunities > Opportunity Details

The PeopleSoft login and logout are not included in the iteration loop.

## Siebel

The Siebel tests are generalized business processes used to create load on most of the application's components. The available sub-steps are:

**Siebel 7.x Portal User Login:** Logs on to the Siebel 7.x portal, and immediately logs off.

## Security

The Security tests simulate Denial of Service (DoS) attacks on servers.

---

**Note:** Refer to the documentation on each script to see the platforms on which it can be run.

---

The available sub-steps are:

**Bonk DoS Attack:** Disrupts the sequence of IP packet fragments by changing the information in the fragment's header information. Large IP packets are frequently broken up into fragments when they pass through routers on the Internet. The fragment's header includes an offset field that specifies the fragment's position in the packet. The attacker creates a sequence of fragments with offset fields greater than the header length. If the destination host cannot handle the overlapping fields, it crashes, freezes or reboots. The script title as it appears in the script list is *ddos\_bonk*. **Note:** This script can be run only on a load generator running Windows XP.

**Ping of Death DoS Attack:** Creates a packet whose size is too great for the destination host to handle. The TCP/IP specification allows creating a packet with a maximum size of 65536 bytes. The attacker uses the Ping utility to create an IP packet with a size greater than 65536 bytes. This can cause the destination system to crash, freeze, or reboot. The script title as it appears in the script list is *ddos\_ping\_of\_death*. **Note:** This script can be run only on a load generator running Windows XP.



**SYN Flood DDoS Attack:** Causes the server's queue to overflow by filling it with half-open (also known as *pending*) connections. To do this, it initiates a TCP/IP connection, using IP spoofing to specify an unreachable or down client as the sender. The server mistakenly identifies each TCP/IP connection as a genuine one coming from a different IP address, and responds by sending a SYN/ACK signal to the client. It then waits for an ACK signal (which never arrives) to complete the three-way handshake. The connection is stored in the server's queue, which keeps growing as the attacker creates more half-open connections, until it eventually overflows, preventing the server from receiving any new connections. Although the connections eventually expire, the attacker can initiate new fake connections at a rate faster than the server can cause them to expire. The script title as it appears in the script list is *ddos\_syn\_flood*. **Note:** Can be run only on machines running Windows 2000 or Windows XP.

**Targa3 DoS Test:** Sends combinations of uncommon IP packets to hosts to generate attacks using invalid fragmentation, protocol, packet size, header values, options, offsets, TCP segments, routing flags, and other unknown or unexpected packet values. This test is useful for testing IP stacks, routers, firewalls, NIDs (Network Intrusion Detectors) and other similar components for stability and reactions to unexpected packets. Some of these packets might not pass through routers with filtering enabled; tests with the source and destination hosts on the same Ethernet segment give the best effects. The script title as it appears in the script list is *ddos\_targa3*. **Note:** This script can be run only on a load generator running Windows 2000 or Windows XP.

**Teardrop DoS Attack:** Disrupts the sequence of IP packet fragments by changing the information in the fragment's header information. Large IP packets are frequently broken up into fragments when they pass through routers on the Internet. The fragment's header includes an offset field that specifies the fragment's position in the packet. The attacker creates a sequence of fragments with offset fields that overlap. If the destination host cannot handle the overlapping fields, it crashes, freezes or reboots. The script title as it appears in the script list is *ddos\_teardrop*. **Note:** This script can be run only on a load generator running Windows XP.


**UDP Echo DDoS Attack:** Sends a UDP datagram to the destination address. In particular, this illustrates the danger of having the UDP echo service turned on in */etc/inetd.conf* on many versions of UNIX. Consider the result if the source address and port are set to localhost and 7 respectively: the inetd in FreeBSD 2.2 seems to detect this denial of service attack, but many UNIX variants do not. The script title as it appears in the script list is *ddos\_udp\_echo*. **Note:** Can be run only on machines running Windows 2000 or Windows XP.

## Adding Session Steps

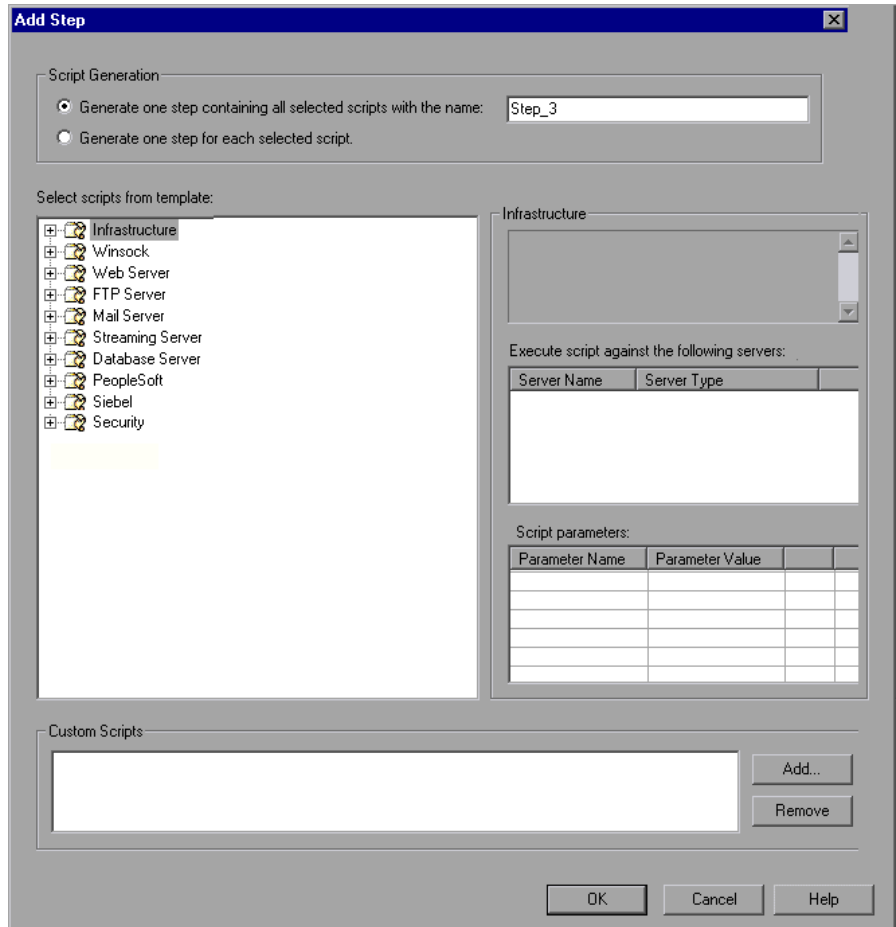
Before creating session steps, you must create or load a topology (see Chapter 2, “Creating a Topology.”) ProTune displays all of the relevant servers and session steps. When you click **OK** in the Topology window, ProTune opens the Session window's Design tab.

You can choose from the canned scripts or add custom scripts created with the ProTune Virtual User Generator.

### To create a new step:

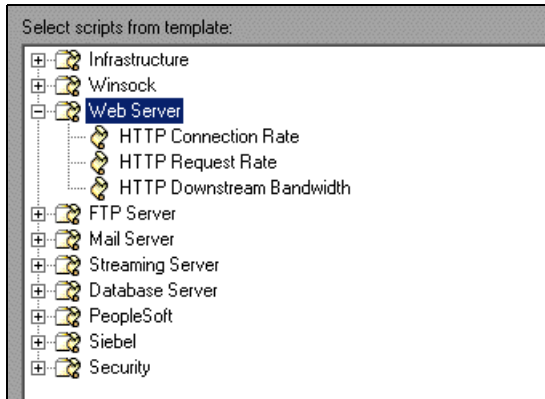


- 1 In the Console window, click the **Add Step** button or choose **Session > Add Step**. The Add Step dialog box opens.




- 2 You now specify whether you want to create a separate step for each selected script, or to include multiple scripts in one step. Click the appropriate radio button in the Step Generation section at the top of the window. If you are creating a step with multiple scripts, enter the step name in the text box, or accept ProTune's default. If you choose to create separate steps for each script, ProTune assigns a name to each step.

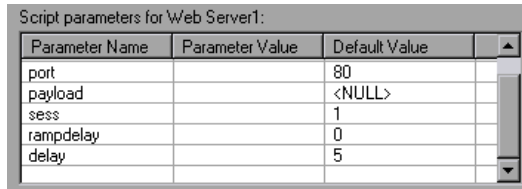
- 3 The left pane lists the types of canned scripts. Clicking a type expands it and shows all of its scripts. For example, clicking Web Server shows you the following scripts: HTTP Connection Rate, HTTP Request Rate and HTTP Downstream Bandwidth.



- 4 Expand the script type that you want to use (if it is not already expanded).
- 5 Click the script that you want to test. Note that ProTune displays information about the script in the upper section of the window's right side.
- 6 In the middle section of the right pane, ProTune displays the servers against which the script can be run. To specify the servers against which to run the script, check the boxes adjacent to those servers.



In the Script Parameters pane, ProTune displays a table containing the arguments for the selected script and their default values.



| Parameter Name | Parameter Value | Default Value |
|----------------|-----------------|---------------|
| port           |                 | 80            |
| payload        |                 | <NULL>        |
| sess           |                 | 1             |
| rampdelay      |                 | 0             |
| delay          |                 | 5             |


- To specify a different value, click the appropriate cell in the **Parameter Value** column and enter the value.

---

**Note:** A command line parameter is limited to one line.

---


- To add a custom script to a test, click the **Add** button in the lower right section of the window, browse to the script that you want to use, and click **Open**. The script name is added to the Custom Scripts pane.



- Repeat steps 4 through 8 for all the scripts and servers that you want to add to your session.
- Click **OK**.

In the left pane, ProTune displays all the steps you defined. It displays separate steps for each server—if the same step is to be run against three servers, the step is displayed three times. You can see which scripts are included in each step (both canned and custom scripts).

In the following example, the session includes three steps: *Check request rate*, *Check connection capacity*, and *Measure bandwidth*.




When you click a step name, the scripts that the step uses are displayed in the lower right pane of the window.

| Groups                              |                     |  |       |                       |
|-------------------------------------|---------------------|--|-------|-----------------------|
|                                     | Name                | Script Path  | %     | Load Generators       |
| <input checked="" type="checkbox"/> | http_rqst_rate_keep | C:\...\ProTune\data\scripts\infrastructure\http\http_rqst_rate_keepalive | 100 % | <All Load Generators> |

The Description section, on the right side of the screen, is useful for entering notes about the session. You can edit this section at any time. The text that you enter is saved when you save the session.

## Managing Steps

After you add steps to your session, you can manage them in several ways. The step commands are available from the right-click menu and the toolbar.



**Add a step:** In the session tree, select the session or a step, and then click **Add Step** or choose **Add Step** from the right-click menu.



**Remove a step:** Select a step, and then click **Remove Step** or choose **Remove Step** from the right-click menu.

**Rename a step:** Select a step and choose **Rename Step** from the right-click menu to enable editing of the step name. Specify a new name for the step.

**Duplicate a step:** To place a copy of a step directly below it, select it and choose **Duplicate Step** from the right-click menu.

**Move a step up:** To move a step up in the list, select it and choose **Step Up** from the right-click menu.

**Move a step down:** To move a step down in the list, select it and choose **Step Down** from the right-click menu.

**Disable a step:** To disable a step from the tuning session, select it and choose **Disable Step** from the right-click menu.

**Enable a step:** To enable a step that was disabled, select the step and choose **Enable Step** from the right-click menu.

## Specifying Step Execution Order

When you create a step, you can specify which step should be executed after the step completes execution. This allows you to define a methodology: if the step succeeds, you can specify that Step A should subsequently be executed; if it fails, the user should execute Step B. You can then specify subsequent steps for Step A and Step B, and so on.

For the purpose of determining which step should subsequently be executed, a failed step is one that ends when an alert causes all the Users to halt execution. If the step ends in any other way, it is considered to be a successful step.

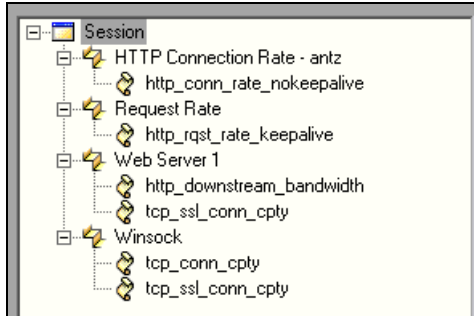
If you do not specify a subsequent step (for success, failure, or both), execution stops when the current step has completed.

### To specify subsequent steps:

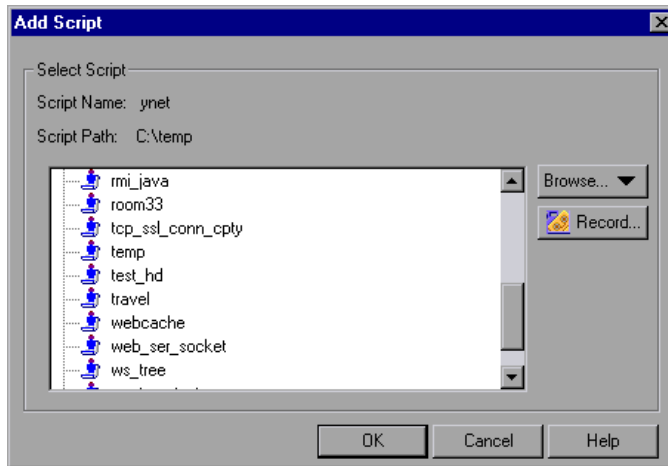
- 1** Create the subsequent steps (see “Adding Session Steps,” on page 40).
- 2** In the **Design** tab, select the earlier step in the left pane.
- 3** In the **If step fails...** box, choose or type the name of the step that should be executed if the current step fails. (Optional)
- 4** In the **If step succeeds...** box, choose or type the name of the step that should be executed if the current step completes successfully. (Optional)

## Managing Scripts

After you add scripts to your steps, you can add, delete, or view them from the step tree in the left pane of the Session window.



**Add a script:** Select the step to which you want to add the script, and choose **Add Script** from the right-click menu or click the **Add Script** button. The Add script dialog box opens.



To add an existing script, select one of the displayed scripts and click **OK**. To change the path, click **Browse** and select an alternate path. To record a new script in the ProTune Virtual User Generator, click **Record**. For more information about recording scripts, refer to the *ProTune Creating Virtual User Scripts* guide.





**Delete a script:** Select the script you want to delete, and choose **Delete Script** from the right-click menu.



**View or Modify a script:** To view or modify a script, select it and click the **Modify the Vuser script** button. ProTune opens the Virtual User Generator, so that you can edit the script. For more information on editing scripts, refer to the *ProTune Creating Vuser Scripts User's Guide*.

## Working with the Script List Window

The Script List shows all the scripts assigned to the current step.

|                                     | Script Name       | Script Path  | %       | Load Generators       |
|-------------------------------------|-------------------|--|---------|-----------------------|
| <input checked="" type="checkbox"/> | web_ser_socket    | R:\LR_TESTS\web_ser_socket                         | 14.29 % | <All Load Generators> |
| <input type="checkbox"/>            | test_hd           | R:\test_hd   | 14.29 % | <All Load Generators> |
| <input type="checkbox"/>            | ws_tree_test      | C:\temp\ws_tree_test                               | 14.29 % | localhost             |
| <input checked="" type="checkbox"/> | browserlevel      | R:\LR_TESTS\browserlevel                           | 14.29 % | <All Load Generators> |
| <input checked="" type="checkbox"/> | tcp_conn_cpty     | E:\...\network\tcp_conn_cpty\tcp_conn_cpty         | 14.29 % | <All Load Generators> |
| <input checked="" type="checkbox"/> | webcache          | R:\LR_TESTS\webcache                               | 14.29 % | <All Load Generators> |
| <input checked="" type="checkbox"/> | tcp_ssl_conn_cpty | E:\...\network\tcp_ssl_conn_cpty\tcp_ssl_conn_cpty | 14.29 % | <All Load Generators> |

You can further manage your scripts from this window:

**Sort the scripts:** To sort the scripts by their name, path, percentage or Load Generator, click the title of the desired column.

**Enable a Script:** Select the check box adjacent to the script name. All selected scripts will be executed during the tuning session.

**Disable a Script:** Clear the check box adjacent to the script name. All disabled scripts will be ignored during the tuning session.


**Rename a Script:** Double-click on an item in the **Script Name** column and enter the desired name.

**Add a Script to the List:** Click in the **Script Name** column in the next available row, and click on the arrow to the right of the box. A list box opens showing all scripts in the most recent path. To add a script from the list, select it and click **OK**. To change the path, click **Browse** and select an alternate path.

**Modify the Vuser Percentage:** When you assign multiple scripts to a step, ProTune runs them simultaneously. By default ProTune distributes the scripts evenly between the Vusers. For example, if you assign two scripts to one step, fifty percent of the Vusers run one script, while the remaining fifty percent run the second script. To modify the script distribution, click in the % column, and modify the percentages accordingly. Note that the total sum or the percentages must equal 100.

**To Add a Load Generator:**

- 1 The Load Generators column automatically contains <All Load Generators> for each script. You can assign specific load generators for each script. Click in the **Load Generators** column in the next available row.
- 2 Click on the arrow to the right of the box. A list box opens showing the available load generator machines.
- 3 Select one or more machines and click **OK**.
- 4 Click **All Generators** to instruct ProTune to run the script on all available machines.
- 5 To add a new load generator, click the first entry in the list box, **Add**. The Add Load Generator dialog box opens:



- 6 Type the name of the load generator in the **Name** box.
- 7 In the **Platform** box, select the type of platform on which the load generator is running.
- 8 By default, ProTune stores temporary files on the load generator during session execution, in a temporary directory specified by the load generator's TEMP or TMP environment variables. To override this default for a specific load generator, type a location in the **Temporary Directory** box.
- 9 To allow the load generator to take part in the session, check **Enable load generator to take part in the session**.

- 10** Click **OK** to close the Add Load Generator dialog box. ProTune adds the new load generator to the Load Generator Name list.

To include the new load generator in your session, select it from the Load Generator Name list, and click **OK**. Note that you can select multiple load generators.

Repeat the above procedure for each load generator you want to add to your session step.

For more information about setting up load generators, see Chapter 4, “Managing Load Generators.”

**To configure a load generator:**

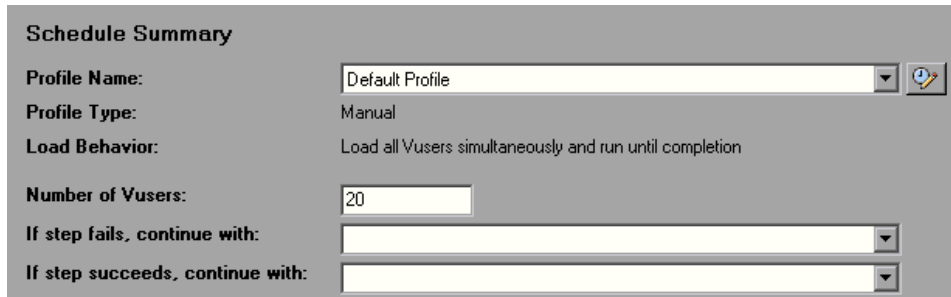
- Use the Load Generators dialog box to set a load generator’s attributes while adding it to the load generator list, or to modify the attributes of an existing load generator at any time.
- You can also use the Load Generators dialog box to indicate which load generators will run Vusers in the session step. For example, if a load generator is unavailable for a particular session step, you can use the Load Generators dialog box to exclude it temporarily instead of removing it entirely from your list of load generators. For instructions on using the Load Generators dialog box, see “Configuring Load Generators” on page 57. To configure additional load generator settings, see “Configuring Load Generator Settings” on page 61.
- To configure global settings for all load generators participating in the session, use ProTune’s Options dialog box. For more information, see Chapter 5, “Configuring Session Steps.”

## Setting an Initial Load (Manual Profiles)

If the profile associated with a step is a manual one, you can specify the number of Vusers to run on your system.

**To define an initial load:**

- 1 Select a session step in the left pane. The Schedule Summary section shows the following fields:
  - Profile Name
  - Profile Type (Manual or Goal-Oriented)
  - Load behavior
  - Number of Vusers



The screenshot shows a configuration window titled "Schedule Summary" with a grey background. It contains several fields for configuring a manual profile:

- Profile Name:** A dropdown menu showing "Default Profile" with a small icon to its right.
- Profile Type:** A dropdown menu showing "Manual".
- Load Behavior:** A text field containing "Load all Vusers simultaneously and run until completion".
- Number of Vusers:** A text input field containing the number "20".
- If step fails, continue with:** A dropdown menu.
- If step succeeds, continue with:** A dropdown menu.

- 2 In the Number of Vusers section, enter the number of Vusers to run during the tuning session.

Note that in the **Execute** tab you can modify this value during the actual tuning. The value in this tab is useful for automatic scheduling—ProTune runs the session steps at the designated times with the number of Vusers specified in the **Total Number of Vusers to Run in Step** box.

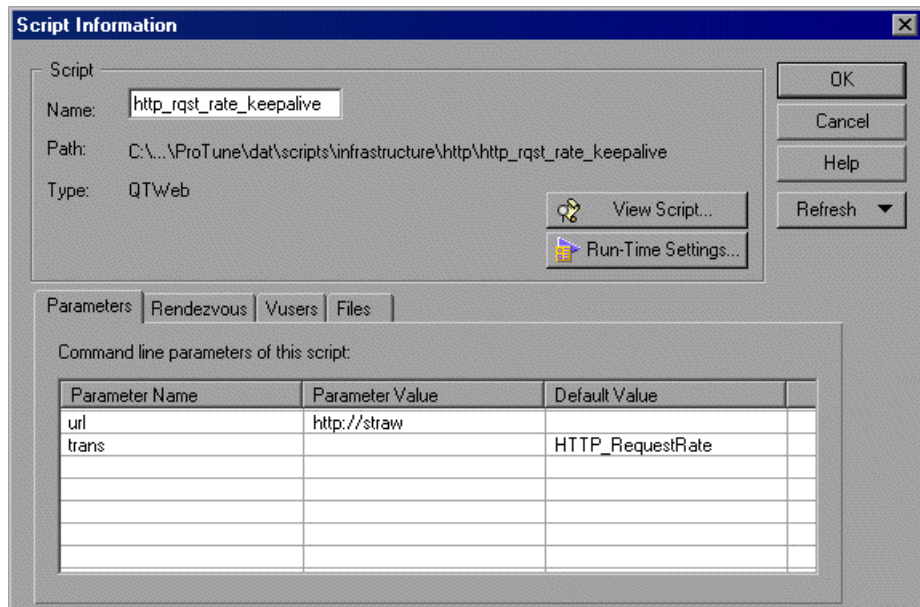
## Configuring Script Details

For scripts that are displayed in your list, you can view the details of the script you selected, edit the script, or change its run-time settings.

**To view script details:**



- 1 Select a script and click the **Script Details** button on the toolbar or double-click on the script in the left pane. The Script Information dialog box opens, displaying the Path, Name, and Type of the selected script. It also displays the script's command line options (for example -url Application Server1), and tabs displaying information about parameters, rendezvous points, Vusers and files.



- 2 Click **Run-Time Settings** to set the script's run-time settings (optional), which allow you to customize the way the Console executes a script. The Run-Time Settings dialog box opens, displaying the settings you previously set using VuGen. If you did not set run-time settings for a script in VuGen, the default VuGen settings are displayed for all but the Log and Think Time tabs, which display the default Console settings. Note that several protocols, such as Web and Java, have specific settings.

For information on configuring the run-time settings, refer to the *ProTune Creating Virtual User Scripts* guide.

---

**Note:** If you modify the run-time settings from the Console, ProTune runs the script using the modified settings. To restore the initial settings, click the **Refresh** button and select **Run-Time Settings**.

---

- 3** To edit the script, click **View Script**. The script generation tool, VuGen, opens. For more information on editing scripts, refer to the *ProTune Creating Virtual User Scripts* guide.
- 

**Note:** If you use VuGen to make changes to a script while the Console is running, click the **Refresh** button and select **Script** to update the script details in the session step.

---

- 4** In the Command Line box, type any command line options to use when running the script. For example: -x value -y value  
For information about passing command line argument values to a script, refer to the *ProTune Creating Virtual User Scripts* guide.
  - 5** Click the **Parameters** tab to view the arguments for the selected script and their default values. To specify your own value, click in the **Parameter Value** column and enter the desired value.
- 

**Tip:** To view information about a parameter, hold your cursor over the parameter name and view the tooltip.

---

- 6** To see the rendezvous points included in the selected script, click the **Rendezvous** tab.
- 7** To see the list of Vusers associated with the selected script, click the **Vusers** tab. If you have not yet created Vusers, the box will be empty.

- 8 To see the list of files used by the script, click the **Files** tab. By default this list shows all files in the script's directory (only after your script has been added to the script list). These files include the configuration settings file, the init, run, and end portions of the script, the parameterization definitions file, and the *usr* file. To add a file to the list, click **Add** and add the file name. Note that you can delete the files that you add, but not the other files listed.
- 9 Click **Refresh > Script** to refresh the script and use the default settings in the tabs (File tab). Click **Refresh > Runtime settings** to use the script's default run-time settings.

---

**Note:** Refreshing the Log run-time settings does not override the **Send messages only when an error occurs** option, even if the original setting was **Always send messages**. Refreshing the Think Time run-time settings does not override the **Replay think time** option, even if the original setting was **Ignore think time**.

---

- 10 Click **OK** to close the Script Information dialog box.

After you add steps and scripts to your session, you can configure run-time options and set a schedule. Refer to “Getting Started with Designing a Session,” on page 30 for a summary of the testing procedure. After you set up your session steps, you run the steps and begin your tuning session. For information on running the tuning session, see Chapter 9, “Running a Session.”

## Using Relative Paths for Scripts

To specify the location of a script, you can either browse to the script or type its relative location into the Script Path column. The location can be relative to the current session directory, or the ProTune installation directory.

You can specify a path relative to the current session directory by typing either of the following notations at the start of the script path:

- .\                                indicates that the path is relative to the location of the session directory.
- ..\                                indicates that the path is relative to the location of the parent directory of the session directory.

For example, if the current session is located at F:\sessions, to specify a script located at F:\sessions\scripts\user1.usr, you could type the following:

```
.\scripts\user1.usr
```

You can specify a path relative to the ProTune installation directory by typing a percent sign (%) at the beginning of the script path. For example, if the ProTune installation directory is located at F:\ProTune, to specify a script located at F:\ProTune\scripts\user1.usr, you could type the following:

```
%\scripts\user1.usr
```

---

**Note:** When specifying a relative path, you can include standard DOS notation (.\ and ..\) inside the path, as shown in the following example: M:\LRALT\my\_tests\..\test.usr.

---

When you run a session, by default, the script is copied to a temporary directory on the Vuser group machine. This enables the Vuser group load generator to access the script locally instead of over a network.



You can instruct the Console to store the script on a shared network drive (see Chapter 5, “Configuring Session Steps.”) If you configure the Console to save the script to a network drive, you must ensure that the Vuser load generator recognizes the drive. The Script window contains a list of all the scripts and their paths. A script’s path is based on the Console load generator’s mapping of that location. If a Vuser load generator maps to the script’s path differently, path translation is required. Path translation converts the Console load generator’s mapping to the Vuser load generator’s mapping. For more information see Appendix C, “Performing Path Translation.”



# 4

---

## Managing Load Generators

After choosing the tests you want to run, and the components that you want to test, you need to specify the computers from which ProTune will run the tests. These computers are called load generators. This chapter describes how to define and manage load generators.

This chapter discusses:

- Configuring Load Generators
- Configuring Load Generator Settings

### About Managing Load Generators

You use load generator machines for running the tests on your components. When you assign a script to a step, you also specify the load generators that will run the script, and can configure their properties.

### Configuring Load Generators

You can set a load generator's attributes while adding it to the load generator list, or modify the attributes of an existing load generator at any time, using the Load Generators dialog box.

To configure global settings for all load generators participating in the session, use the **Timeout** tab in the **Tools > Options** dialog box. For more information, see Chapter 5, "Configuring Session Steps." To set properties specific to each load generator, use the Load Generators dialog box as described below.


You can also indicate which load generators will run Vusers in the session. For example, if a load generator is unavailable for a particular session run, you can exclude it temporarily instead of removing it entirely from your list of load generators.

You select which load generators will take part in the session by using the Enable and Disable commands. Disabling a load generator temporarily removes it from the list. Enabling a load generator reinstates it. Disabling load generators is particularly useful if you want to isolate a specific machine to test its performance.

### To configure a load generator:




- 1 Click the **Generators** button, or select **Session > Load Generators**. The Load Generators dialog box opens. The **Name** of the load generator, its **Status**, **Platform**, and **Details** are displayed.



- 2 Click **Connect** to change the Status of the load generator from DOWN to READY. Click **Disconnect** to change the Status of the load generator from READY to DOWN.
- 3 To disable a load generator, select the load generator and click **Disable**. The load generator name changes from blue to gray, and the load generator is disabled. To enable a load generator, select the load generator and click **Enable**. The load generator name changes from gray to blue, and the load generator is enabled.

- 4 To view details of a load generator, select the load generator and click **Details**. The Load Generator Information dialog box opens with information about the load generator you selected.
- 5 To add a load generator, or modify information for an existing load generator, click **Add**. The Add Load Generator dialog box opens.



Type the name of the load generator in the **Name** box. In the Platform box, select the type of platform on which the load generator is running.

- 6 By default, ProTune stores temporary files on the load generator during session execution, in a temporary directory specified by the load generator's TEMP or TMP environment variables. To override this default for a specific load generator, type a location in the Temporary Directory box.
- 7 To allow the load generator to take part in the session, check **Enable load generator to take part in the session**.
- 8 Click **More** to expand the dialog box and show the following additional tabs where you can configure load generator settings:
  - **Status**
  - **Run-Time Quota**
  - **Firewall**
  - **Run-Time File Storage**
  - **Unix Environment**
  - **Vuser Limits**

**Note:** For information on configuring these settings, see “Configuring Load Generator Settings” on page 61.

---

- 9** Click **OK** to close the Add Load Generator dialog box.
- 10** To remove a load generator, select it and click **Delete**.
- 11** Click **Close** to close the Load Generators dialog box. The load generator name you entered appears in the Load Generators list; its status is set to Down.

---

**Note:** The ProTune Console monitors a Windows load generator machine's CPU usage and automatically stops loading Vusers on a load generator when it becomes overloaded. You can monitor the status of a machine's CPU usage using the icons in the Load Generators dialog box. When the CPU usage of a load generator becomes problematic, the icon to the left of the load generator name contains a yellow bar. When the machine becomes overloaded, the icon contains a red bar.

---

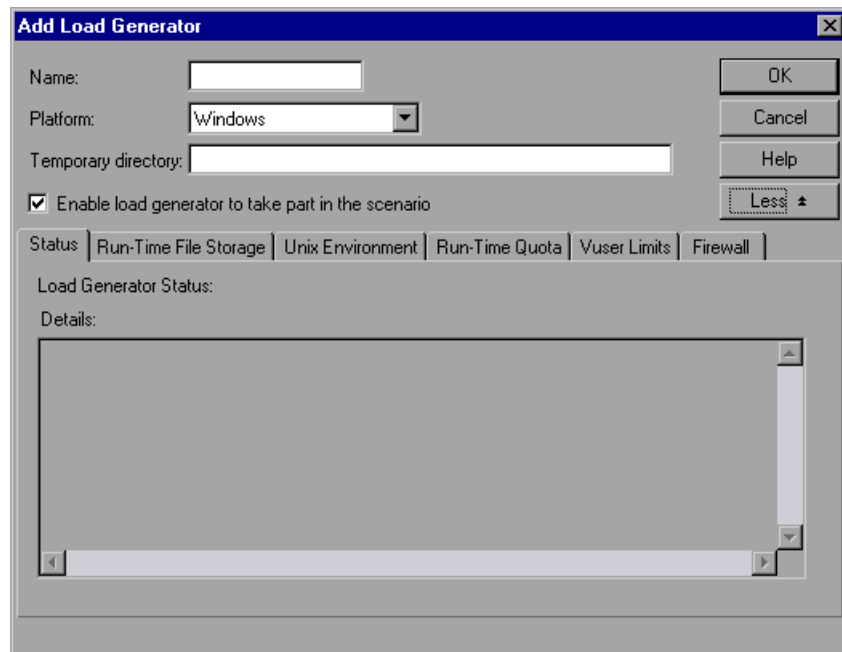
## Configuring Load Generator Settings

You can configure additional settings for individual load generators using the tabs in the Add Load Generator or Load Generator Information dialog boxes. The settings that can be configured are: Run-Time File Storage, UNIX Environment, Run-Time Quota, Vuser Limits, Connection Log (Expert mode), and Firewall.

You can configure global settings for all load generators participating in the session, using the Options dialog box. For more information, see Chapter 5, “Configuring Session Steps.”

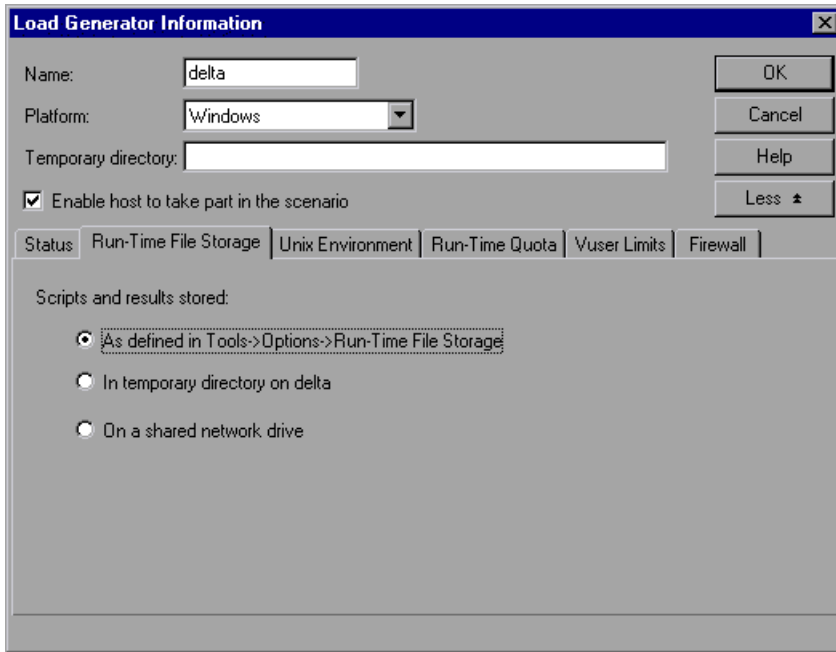
### To configure load generator settings:

- 1 From the Add Load Generator or Load Generator Information dialog box, click **More** to expand the box and show the Status, Run-Time File Storage, UNIX Environment, Run-Time Quota, Vuser Limits, and Firewall (when the load generator is not the localhost) tabs.



- 2 Select the **Status** tab to display the Load Generator’s status.

- 3 Select the **Run-Time File Storage** tab to specify the result directory for the performance data that ProTune gathers from each load generator during a session.



To store the results as specified in the global settings, click **As defined in Tools > Options > Run-Time File Storage**. To store the results temporarily on a hard drive of the load generator computer, click **In temporary directory on <load generator name>**. To store the session scripts or results on a shared network drive, click **On a shared network drive**. To set the network location for the results, see Chapter 8, “Preparing to Run a Session Step.”

---

**Note:** If the load generator is *localhost*, ProTune stores the scripts and results on a shared network drive, and the checkboxes and radio buttons for setting the location are all disabled.

---



- 4 Select the **UNIX Environment** tab to configure the login parameters and shell type for each UNIX load generator.

The screenshot shows the 'Load Generator Information' dialog box with the following settings:

- Name:** localhost
- Platform:** UNIX
- Temporary directory:** (empty)
- Enable load generator to take part in the scenario
- Tab:** Unix Environment
- Login as:**
  - Name: (empty)
  - Use lower case for login names
- Shell Settings:**
  - Default shell:** csh
  - Initialization command:** (empty)
- Buttons:** OK, Cancel, Help, Defaults

To specify a login name other than the current Windows user, select the **Name** check box and specify the desired UNIX login name. To login with lower case characters, select the **Use lower case for login names** check box.

---

**Note:** For information on the Local User setting available in Expert mode, see “Working in Expert Mode” on page 645.

---

From the Default Shell box, select **csh** (C Shell—the default), **bsh** (Bourne Shell), or **ksh** (Korn Shell).

To allow ProTune to run your application under the Korn shell, you first need to make sure that the *.profile* file contains all of the ProTune environment settings—for example, the `M_LROOT` definition and the `LicenseManager` variable. These environment settings already exist in your *.cshrc* file. Your UNIX `$M_LROOT/templates` directory contains a template for the *.profile* file, called *dot profile*. Use the template as a guide for modifying your *.profile* file with the ProTune environment settings.

---

**Note:** If you are using a Korn shell (ksh), you must delete all ProTune settings from the *.cshrc* file (e.g. `M_LROOT`) before executing the session.

---

In the Initialization Command box, enter any command line options that ProTune will use when logging on to a UNIX system. This initialization command will run as soon as the shell opens.

For example, you could select *ksh* and use the following initialization command:

```
. .profile;
```

- 5 Select the **Run-Time Quota** tab to specify the maximum number of Vuser types that the load generator will initialize or stop simultaneously.

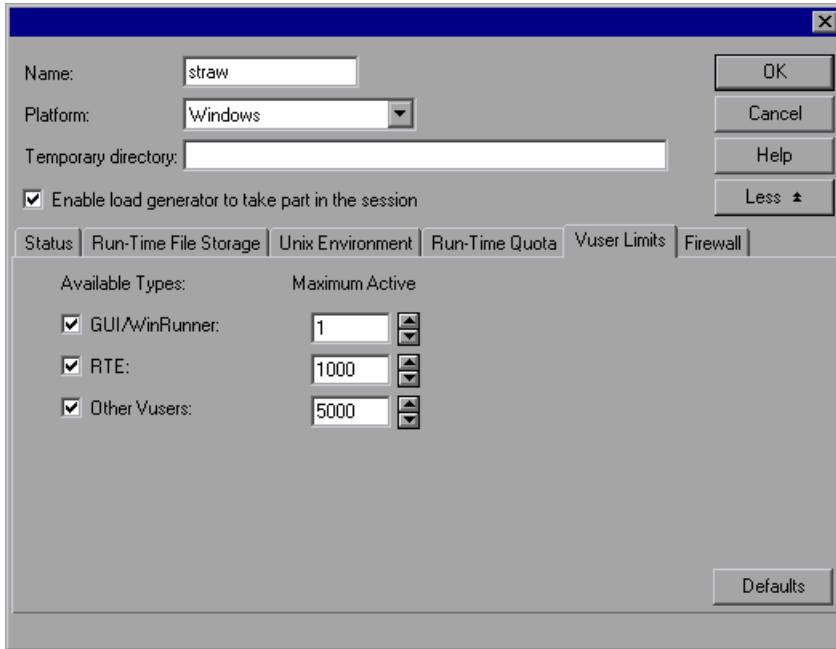
The screenshot shows the 'Load Generator Information' dialog box with the 'Run-Time Quota' tab selected. The dialog has a title bar with a close button. Below the title bar are fields for 'Name' (delta), 'Platform' (Windows), and 'Temporary directory'. There are 'OK', 'Cancel', and 'Help' buttons on the right. A checkbox labeled 'Enable load generator to take part in the scenario' is checked. Below these are tabs for 'Status', 'Run-Time File Storage', 'Unix Environment', 'Run-Time Quota', 'Vuser Limits', and 'Firewall'. The 'Run-Time Quota' tab is active, showing a 'Vuser Quota' section with two spinners, both set to 50. The first spinner is labeled 'Number of Vusers that may be initialized at one time - delta' with a note '( Total number of vusers that can be initialized at one time - 999 )'. The second spinner is labeled 'Limit the number of users that may be stopped at one time to:'. A 'Defaults' button is at the bottom right. A 'NOTE' at the bottom left reads: 'NOTE: to change the settings for the total number go to Tools->Options->Run-Time Settings'.

Click **Defaults** to use the Default values.

Initializing or stopping a large number of Vusers simultaneously places large stress on a load generator. To reduce stress on a load generator, you can initialize or stop smaller batches of Vusers.

You can set run-time quotas for an entire session using the Run-Time Settings tab in the Options dialog box. For information on setting quotas globally for an entire session, see Chapter 5, “Configuring Session Steps.”

- 6 Select the **Vuser Limits** tab to modify the maximum number of GUI, RTE, and other Vusers that a load generator can run.



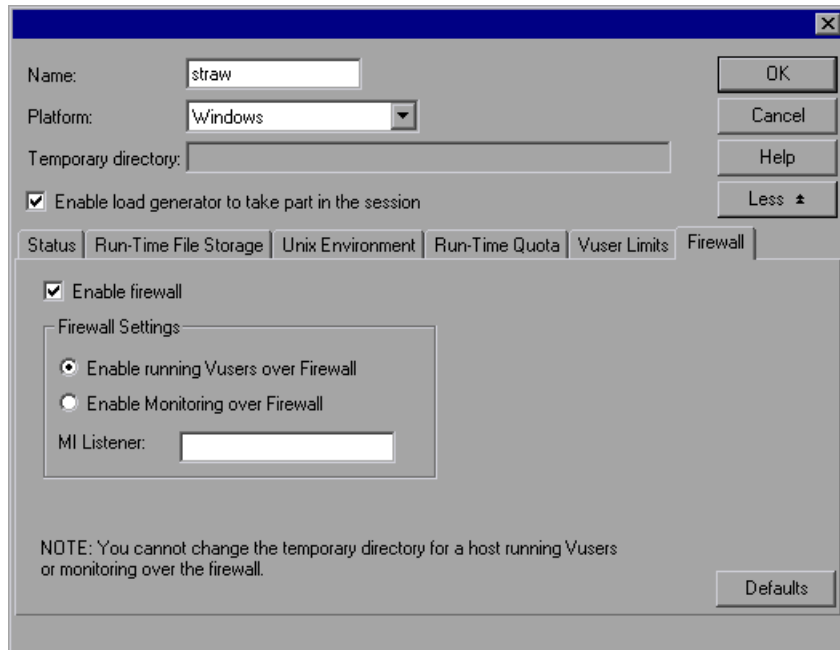
In the Maximum Active boxes, enter the maximum number of Vusers of each type that the load generator can run.

---

**Note:** The maximum number of active Vusers that you specify must not exceed the number of Vusers that you are licensed to run. To check your Vuser licensing limitations, choose **Help > About ProTune**.

---

- 7 Select the **Firewall** tab to enable monitoring or running Vusers through a firewall.



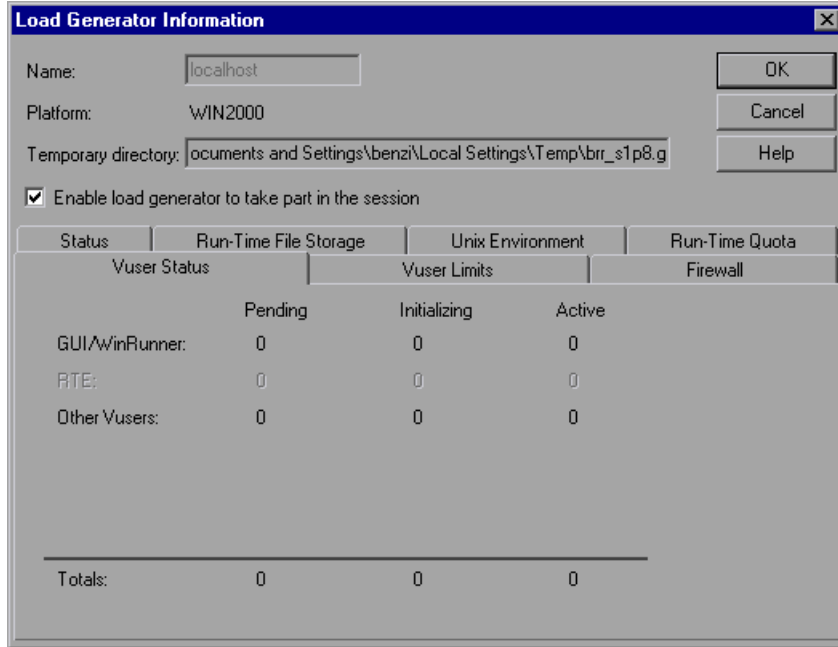
You can enable or disable ProTune's firewall functionality, specify the name of the MI Listener the load generator is using, and click either the **Enable Monitoring over Firewall** radio button, or the **Enable running Vusers over Firewall** radio button.

---

**Note:** If the load generator is connected, you cannot change values in the **Firewall** tab, or if you change the name of the load generator to the name of the local host, you cannot set any values in the **Firewall** tab.

---

- 8 If the load generator machine is connected, you can view the **Vuser Status** tab, which displays the number of GUI/WinRunner, RTE, and other Vusers that are *Pending*, *Initializing*, and *Active* on the selected load generator machine.




---

**Note:** For information on the Connection Log tab available in Expert mode, see “Working in Expert Mode” on page 645.

---

- 9 Click **OK** to close the Add Load Generator or Load Generator Information dialog box and save your settings.

# 5

---

## Configuring Session Steps

You can configure how load generators and Vusers behave when you run a session so that the session accurately emulates your working environment.

This chapter describes:

- ▶ Configuring Session Run-Time Settings
- ▶ Setting Timeout Intervals
- ▶ Setting the Run-Time File Location
- ▶ Specifying Path Translation

### About Configuring a Session

Before you run a session, you can configure both the load generator and Vuser behaviors for the session. Although the default settings correspond to most environments, ProTune allows you to modify the settings in order to customize the session behavior. The settings apply to all future session runs and generally only need to be set once.

The settings described in this chapter apply to all the load generators in a session. To change the settings for individual load generator machines, refer to Chapter 2, “Creating a Topology.” If the global session settings differ from those of an individual load generator, the load generator settings override them.

The settings discussed in this chapter are unrelated to the Vuser run-time settings. These settings, which apply to individual Vusers or scripts, contain information about logging, think time, and the network, the number of iterations, and the browser. For information on setting the run-time settings, refer to the *ProTune Creating Virtual User Scripts* guide.

For information on setting the options for online monitors, see Chapter 14, “Online Monitoring.”

The ProTune Expert mode allows you to configure additional settings for the ProTune agent and other ProTune components. For more information, see Appendix B, “Working in Expert Mode.”

## Configuring Session Run-Time Settings

The session run-time settings relate to:

- ▶ Vuser Quotas
- ▶ Stopping Vusers
- ▶ Random Sequence Seed

**Vuser quotas:** To prevent your system from overloading, you can set quotas for Vuser activity. The Vuser quotas apply to Vusers on all load generators. You can limit the number of Vusers initialized at one time (when you send an Initialize command).

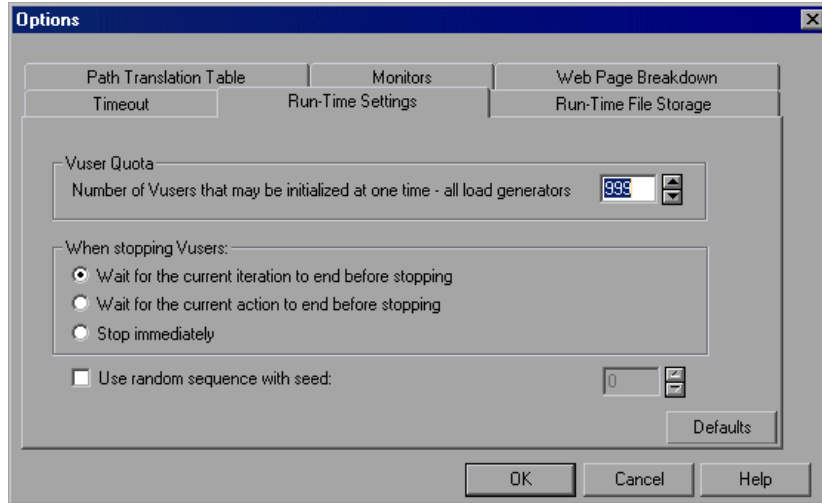
**Stopping Vusers:** ProTune lets you control the way in which Vusers stop running when you click the Stop button. You can instruct ProTune to allow a Vuser to complete the iteration it is running before stopping, to complete the action it is running before stopping, or to stop running immediately.

**Random sequence seed:** ProTune lets you set a seed number for random sequencing. Each seed value represents one sequence of random values used for test execution. Whenever you use this seed value, the same sequence of values is assigned to the Vusers in the session. This setting applies to parameterized scripts using the Random method for assigning values from a data file. Enable this option if you discover a problem in the test execution and want to repeat the test using the same sequence of random values.



### To set the session run-time settings:

- 1 Choose **Tools > Options**. The Options dialog box opens. Click the **Run-Time Settings** tab.



- 2 To set a Vuser quota, specify the desired value.
- 3 Select the way in which you want ProTune to stop running Vusers.
- 4 To specify a seed value for a random sequence, select the **Use random sequence with seed** check box and enter the desired seed value.

## Setting Timeout Intervals

ProTune enables you to set the timeout interval for commands and Vuser elapsed time.

The command timeouts are the maximum time limits for various ProTune commands. When a command is issued by the Console, you set a maximum time for the load generator or Vuser to execute the command. If it does not complete the command within the timeout interval, the Console issues an error message.

The command timeouts relate to load generators and Vusers. The load generator commands for which you can specify a timeout interval are Connect and Disconnect. The Vuser commands for which you can specify a timeout interval are *Init*, *Run*, *Pause*, and *Stop*.

For example, the default *Init* timeout is 180 seconds. If you select a Vuser and click the **Initialize** button, ProTune checks whether the Vuser reaches the READY state within 180 seconds; if it does not, the Console issues a message indicating that the *Init* command timed out.

In the Vuser view, the *Elapsed* column (the last column) indicates the amount of time that elapsed from the beginning of the session. You can specify the frequency in which ProTune updates this value. The default is 4 seconds.

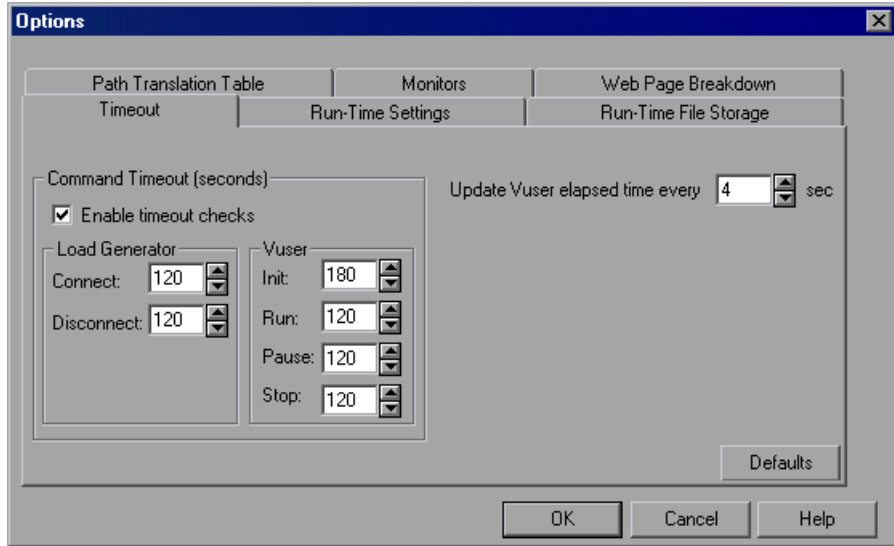
---

**Note:** ProTune's calculations consider the number of active Vusers and their influence on the timeout values. For example, 1000 Vusers trying to initialize will take much longer than 10 Vusers. ProTune adds an internal value, based on the number of active Vusers, to the specified timeout value.

---

**To set timeout intervals:**

- 1 Choose **Tools > Options**. The Options dialog box opens. Click the **Timeout** tab.



- 2 Clear the **Enable timeout checks** check box to disable the timeout test. ProTune waits an unlimited time for the load generators to connect and disconnect, and for the Initialize, Run, Pause, and Stop commands to be executed.
- 3 To specify a command timeout interval, select the **Enable timeout checks** check box and specify the appropriate timeouts.
- 4 Specify the frequency at which ProTune updates the Elapsed time, in the **Update Vuser elapsed time every** box.

## Setting the Run-Time File Location

When you run a session, by default the run-time files are stored locally on each Vuser load generator (the machine running the script). The default location of the files is under the temporary directory specified by the load generator's environment variables (on Windows, TEMP or TMP and on UNIX, \$TMPDIR or \$TMP). If no environment variable is defined, the files are saved to the /tmp directory.

---

**Note:** The run-time file storage settings that are described in this apply to all the load generators in a session. You can change the settings for individual load generator machines as described in “Configuring Load Generators” on page 57.

---

The primary run-time files are script and result files:

*Script files:*

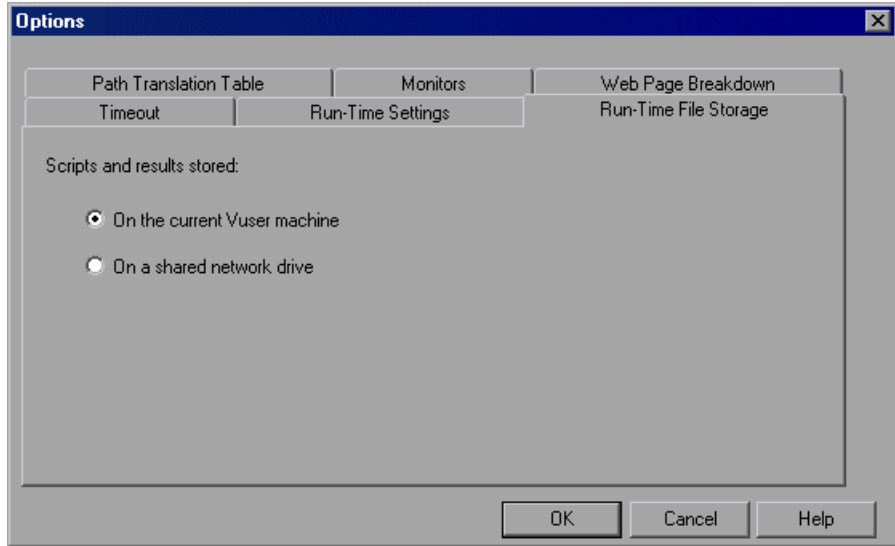
When you run a Vuser, the Console sends a copy of the associated script to the Vuser load generator. The script is stored in the load generator's temporary run-time directory.

*Result files:*

While you run a session, the participating Vusers write their results to the temporary run-time file directory. After session execution, these result files are collated or consolidated—results from all of the load generators are transferred to the results directory. You set the location of the results directory as described in Chapter 9, “Running a Session.” After collating the results, the temporary run-time directory is deleted.

To specify where ProTune stores run-time files:

- 1 Choose **Tools > Options**. The Options dialog box opens. Click the **Run-Time File Storage** tab.



By default, the **On the current Vuser machine** option is selected. This means that all run-time files—including result files and script files—are stored on the Vuser load generators. The only exception is for Vusers running on the local load generator (Console machine), where you must use the shared drive option.

- 2 To store script and result files on a shared network drive, click **On a shared network drive**. To set the exact location on the network drive, see Chapter 8, “Preparing to Run a Session Step.”

If you select to save results to a shared network drive, you may need to perform path translation. Path translation ensures that the specified results directory is recognized by the remote load generator. For information about path translation see Appendix C, “Performing Path Translation.”

If you specify that all Vusers access their scripts directly at some shared location, no transfer of script files occurs at run time. This alternative method may be useful in either of the following situations:

- The file transfer facility does not work.
- The script files are large and therefore take a long time to transfer. Remember that script files are transferred only once during a session.

This alternate method often necessitates path translation. For details, see Appendix C, “Performing Path Translation.”

- 3 Click **OK** to close the dialog box.

---

**Note:** If you choose to save result files on the Vuser load generators, you must collate the results before you can perform any analysis. You can wait for ProTune to collate the results when you launch the Analysis tool, or you can collate results by selecting **Results > Collate Results**. Alternatively, select **Results > Auto Collate Results** to automatically collate the results at the end of each session run.

---

## Specifying Path Translation

If you specified a shared network drive for run-time file storage, (see “Setting the Run-Time File Location” on page 74), you may need to perform *path translation*. Path translation is a mechanism used by ProTune to convert a remote path names. A typical session may contain several load generator machines that map the shared network drive differently. For more information, see the Appendix C, “Performing Path Translation.”

# 6

---

## Defining Alerts

Before tuning your session, you set up alerts that define what actions ProTune should take when server performance problems occur.

This chapter describes:

- ▶ Types of Alerts
- ▶ Specifying Alert Conditions
- ▶ Specifying Alert Actions
- ▶ Viewing Alert Descriptions
- ▶ Creating, Configuring, and Deleting Alerts
- ▶ Enabling and Disabling the Alert Mechanism
- ▶ Enabling and Disabling the Alert Mechanism
- ▶ Viewing Alerts in the Output Window

### About Defining Alert Schemes

ProTune uses alerts to let you know when server performance problems occur.

Before running scripts, you define alerts for one or more measurements. This includes specifying what conditions should trigger an alert, and the action that ProTune should take when the need to issue an alert is detected.

You use the Alerts window to specify alert conditions and an alert scheme. You specify separate alert conditions for each session step.

**Note:** Although you select measurements to monitor topology elements (not physical components), you assign alerts to physical machines. This means that if an alert is triggered on a measurement of one topology element, it affects all the other elements mapped to the same physical machine. For example, if the alert action is to stop Vusers or stop the ramp-up, the action will affect all the Vusers, regardless of the load generators on which they run and the machines against which they are running.

---

## Types of Alerts

You can create alerts that will be triggered by specific values occurring in the following measurements:

---

**Note:** The following is a partial list of the measurements that you can use for triggering alerts. All measurements that are available in the Console can be used.

---

► **Running Vusers**

Informs you when the number of Running Vusers in the Running, Ready, Finished, or Error state, reaches a specific value.

► **Error Statistics**

Informs you when the number of errors reaches a specified number.

► **Vusers with Errors**

Informs you when the number of Vusers with errors reaches a specified number.

► **Transaction Response Time (Passed)**

Informs you when the transaction response time reaches a specified value, for the *vuser\_int*, *Actions*, or *vuser\_end* sections of the script.



➤ **Transactions Per Second (Passed)**

Informs you when the specified number of transactions per second is reached for the *vuser\_int*, *Actions*, or *vuser\_end* sections of the script.

➤ **Total Transactions Per Second (Passed)**

Informs you when the specified number of transactions per second is reached for the *vuser\_int*, *Actions*, or *vuser\_end* sections of the script.

➤ **Hits per Second**

Informs you when the specified number of hits per second is reached.

➤ **Throughput**

Informs you when your server's throughput reaches a specific level.

➤ **Pages Downloaded per Second (Passed)**

Informs you when a specific number of pages has been downloaded per second by the server.

## Specifying Alert Conditions

ProTune allows you to specify when you want it to issue alerts for the measurement that you are monitoring. ProTune provides several types of alert schemes. The alert schemes instruct ProTune to issue an alert when it detects a specific value, a value for a specific duration, an out of range value, a change in value, or a standardized value.

| Scheme                       | Additional fields           | Meaning   |
|------------------------------|-----------------------------|---|
| value                        | none                        | Issue an alert if the measurement compares with the specified value in one of the following ways:<br>><br>>=<br><<br><=<br>=                          |
| value for a duration of time | for a period of $n$ seconds | Issue an alert if the selected measurement compares with the specified value in one of the following ways for $n$ seconds:<br>><br>>=<br><<br><=<br>= |
| value out of range           | range                       | Issue an alert if the measurement deviates from the specified range.  |

| Scheme             | Additional fields           | Meaning  |
|--------------------|-----------------------------|--|
| value change       | for a period of $n$ seconds | Issue an alert if the change in the selected measurement compares with the specified value in one of the following ways over the last $n$ seconds:<br>><br>>=<br><<br><=<br>=                                    |
| standardized value | for a period of $n$ seconds | Issue an alert if the standardized value of the selected measurement compares with the specified value in one of the following ways, considering the period of the last $n$ seconds :<br>><br>>=<br><<br><=<br>= |

You can choose one of the following conditions for your alert scheme:

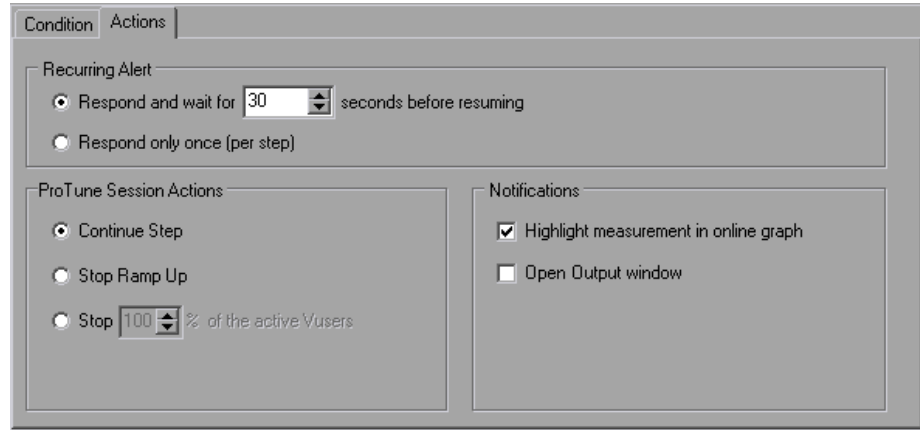
- > greater than
- >= greater than or equal to
- < less than
- <= less than or equal to
- = equal to

You define an alert by selecting a measurement, scheme, and condition in the Alerts Window's **Condition** tab. In the following example, ProTune is instructed to issue an alert when it detects a change of 10 or more in the value of the *% Processor Time* measurement, over a period of 20 seconds.

The screenshot shows the 'Condition' tab of the Alerts Window. At the top, there are two tabs: 'Condition' and 'Actions'. Below the tabs, the 'Measurement' field is set to '% Processor Time (Processor\_Total) (localhost)'. Under the heading 'Send alert under the following condition:', there are three dropdown menus: the first is set to 'value change', the second to '>=', and the third to '10'. Below these, there is a spinner box set to '20' followed by the text 'seconds'. A 'Condition Description' box contains the text: 'Issue an alert if the change of the selected measurement is greater than or equal to the specified value for the specified period.' At the bottom, the 'Alert Description (click on the underlined value to edit it)' box contains the following text: 'Alert [cpu](#) will be issued if the [value change](#) of measurement [% Processor Time \(Processor\\_Total\) \(localhost\)](#) in the last [20 seconds](#) [is greater than or equal to](#) [10](#). If the condition is satisfied, issue an alert (wait for [5](#) seconds before resuming) and [highlight measurement in online graph](#)'.

## Specifying Alert Actions

The Alerts Window's Actions tab allows you to specify the actions ProTune takes when an alert is triggered.



You can specify settings for the following:

- Recurring Alerts
- Session Actions
- Notification Types

### Recurring Alerts

The first notification setting relates to the frequency of responses for a recurring alert in a single step. You can instruct ProTune to respond in one of the following ways:

- **Respond and Wait *n* seconds before resuming**
- **Respond only once** (per ProTune step)

Default: respond and wait 30 seconds before resuming the test.

## Session Actions

ProTune lets you specify what action to perform when encountering an alert condition:

- **Continue Step:** Continue executing the current step after the alert is triggered.
- **Stop Ramp Up:** Stop adding additional Vusers to the session step.
- **Stop a percentage of the active Vusers:** When an alert is triggered, stop a percentage of the active Vusers. You can set a percentage from 0 to 100.

## Notification Types

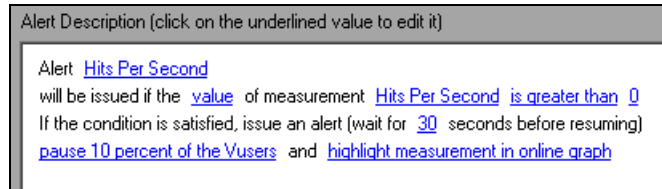
You can specify the type of notification to issue when an alert is triggered. The available notification types are:

- **Highlight measurement in online graph:** On the online graph, highlight the measurement that triggered the alert. This graph can be viewed on the Execute tab.
- **Open Output window:** Open the Output window to display alert messages when an alert trigger occurs. For more information about the viewing alert messages in the Output Window, see “Viewing Alerts in the Output Window,” on page 92.

You can enable either notification mechanism, both of them, or neither. Even if you do not enable a notification mechanism, you can still open the Alerts Output window to view the alerts that occurred.

## Viewing Alert Descriptions

The Alert Description pane contains a description of the alert in plain language. It includes hyperlinks that help you to change the alert's condition and actions.



When you click a hyperlink, ProTune positions the cursor at the field that you want to change. In some cases, clicking a hyperlink causes ProTune to open a dialog box so you can change the relevant value.

## Creating, Configuring, and Deleting Alerts


After designing the session steps, you set alerts for your tuning session. The alerts provide real-time information about your server performance and inform you when specific thresholds are reached.

After creating an alert, you specify the alert conditions and actions.


### Creating an Alert

Creating an alert includes specifying the alert name and choosing the measurement that will trigger the alert.


**To create an alert:**

-  1 Click the **Alerts Definition** button in the ProTune Console window. The Alerts window opens.

- 2 Click **Add** in the bottom left corner. The Online dialog box opens.



- 3 Enter a name for the alert and click **OK**. The Select Measurement dialog box opens, displaying categories of measurements that you can monitor.



The following categories always appear, regardless of your topology and monitors:


- Running Vusers
- Total Transactions per Second (Passed)
- Hits per Second
- Throughput
- Pages Downloaded per Second
- Connections



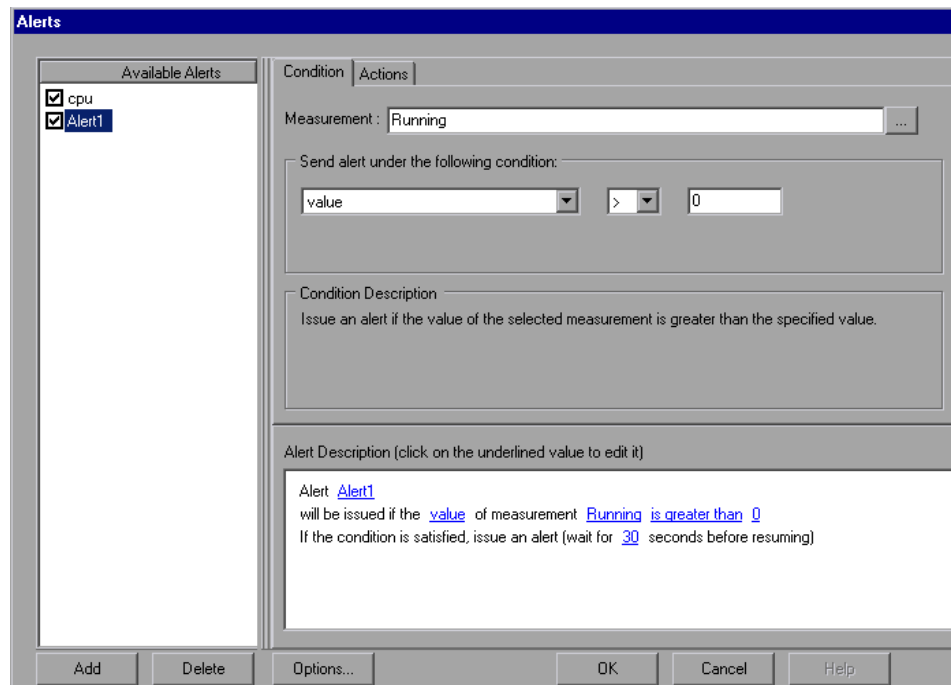
➤ Connections per second

Additional categories may appear, depending on your topology and the monitors you assigned to various elements.

4 Expand the required category to display its measurements.



5 Select the desired measurement and click **OK**. ProTune creates an alert with default settings. The alert name you supplied appears in the Available Alerts pane, the Condition and Actions tabs show the default values, and the Alert Description pane displays a plain language description of the alert.



6 To delete an alert, select it in the Available Alerts pane and click **Delete**.

## Specifying Alert Conditions

You can change the existing settings to suit your needs.

To specify alert conditions:

- 1** In the Condition tab, choose an alert scheme in the **Send alert under the following condition** box.
- 2** Choose an operator: >, >=, <, <=, or =.
- 3** Specify a condition:
  - For the **value** scheme, specify a trigger value.
  - For the **value for a duration of time** scheme, specify a trigger value and a duration.
  - For the **value out of range** scheme, specify start and end range trigger values.
  - For the **value change** scheme, specify a trigger value and a watch time—the alert is only triggered if the condition is reached within the watch time.

---

**Note:** If you want ProTune to issue an alert when the value decreases by more than 5, specify the "<" operator and a trigger value of -5.

---

- For the **standardized value** scheme, specify a trigger value and a watch time—the alert is only triggered if the condition is reached within the watch time.

---

**Note:** You can use the Alert Description pane to change alert condition settings (see “Viewing Alert Descriptions,” on page 85).

---

## Specifying Alert Actions

ProTune also assigns default settings for the actions that are triggered by an alert. ProTune displays these settings in the Actions tab. You can change the default settings to suit your needs.

The Actions tab lets you specify the following:

- Recurring alert behavior
- ProTune session actions
- Notifications

### To specify alert actions:

- 1** In the **Recurring Alert** section, select one of the following response schemes:
  - **Respond and wait for  $n$  seconds before resuming**—specify a value for  $n$
  - **Respond only once** (per ProTune step)
- 2** In the **ProTune Session Actions** section, choose one of the following actions:
  - **Continue current step**
  - **Stop Ramp Up**
  - **Stop  $n$  % of active Vusers**—specify a value for  $n$ . The default is 100%
- 3** In the **Notifications** section, select the desired notification method(s):
  - **Highlight measurement in online graph**
  - **Open Alerts Output window**

---

**Note:** You can use the Alert Description pane to change alert action settings (see “Viewing Alert Descriptions,” on page 85).

---

- 4** To set the alert triggers or enable/disable alerts, click **Options**. The Alerts options dialog box opens. Select a trigger option and click **OK**. For more information, see the section “Enabling and Disabling the Alert Mechanism,” on page 91.
- 5** Click **OK** to close the Alerts dialog box.

## **Alert Actions for Multiple Elements on the Same Host**

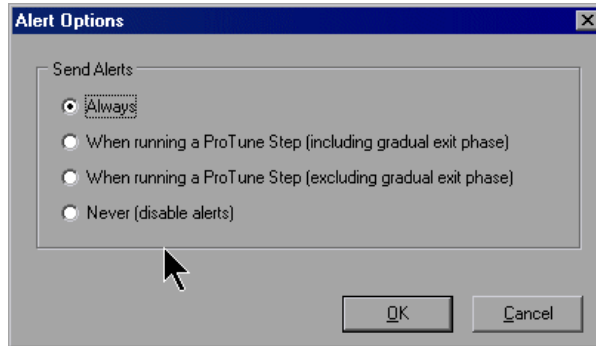
As mentioned above, an alert that is triggered by a measurement on one topology element also affects other elements on the same physical host, so if the alert action is stopping Vusers or stopping ramp-up, all the Vusers on the physical machine are affected.

However, if you specify the action **Highlight measurement in online graph**, the Console indicates the affected elements as follows when the alert is triggered:

- If the measurement that triggered the alert exists in all the topology elements on the host, the alert causes all the elements to blink in the topology window and the relevant graphs. For example, if the alert was defined on the CPU Utilization measurement, all the elements on the host will blink, along with the CPU Utilization graphs for each element.
- If the measurement that triggered the alert exists in only one of the elements, only that element (and its graph) will blink when the alert is triggered. For example, if the alert was defined for an Oracle measurement, only the element representing the Oracle database will blink.

## Enabling and Disabling the Alert Mechanism

You can enable, disable, or limit ProTune alerts, using the Alert Options dialog box.



You can choose one of the following trigger options:

**Always:** Always issue alerts when the alert condition is met, even when ProTune is not executing a step.

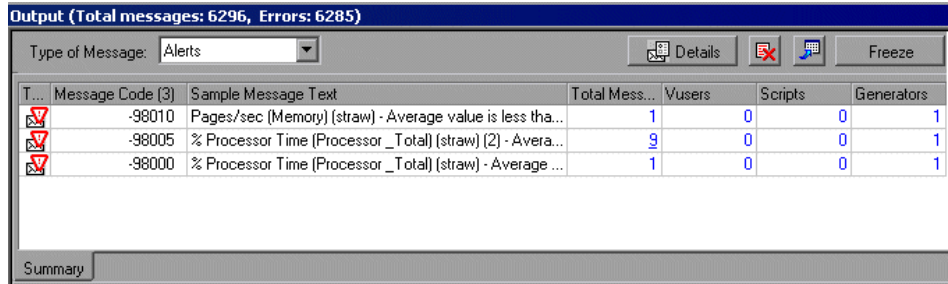
**When Running a ProTune step (including gradual exit):** Issue alerts only during test execution, even during the Vuser's gradual exit stage.

**When Running a ProTune step (excluding gradual exit phase):** Issue alerts only during test execution, except during the Vuser's gradual exit stage. This is the default option.

**Never:** Never issue alerts. This disables the alert mechanism for the active session.

## Viewing Alerts in the Output Window

The Output window in the Execute tab lets you view information about all the alerts triggered during a tuning session.



To open the Output window:

- 1 Click the **Execute** tab.
- 2 In the statistics table (the middle part of the upper section) locate the Alerts Output row.


|                     |          |   |
|---------------------|----------|---|
| Step Status         | Running  |   |
| Active Vusers       | 40       |   |
| Elapsed Time        | 00:03:24 |   |
| Passed Transactions | 7359     | 🔍 |
| Failed Transactions | 0        | 🔍 |
| Errors              | 0        | 🔍 |
| Alerts              | 33       | 🔍 |

The Alerts Output row shows the number of alerts that have been triggered during the session.

Note that alert information is only saved if alerts were enabled during the script's execution. If you enabled alerts after script execution, you will have to run the step again in order to generate alerts.

- 3 Click the magnifying glass in the Alerts Output row. The Output Window opens. Note that the Type of Message box shows "Alerts".

- 4 To show alert details, click **Details**. The Detailed Message Text section opens in the lower part of the Output window.



For more details on using the Output window, see “Viewing the Output Window,” on page 140.





# 7

---

## Scheduling Session Steps

After you create a step, you use the Schedule Builder to specify when the step should begin running. In addition, you can set the duration of the step, or specify that a step should be run till a goal is reached.

This chapter describes:

- Specifying Execution Time
- Creating and Selecting a Profile
- Creating a Manual Profile
- Creating a Goal-Oriented Profile

### About Scheduling Session Steps

An important factor in the creation of a test step, is developing a step that accurately portrays user behavior—the types of actions and the timing of those actions, represented by the scripts.

The Schedule Builder allows you to specify when to run steps, and for how long. You do this by creating *profiles* and associating them with the steps. Each profile defines a specific way of testing. By creating multiple profiles you can run the same step under different conditions.

The Design tab's Schedule Summary section displays information about the profile currently associated with the selected step. It allows you to choose a different profile from the list of available profiles in the Schedule Name box.

The Schedule Builder allows you to create the following types of profile: *manual* and *goal-oriented*.

In a manual profile, you specify:

- start time of a test
- duration
- number of Vusers that will run the test
- the ramp up and ramp down processes

Manual profiles include benchmark profiles, which cause ProTune to run all the enabled scripts on all the Load Generators.

In a goal-oriented profile, you specify the goal you want to reach (for example, the response time that is considered unacceptable). ProTune runs the step, adding Vusers, till the step reaches the goal.

The profile you define is visually displayed in the Schedule Builder's Load Preview graph.

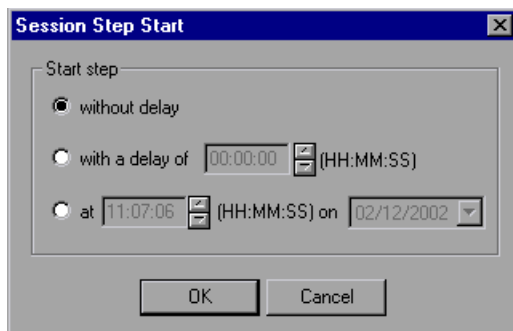
## Specifying Execution Time

By default, when you issue an *Execute* command, ProTune starts the session step immediately. You can instruct ProTune to run the step at a later point in time. You do this by specifying one of the following:

- The period that you want ProTune to delay execution after an *Execute* command is issued.
- The specific time at which you want execution to begin.

**To specify when a session step should be executed:**

- 1 Select **Session > Start Time**. The Session Step Start dialog box opens, with the default option—**without delay**—selected.




- 2 Select **with a delay of xxx (HH:MM:SS)** and enter the period (in hours:minutes:seconds format) by which you want to delay step execution. Alternatively, you can select **at xxx (HH:MM:SS) on xxx** and specify the time (in hours:minutes:seconds format) and date for executing the step.
- 3 Click **OK** to close the dialog box and save your settings.

The next time you execute a session step, the start time will be delayed as specified.

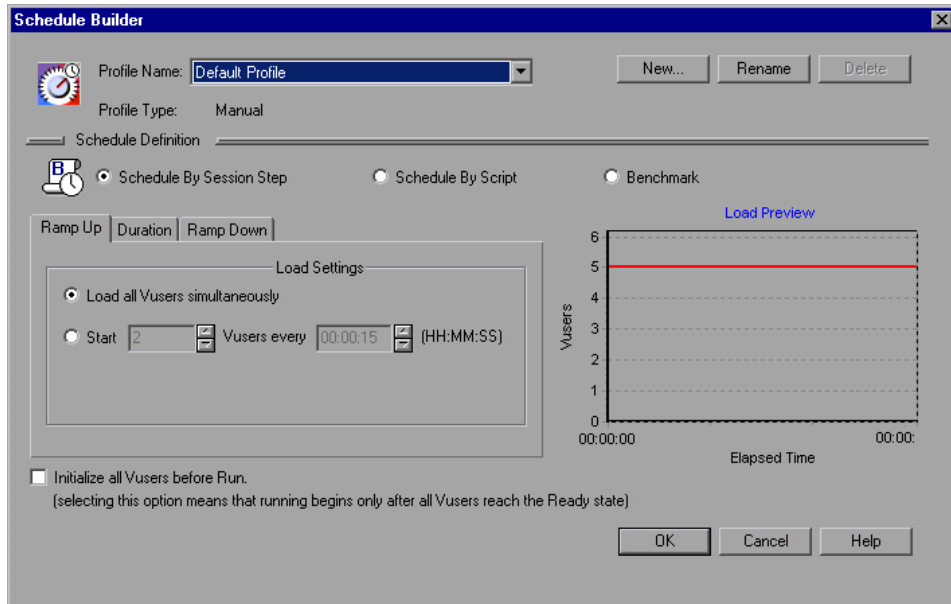
## Creating and Selecting a Profile

You use the Schedule Builder for choosing the profile for your step, creating new profiles, and modifying existing ones. You select the profile that you want to use for your session step from the Profile Name list box in the Schedule Builder window. The default profile loads all Vusers simultaneously.

### To invoke the Schedule Builder:

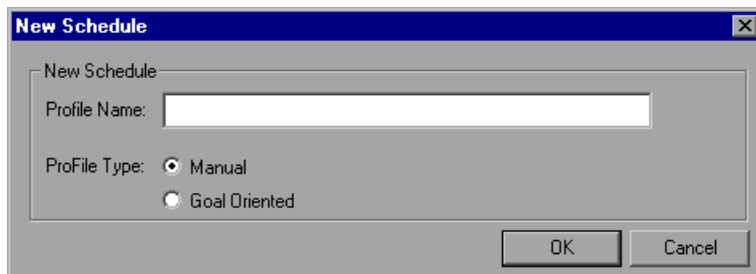


Choose **Session > Schedule Builder** or click one of the Schedule Builder icons (in the toolbar or in the **Design** tab's Schedule Summary section). The Schedule Builder window opens, displaying the last profile that was associated with the step. The profile name is displayed in the Profile Name list box. If no profile has been associated with the step, ProTune by default uses a manual profile and assigns it the name *Default Profile*.



**To create a new profile:**

- 1 Invoke the Schedule Builder.
- 2 Click the **New** button in the Schedule Builder window. The New Schedule dialog box opens.



- 3 In the Profile Name text box, enter the name of the new profile.
- 4 Choose the profile type—Manual or Goal-Oriented—by clicking the appropriate radio button, and click **OK**.

The new profile name appears in the Profile Name list box in the Schedule Builder window.

- 5 Set the values for your profile (see “Creating a Manual Profile,” on page 101 or “Creating a Goal-Oriented Profile,” on page 107).

**To choose an existing profile for your step:**


- 1 Invoke the Schedule Builder.
- 2 Choose a profile from the Profile Name list.

**To modify the properties of an existing profile:**

- 1 Invoke the Schedule Builder.
- 2 Choose the profile you want to modify from the Profile Name list.
- 3 In the Schedule Builder dialog box, modify the profile as required.

**To rename a profile:**

- 1** In the Schedule Builder window, select the profile you want to rename from the Profile Name list.
- 2** Click **Rename**. The Rename Schedule dialog box appears.



- 3** Enter a new name for the selected profile, and click OK. The new name appears in the Profile Name list.

**To delete a profile:**

- 1** In the Schedule Builder window, select the profile you want to delete from the Profile Name list box.
- 2** Click **Delete**. The profile is deleted, and no longer appears in the Profile Name list box.

## Creating a Manual Profile

A manual profile can use the following types of scheduling:

- Session Step Scheduling
- Script Scheduling
- Benchmark Scheduling

### Session Step Scheduling

Session step scheduling means creating a profile that specifies the following:

- how the step should be started (ramped up). This allows you to choose between gradually adding Vusers to the running test, or starting all the Vusers simultaneously when the test starts.
- the step duration
- how the step should be stopped (ramped down). This allows you to choose between gradually stopping Vusers that are running, and stopping them all simultaneously.

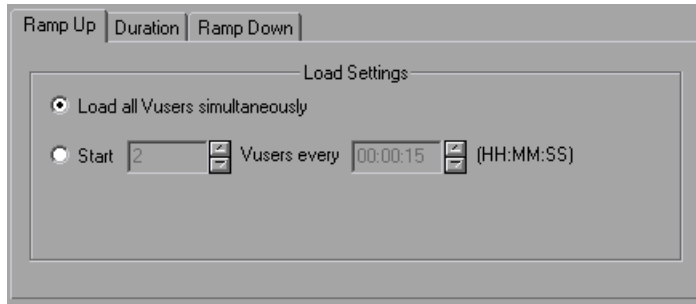
---

**Note:** When you schedule by session step, your settings apply to all the scripts included in the step. For example, if you specify a duration, all the scripts will be executed for the specified period.

---

**To specify how the step should be started:**

- 1** In the Schedule Builder window, click **Schedule by Session Step**.
- 2** Click the **Ramp Up** tab:



- To start all the Vusers running at the same time when the test starts, click the Load all Vusers simultaneously radio button.
- To gradually run the Vusers, enter the number of Vusers you want to begin running concurrently, and the period that you want ProTune to wait before adding more Vusers.

---

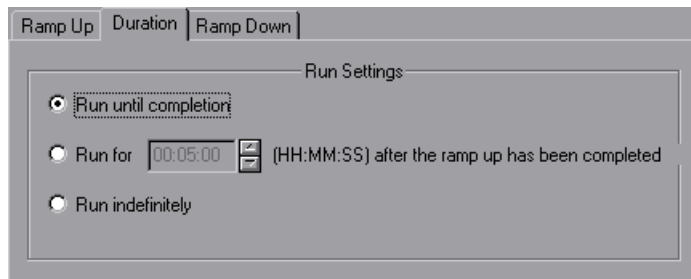
**Note:** While a step is running, you can add scripts to the step and enable them, as long as not all of the Vusers have been ramped up. However, if you add a script after all the Vusers have been ramped up, the new script will not run while the current step is executing. To enable running the new script in the step, you need to stop the step and restart it.

---



**To specify the step duration:**

- 1 In the Schedule Builder window, click the **Duration** tab.



- 2 Choose one of the following options:

- **Run until completion:** Run the step until all its scripts have finished executing.
- **Run for ... after the ramp up has been completed:** Run the step for the specified period after all the Vusers have been ramped up.
- **Run indefinitely:** Run the step until it is manually stopped.

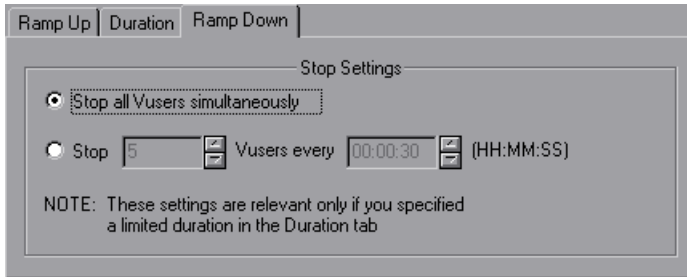
---

**Note:** The duration setting overrides the Vuser iteration run-time setting. For example, if you specify the duration here as five minutes, the Vusers will continue to run as many iterations as required in five minutes, even if the run-time settings specify only one iteration.

---

**To specify how the step should be stopped:**

- 1 Click the **Ramp Down** tab.



- 2 Choose one of the following options:

- **Stop all Vusers simultaneously:** Stops all the Vusers in the session step at once.
- **Stop X Vusers every X (HH:MM:SS):** Stops a certain number of Vusers within a specified time frame. For example, you might want to stop 5 Vusers every 30 seconds.

---

**Note:** The Ramp Down tab settings are enabled only if you select the second option (**Run for ... after the ramp up has been completed**) in the **Duration** tab.

---

**To instruct ProTune to initialize Vusers before beginning to load them:**

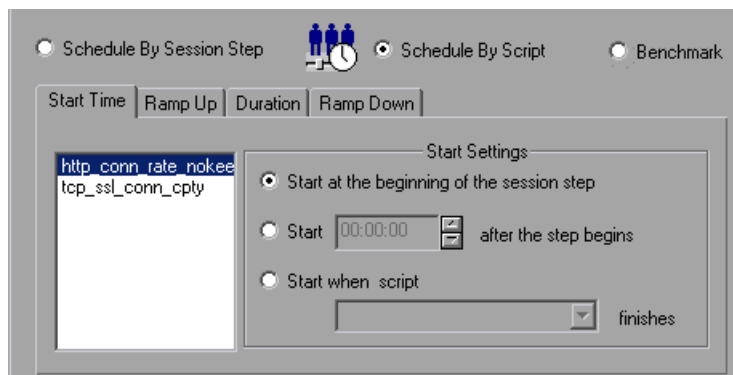
- 1 Check the **Initialize all Vusers before run** box. ProTune will begin to load the Vusers only after they have all reached the READY state.
- 2 Click **OK** to close the Schedule Builder and save your settings.

## Script Scheduling

Script scheduling allows you to specify settings for each script in the step separately. For example, you can schedule a different duration for each script.

### To schedule scripts:

- 1 In the Schedule Builder window, click the Schedule by Script radio button. ProTune displays a **Start Time** tab, in addition to the tabs displayed when you schedule by session step (see Session Step Scheduling on page 101).



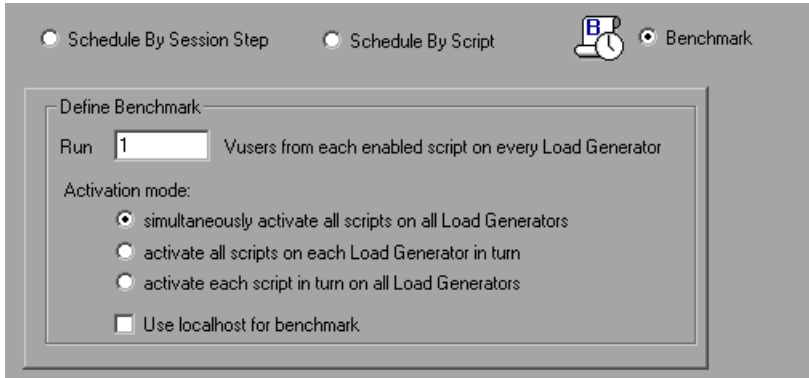
- 2 The **Start Time** tab includes a list of the scripts contained by the step. To specify settings for a particular script, click the step name to select it and choose one of the following settings:
  - To run the script at the beginning of the session step, click the appropriate radio button.
  - To delay running the script, click the **Start HH:MM:SS after the step begins** radio button, and specify the delay period in HH:MM:SS format. ProTune will run the script when the specified period has passed after the step has begun executing.
  - To make the script dependent on another script in the step finishing execution, click the **Start when script finishes** radio button and choose a script from the list box. ProTune will run your script only after the specified script has been run.
- 3 Specify the settings in the other tabs (ramp up, duration and ramp down) as described in “Scheduling Session Steps,” on page 95.

## Benchmark Scheduling

Benchmark scheduling allows you to test all of the enabled scripts on all the Load Generators (instead of only testing each script on the machine to which it is assigned). Benchmark scheduling is particularly useful when dealing with a complicated system topology that has a large number of servers and hosts. Running a benchmark test before tuning lets you verify that all the hosts and scripts are valid.

**To schedule by benchmark:**

- 1 In the Schedule Builder window, click the Schedule by Benchmark radio button. ProTune displays the Define Benchmark pane:



- 2 In the Run Vusers text box, enter the number of Vusers that ProTune should run from each enabled script or Load Generator. For example, if you enter the number 5, each script will be run by five Vusers.

- 3 Click the appropriate radio button to choose one of the following benchmark modes:
  - ▶ Simultaneously activate all the enabled scripts on all of the Load Generators.
  - ▶ Activate all the enabled scripts on each Load Generator in turn. For example, ProTune first runs all the scripts on Load Generator 1, next on Load Generator 2, and so on.
  - ▶ Activate each enabled script in turn on all the Load Generators. For example, ProTune runs script 1 on Load Generator 1, next on Load Generator 2, and so on until the script has been run on all the Load Generators. ProTune next runs script 2 on all of the Load Generators, and so on until all the scripts have been run on all the Load Generators.
- 4 If you want to use the local machine as one of the Load Generators, check the **Use localhost as one of the Load Generators** box.

## Creating a Goal-Oriented Profile

You specify a goal for a session step—not for a script.

When you run a session step, the goal you defined is displayed in the appropriate graph, along with the session results. This enables you to compare the results with your target goal and determine if your goal was reached. If your goal is not reached, you reconfigure your applications and servers accordingly, in order to reach the desired goal.

Creating a goal-oriented profile includes defining the following:

- ▶ Step goal—the goal that you want to achieve before terminating execution of the step. This includes the value that you want to achieve (for example, the number of concurrent transactions) and the maximum number of Vusers that will participate in the session step.
- ▶ Step settings—when ProTune should run the step and what it should do if the goal is not reached. **Note:** Times in the scheduler are approximate; although ProTune attempts to reach the target within the specified period, the actual period may be affected by external factors (for example, Internet connections).

- Load behavior—how and when you want ProTune to reach your target. ProTune starts each step by ramping up.

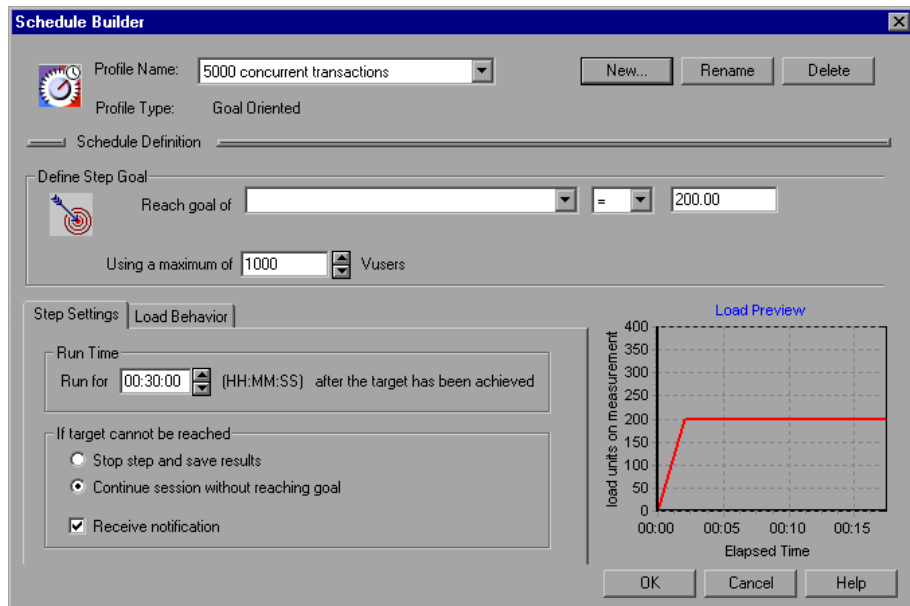
When ProTune starts executing a goal-oriented step, it first executes the step for two minutes with one Vuser from each group. Based on the performance that it measures during the two minutes, ProTune calculates the number of Vusers needed to reach the goal. The subsequent behavior depends on the selected Ramp Up option (*Automatic* or *Reach target after*) in the **Load Behavior** tab.

- If you select *Automatic*, the rest of the Vusers needed are run immediately. For example, if ProTune was running one Vuser during the initial two minutes, and it calculates that a total of ten Vusers are needed to reach the goal, it starts another nine Vusers and runs them all for another two minutes. If the goal is not reached by the end of this period, ProTune calculates the difference between the value that was reached and the goal, and will either start another batch of virtual users or stop some virtual users, depending on whether the value is below the goal or above it, respectively. Note that if you use the equals operator (“=”) to specify an exact value for the goal, ProTune considers the goal to have been reached only if the actual value stays within 2% of the goal for 60 seconds. For example, if you specify a goal of exactly 100 connections, and the actual number of connections is 110, ProTune does not consider the goal to have been reached.
- If you select *Reach target after*, ProTune spreads the starting of the Vusers over the time period you specify. For example, if you specify ten minutes and ProTune calculates that it needs five Vusers, one Vuser will be started every two minutes. Note that ProTune adjusts the number of Vusers that it starts each time (the batch) so that adding each batch results in the same change of measurement over the same time period. If the current batch does not cause the measurement to increase as much as the previous batch, ProTune increases the number of Vusers in the batch. Similarly, if the current batch causes a greater change than the previous one, ProTune decreases the number of Vusers in the batch.

Once the goal has been reached, ProTune keeps running it for the period specified in the Step Settings section, adding or subtracting Vusers as needed to keep the actual measurement within 6% of the specified target. For example, if the Step Goal is specified as 100 hits/sec, ProTune adds or subtracts Vusers as needed to keep between 94 and 106 hits per second.

You can define multiple goal-oriented profiles for each step, and execute the step using a different profile each time.

When you create a goal-oriented profile, the Schedule Builder displays the following window:



### To define the step goal:

- 1 Choose the type of goal for your step from the **Reach goal of** list box.

The list of goal types includes:


- basic HTTP-related goals. These goal types are displayed when you open the list box.
- additional goals that are related to the monitors and measurements you have specified. To view these goal types, choose **<more>** in the **Reach goal of** box.

Following is the list of basic HTTP-related goals:

- ▶ **Throughput**—target downstream bandwidth.
  - ▶ **Hits / Second**—target number of hits per second (HTTP requests per second) that you would like your step to reach. When you choose this goal, you also need to enter the maximum number of Vusers for the session step.
  - ▶ **Number of Connections**—target number of connections that you would like the server to host.
  - ▶ **Connections / Second**—target number of connections per second you would like the server to handle.
- 2 Choose an operator (either “=” or “>=”)from the middle list box.
  - 3 Enter a value in the text box on the right side.
  - 4 Specify the maximum number of Vusers in the **Using a Maximum of ... Vusers** box.

**To define step settings:**

- 1 Click the Step Settings tab.




- 2 Enter a value (in HH:MM:SS format) in the **Run For** box. This value specifies how long the step will run after the goal has been reached.
- 3 Specify what ProTune should do if the goal is not reached: click the appropriate radio button to either stop the session and save the results, or continue the session despite not reaching the goal.
- 4 Check the Receive Notification box if you want ProTune to display a message when it determines that the goal cannot be reached.



**To define load behavior:**

- 1 Click the Load Behavior tab.



- 2 Specify the ramp up as follows:
  - To start all the required Vusers simultaneously, click the **Automatic** radio button.
  - To specify the total time required for starting all the batches of Vusers, click the **Reach target after** radio button and enter the period in the box in HH:MM:SS format. (Minimum: 2 minutes.) ProTune starts a batch of Vusers approximately every 2 minutes, and attempts to make the batches equal in size.

---

**Note:** Reaching the goal depends on the maximum number of Vusers specified in the Define Step Goal settings.

---



# 8

---

## Preparing to Run a Session Step

Before you run a session step, you specify a location for the session step results and other run-time related settings.

This chapter describes:

- ▶ Specifying a Results Location
- ▶ Results Directory File Structure
- ▶ Collating Results

### About Preparing to Run a Session Step

Before you run a session step, you need to specify the location of the results (mandatory), assign a name to the results, schedule the session step, and provide session step summary information. In addition, you can specify the applications to invoke at the start of a session step.

Although most of the pre-session step settings are optional, by using them you can enhance the testing process. These values are session step specific—you can set different values for each ProTune session step.


For information on one-time configuration settings such as timeout, output, and quotas, see Chapter 5, “Configuring Session Steps.”

## Specifying a Results Location

When you run a session step, by default the run-time files are stored locally on each load generator. After the session step, the results are collated together and processed on the Console machine. Alternatively, you can instruct ProTune to save the results on a shared network drive. For information about specifying a file storage method, see the Run-Time File Storage settings in Chapter 5, “Configuring Session Steps.”

ProTune allows you to give descriptive names to each result set. This is especially useful for cross results analysis, in which ProTune superimposes the results of several session step runs in a single graph and lets you compare the results of multiple session step runs. The descriptive graph names enable you to distinguish between the results of the multiple runs.

In the example below, the results of two session step runs are superimposed. The result sets are *res12*, and *res15*.

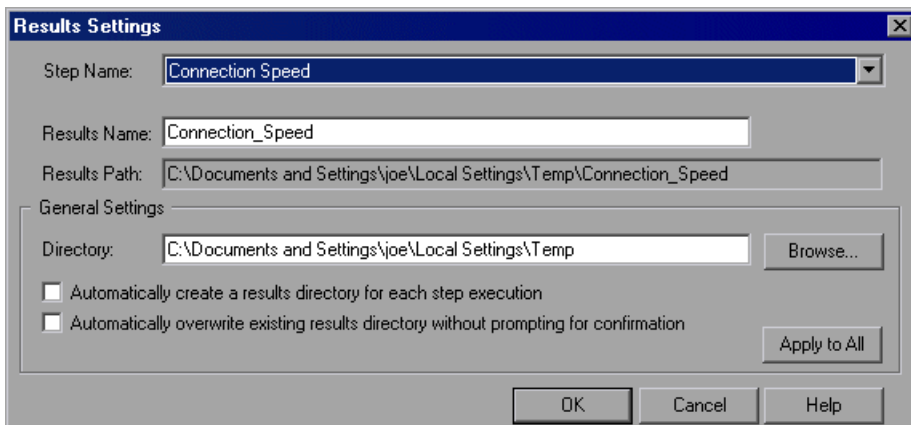


For more details on cross result graphs, refer to the *ProTune Analysis User's Guide*.

You can specify separate result settings for each session step.

**To specify where results are stored:**

- 1 Choose **Results > Results Settings**. The Set Results Directory dialog box opens.



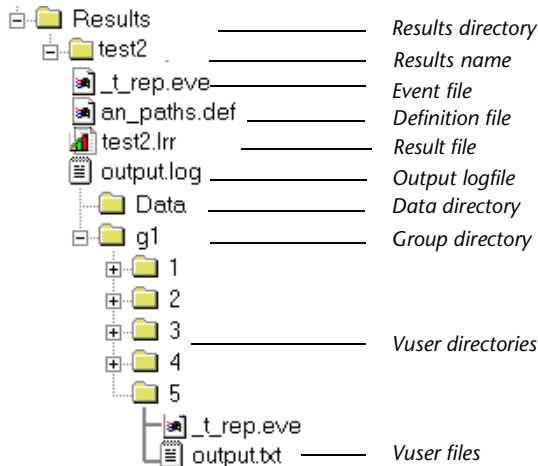
- 2 The Step Name box shows the step to which the settings apply. To specify result settings for a different step, choose the step from the list in the box.
- 3 By default, ProTune uses the step name as a basis for the results name. For example, if the step name is "Connection Speed", ProTune adds an underscore and uses "Connection\_Speed" as the results name.
- 4 To specify a different results name, enter the new name in the Results Name box. Avoid using the same name with different paths, since the names will appear identical on the graphs.
- 5 The Directory box shows the full path to the results directory. To change the results directory, enter the new path in the box. If you are using the default file storage setting (local machine), specify a directory in which to store all of the collated results after the session step run. If you specified a shared network drive as the file storage method, specify the directory to which Vuser groups should write during session step execution.
- 6 Using the results name from step 2, the Console creates a subdirectory within the results directory. All results from this step are saved within this subdirectory. The Results Path field shows the full path to the subdirectory.

- 7 Select the appropriate check box for subsequent executions: **Automatically create a results directory for each session step execution** or **Automatically overwrite existing results directory without prompting for confirmation**.
- 8 To apply the settings specified in the dialog box to all the steps in your session, click **Apply to All**.
- 9 Click **OK** to save the results directory setting.

## Results Directory File Structure

When you set the results directory, you also specify a results name. ProTune creates a subdirectory using the results name, and places all of the data it gathers in that directory. Every set of results contains general information about the session step in a result file (*.lrr*) and an event (*.eve*) file.

During session step execution, ProTune also gathers data from each Vuser and stores it in an event file *\_t\_rep.eve* and an output file *output.txt*. ProTune creates a directory for each group in the session step and a subdirectory for each Vuser. A typical result directory has the following structure:



- *t\_rep.eve* in the main result directory contains Vuser and rendezvous information.
- *\*.def* are definition files for graphs that describe the online and other custom monitors.
- *results\_name.lrr* is the ProTune Analysis document file.
- *output.log* contains output information about the session step generated during test execution.
- The *Data* directory contains the database created by the Analysis (from the results files).
- *gl* is a group directory. A separate directory exists for each Vuser group that runs in the session step. Each group directory consists of Vusers subdirectories.
- *t\_rep.eve* in each Vuser directory contains transaction information.
- *output.txt* in each Vuser directory contains output information generated during replay.

When you generate analysis graphs and reports, the ProTune Analysis engine copies all of the session step result files (*.eve* and *.lrr*) to a database. Once the database is created (and stored in the *Data* directory), the Analysis works directly with the database and does not use the result files.

For information on ProTune Analysis, refer to the *ProTune Analysis User's Guide*

## Collating Results

When you run a session step, by default all Vuser information is stored locally on each load generator. After session step execution, the results are automatically collated or consolidated—results from all of the load generators are transferred to the results directory. You set the location of the results directory as described in “Specifying a Results Location,” on page 114.

**Note:** If you have selected to store all the session step results directly to a shared network drive, then collation of the results is not required. See “About Configuring a Session,” on page 69 for details on changing how results are stored.

---

To disable automatic collation and clear the check mark adjacent to the option, choose **Results > Auto Collate Results**. To manually collate results, choose **Results > Collate Results > Collate**. The Collating Files dialog box opens, displaying the progress of result and log file collation from each load generator. To stop collating the results and close the dialog box, click **Stop** and then **Close**. To resume collating the results, select **Results > Collate Results > Continue stopped collation**.

---

**Note:** You can choose to disable log file collation. For more information, see “Options - General Settings,” on page 647.

---

The log and result directories are only deleted from a load generator once ProTune successfully collates the results from the machine. You can therefore close the Console after saving a session step, and collate the results once you reopen the session step in the Console.

If collation fails due to a lack of disk space, select **Results > Collate Results > Recollate**. ProTune attempts to collate the results again.

Before generating the analysis data, ProTune automatically collates the results if they have not previously been collated.

---

**Note:** If you enabled the **Auto Load Analysis** option in the Results menu, the Analysis may open during a lengthy collation process, displaying Analysis summary data.

---



# Part IV

---

## Executing a Tuning Session



# 9

---

## Running a Session

When you tune a session, ProTune simulates your environment and measures the system's performance.

This chapter describes:

- ▶ Running an Entire Session
- ▶ Controlling a Specific Number of Vusers
- ▶ Continuing With Subsequent Steps
- ▶ Adding Vusers to a Running Session
- ▶ Viewing and Controlling Vusers
- ▶ Invoking the System Topology Window

### About Running a Session Step

When you run a session step, the Vusers are assigned to their load generators and execute their scripts. During session step execution, ProTune:

- ▶ records the durations of the transactions you defined in the scripts
- ▶ performs the rendezvous included in the scripts
- ▶ collects error, warning, and notification messages generated by the Vusers

You can run an entire session step unattended, or you can interactively select the Vusers that you want to run. When the session step starts running, the Console first checks the session configuration information. Next, it invokes the scripts that you selected to run with the step. Then, it distributes each script to its designated load generator. When the Vusers are ready, they start executing their scripts.

While the step runs, you can monitor each Vuser, view error, warning, and notification messages generated by the Vusers, and stop both Vuser groups and individual Vusers. You can instruct ProTune to allow an individual Vuser or the Vusers in a group to complete the iterations they are running before stopping, to complete the actions they are running before stopping, or to stop running immediately. For more information, see “Configuring Session Run-Time Settings” on page 70.

You can also activate additional Vusers while the step is running, using the Run/Stop Vusers dialog box. For more information, see “Adding Vusers to a Running Session” on page 126.

The session step ends when all the Vusers have completed their scripts, when the duration runs out, or when you terminate it.

If you specified subsequent steps for execution (see “Specifying Step Execution Order,” on page 45), depending on whether or not the current step was successful, ProTune tells you the name of the next step and gives you the option of executing it immediately.

**The following procedure outlines how to run a session step:**

- 1** Open an existing session or create a new one.
- 2** Configure and schedule the session.
- 3** Set the results directory.
- 4** Run and monitor the session.

---

**Note:** When creating and running a script in VuGen, the full browser is used. This differs from a test run in the Console, where only the browser basics are used. There may be occasions when a test passes its run in VuGen, but fails when it is run in the Console. Before running a session in the Console with multiple Vusers, run a single Vuser to ensure the script is bug-free.

---


## Running an Entire Session

You can run all the Vusers in a session, or you can select the number of Vusers that you want to run. Note that when you run your session, ProTune runs them as soon as they reach the READY state.

The following section describes how to run an entire session. “Viewing and Controlling Vusers,” on page 129 describes how to manipulate individual Vusers.

**To run a session step:**

- 1 Open an existing session or create a new one.
- 2 Click the **Execute** button in the bottom right of the Design tab, or click the Execute tab and then click the **Start the current Session Step** button. The ProTune Console Session window displays the selected monitors.



The screenshot shows the ProTune Console interface during a session execution. The title bar reads "ProTune Console - Session1 - [Execute]". The menu bar includes File, View, Session, Monitors, Results, Tools, and Help. The toolbar contains buttons for Scheduler, Generators, Monitors, Alerts Definition, Graphs, Topology, Analyze Results, and a Help icon. Below the toolbar are buttons for User Management, Reset Step, Session Summary, and Session Report.

The main area is titled "Current Step: New Step". On the left, there are controls for "Vusers:" (a text box with "20" and a play button) and "Step:" (navigation buttons). On the right, a table displays session statistics:

|                     |          |
|---------------------|----------|
| Step Status         | Running  |
| Active Vusers       | 10       |
| Elapsed Time        | 00:03:25 |
| Passed Transactions | 40       |
| Failed Transactions | 0        |
| Errors              | 1838     |
| Alerts              | 0        |

Below the table are two graphs. The left graph is "Error Statistics - Whole Step", showing the number of errors over time. The right graph is "Windows Resources on localhost - Last 60 sec", showing resource usage over time.

At the bottom, there is a table with columns: Color, Scale, Error Location, Max, Min, Avg, Std, Last. The data row shows:

| Color | Scale | Error Location           | Max      | Min   | Avg      | Std     | Last     |
|-------|-------|--------------------------|----------|-------|----------|---------|----------|
|       | 1     | tcp_conn_cply:Actions:22 | 1862.000 | 1.000 | 1367.968 | 485.455 | 1738.000 |

The bottom of the window has tabs for Design, Execute, and Tune. At the very bottom right are buttons for "Auto Load Analysis" and "Auto Collate Results".



- 3 Choose **Session > Stop** or click the **Stop** button to terminate session step execution. If you selected the **Exit immediately** option in the Run-Time Settings tab of the Options dialog box, all of the Vusers in the session move to the `EXITING` status.

If you selected the **Wait for the current iteration to end before exiting** or **Wait for the current action to end before exiting** options in the Run-Time Settings tab of the Options dialog box, the Vusers in the session move to the `GRADUAL EXITING` status and exit the session gradually. To stop the Vusers immediately, click **Stop Now**.



- 4 Click the **First Step** button to load the first session step.



- 5 Click the **Previous Step** button to load the previous session step.



- 6 Click the **Current Step** button to run the Vusers in the first session step.



- 7 Click the **Next Step** button to load the next session step.



- 8 Click the **Reset Step** button to reset the session step. This initializes all the Vusers, returning them to `DOWN` status, and clears the step statistics.

You can also add Vusers during the session run. For more information, see “Adding Vusers to a Running Session” on page 126.

## Controlling a Specific Number of Vusers

You can instruct ProTune to run or stop a specific number of Vusers. It distributes these Vusers between the load generators as you defined in the design stage.

To control a specific number of Vusers:

- 1 Open an existing session or create a new one in the **Design** tab.
- 2 Select the **Execute** tab. The ProTune Console Session window opens with the selected monitors.

- 3 In the upper left section, enter the number of Vusers you want to manipulate in the Vusers box

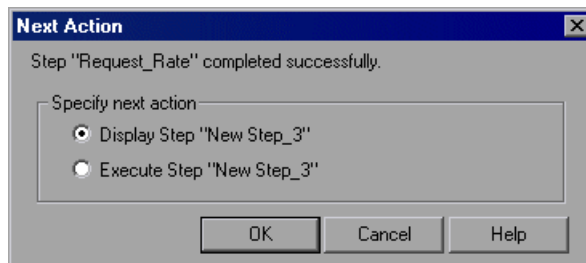


- 4 To run a specific number of Vusers, click the **Run Specific Vusers** button. ProTune begins running those Vusers.
- 5 To stop a specific number of Vusers, click the **Stop Specific Vusers** button. ProTune stops those Vusers.

## Continuing With Subsequent Steps

When you defined your step, you had the option of specifying the steps that should be executed after the current step completes execution (see “Specifying Step Execution Order,” on page 45).

If you specified the subsequent step or steps, ProTune displays the Next Action dialog box when the current step completes execution.



The Next Action dialog box tells you whether the step succeeded or failed, displays the name of the appropriate subsequent step, and lets you choose between the following actions:

- ▶ **Display Step**—Display the subsequent step in the **Design** tab, but does not execute it.
- ▶ **Execute Step**—Executes the subsequent step immediately.

Choose the appropriate radio button and click **OK**, or click **Cancel** to remain in the **Execute** tab without executing any step.

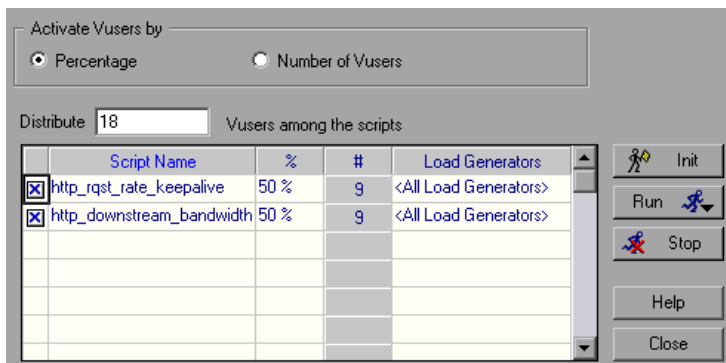
## Adding Vusers to a Running Session

You can add Vusers to a running session to manually ramp up the load and see how your system performs when users are added. You can either distribute Vusers among scripts by percentage, or specify the number of Vusers for each script.

To add Vusers to a running session:



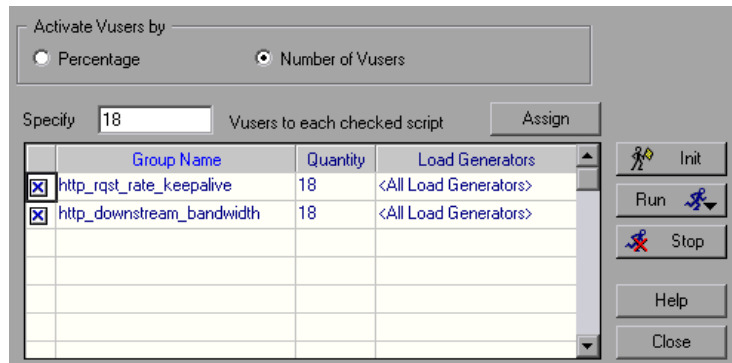
- 1 Click the **Add Vusers to a Running Step** button in the top left section of the Execute tab. The Run/Stop Vusers dialog box opens.
- 2 Click **Percentage** in the **Activate Vusers by** section to display the Vusers by their percentage distribution.





The % column specifies the percentage of Vusers assigned to each script. The # column specifies the number of Vusers assigned to each script.

- ▶ To automatically distribute the total number of Vusers among the checked scripts, enter the total number in the **Distribute ... Vusers** field. ProTune distributes the Vusers among the scripts, according to the percentage specified for each script.
  - ▶ To manually change the percentage of Vusers assigned to a script, enter the new percentage in the script's % column. ProTune automatically adjusts the percentages of the other scripts.
- 3** Click **Number of Vusers** in the **Activate Vusers by** section to display the Vusers by their quantity distribution.



To manually change the number of Vusers assigned to a script, enter the new number in the script's Quantity column.

---

**Note:** If more than one load generator is defined for a script, the added Vusers are proportionally distributed among the defined load generators.

---

- 4** To disable a script, clear the check box to the left of the script name. Note that a script will automatically appear disabled if it is disabled in the Design view.

**Note:** If you disable a script, no Vusers will be distributed to it. However, 100 percent of the Vusers will not be distributed among the remaining scripts, unless you define a 0 percent value for the disabled script.

---

- 5** Click the **Init** button to initialize the number of Vusers you added. The Console first initializes the Vusers in your session that have not yet been run, on the load generator(s) defined in the Run/Stop Vusers dialog box. It then adds additional Vusers, as required, to reach the quantity defined in the Run/Stop Vusers dialog box.
- 6** Click the **Run** button, and select one of the following options:
  - **Run Initialized:** Runs the Vusers that have already been initialized on the load generators defined in the Run/Stop Vusers dialog box. The Console runs only those Vusers that have already been initialized, regardless of their quantity.
  - **Run New:** Runs the number of Vusers you specified. The Console first runs the Vusers in your session that have not yet been run, on the load generator(s) defined in the Run/Stop Vusers dialog box. It then adds additional Vusers, as required, to reach the quantity defined in the Run/Stop Vusers dialog box.
- 7** Click **Stop** to stop the Vusers that are running on the load generator(s) defined in the Run/Stop Vusers dialog box. The Console stops the Vusers according to the settings you defined in the Run-Time Settings tab of the Options dialog box.
- 8** Click **Close** to close the Run/Stop Vusers dialog box.

## Viewing and Controlling Vusers

ProTune allows you different views of the Vusers running in a step:

- ▶ By group—This shows you the Vusers in each group, broken down by status (for example, Down, Pending, and so on).
- ▶ By load generator—This shows you the number of Vusers that each group is running on each load generator.

You can also manipulate individual Vusers. This includes initializing, running, and stopping individual Vusers.

### To view Vuser groups:

- ▶ Choose **Session > Vuser Management** or click the **Vuser Management** button to open the Vuser Management—Groups window. (If ProTune opens the Vuser Management—Groups by Load Generator window instead, click the **Show Groups** button in the window's task bar to display the Vuser Management—Groups window.)

 Vuser Management



| Group Name       | Down | Pending | Init | Ready | Run | Rendez Passed | Failed | Error | Gradual Exiting | Exiting | Stopped |
|------------------|------|---------|------|-------|-----|---------------|--------|-------|-----------------|---------|---------|
| 3                | 5    | 0       | 2    | 0     | 8   | 0             | 0      | 0     | 0               | 0       | 0       |
| http_rqst_rate_1 | 4    |         | 1    |       | 3   |               |        |       |                 |         |         |
| http_rqst_rate_1 |      |         |      |       | 3   |               |        |       |                 |         |         |
| http_conn_rate_1 | 1    |         | 1    |       | 2   |               |        |       |                 |         |         |
|                  |      |         |      |       |     |               |        |       |                 |         |         |
|                  |      |         |      |       |     |               |        |       |                 |         |         |
|                  |      |         |      |       |     |               |        |       |                 |         |         |
|                  |      |         |      |       |     |               |        |       |                 |         |         |
|                  |      |         |      |       |     |               |        |       |                 |         |         |

Each row of the table lists a group of Vusers in the session step. The columns show how many of the group's Vusers have each status.

**To view Vusers by load generator:**



- 1 In the Vuser Management—Groups window, click the **Show Groups by Load Generator** button. The Vuser Management—Groups by Load Generator window opens.

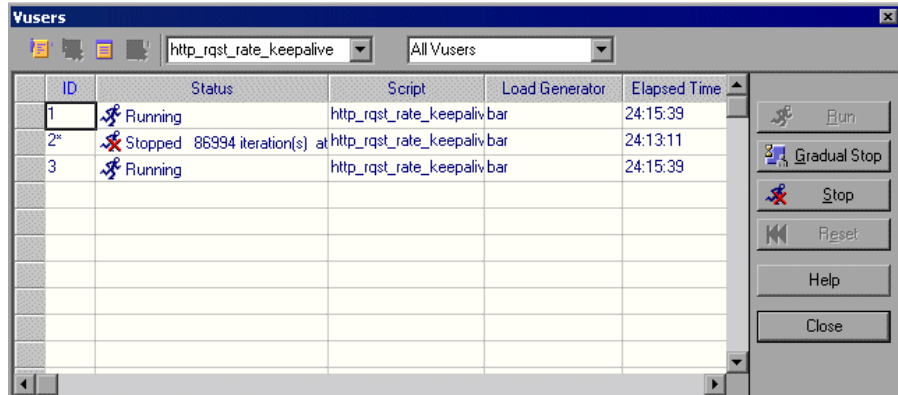
| Vuser Management |           |     |       |
|------------------|-----------|-----|-------|
| All Active       |           |     |       |
| Group Name       | localhost | bar | Total |
| 3                | 3         | 9   | 12    |
| .rqst_rate_keep  | 2         | 3   | 5     |
| qst_rate_keepa   | 1         | 2   | 3     |
| onn_rate_nokee   |           | 4   | 4     |
|                  |           |     |       |
|                  |           |     |       |
|                  |           |     |       |
|                  |           |     |       |
|                  |           |     |       |

- 2 From the window's drop-down list, specify which Vusers you want to display (that is, Vusers with a particular status).

Each row of the table lists a group of Vusers in the session step. The columns show how many Vusers in each group, with the status you specified, are running on each load generator.

**To control an individual Vuser:**

- 1 In the Groups window or the Groups by Load Generator window, double-click a group, or select the group and click the **Show Vusers in Group** button. The Vusers in Group dialog box opens, showing the ID, Status, Script, Load Generator, and Elapsed Time (since the beginning of the session) for each of the Vusers in the group.




---

**Note:** To choose a different Vuser group, select it from the left list box.

---

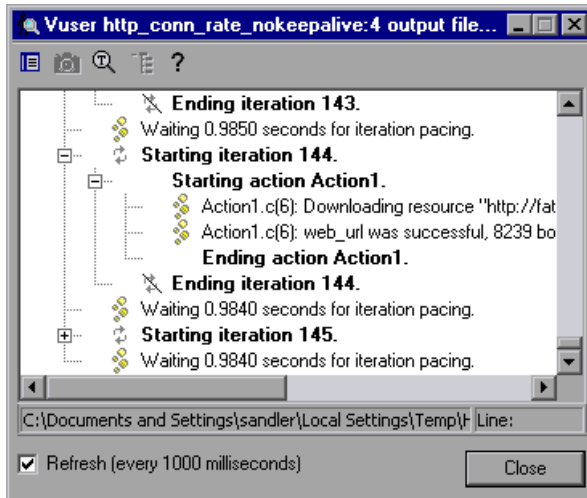
You can control an individual Vuser by doing the following:

- To run a Vuser, select it and click **Run**.
- To immediately stop a running Vuser, select it and click **Stop**.

If you selected the **Wait for the current iteration to end before exiting** or **Wait for the current action to end before exiting** options in the **Run-Time Settings** tab of the Options dialog box, and want to gradually stop a Vuser in the RUN state, click **Gradual Stop**. The Vuser moves to the GRADUAL EXITING status and gradually exits the session.

- 2 To pause a Vuser, right-click it and select **Pause**.
- Select a Vuser and click **Reset** to change its status to DOWN.
  - To initialize a Vuser, right-click it and select **Initialize Vuser/s**.

- ▶ To filter the Vusers listed, right-click in one of the columns and select **Filter Vusers**. Select the way in which you want to filter the Vusers. Alternatively, you can select the filter option you want to use from the right-hand filter selector at the top of the Vusers dialog box.
- ▶ To sort the listed Vusers, right-click in one of the columns and select **Sort Vusers**. Select the way in which you want to sort the Vusers.
- ▶ To view a Vuser executing its assigned script, select the Vuser and click the **Show the selected Vusers** button. The Run-Time Viewer opens, allowing you to see the Vuser executing the script.
- ▶ To close the Run-Time Viewer, click the **Hide the selected Vusers** button.
- ▶ To view the script log, click the **Show Vuser log** button. A script log, such as the following, appears.



To close the script log, click **Close** in the log window or click **Hide Vuser log** in the Vusers dialog box. For more information on the script log, see page 143.

- 3 Click **Close** to close the Vusers dialog box.

## Invoking the System Topology Window

The Execute view's topology pane displays the system topology you defined (see "Creating a Topology," on page 9).

During a tuning session, you may want to make adjustments to the topology or view its details. You can open the System Topology window either by clicking the Topology button on the main toolbar, or double-clicking the white space in the Execute View's topology pane.





# 10

---

## Viewing Vusers During Execution

During session execution, you can view the actions that are performed by Vusers.

This chapter describes:

- Monitoring Vuser Status
- Viewing the Output Window
- Viewing the Script Log
- Viewing the Agent Summary

### About Viewing Vusers During Execution

ProTune lets you view Vuser activity during a session:

- On the Console load generator machines, you can view the Output window, monitor Vuser performance online, and check the status of Vusers executing the session.
- On remote machines, you can view the Agent summary with information about the active Vusers.





## Monitoring Vuser Status

During session execution, you can use the Session Groups window in the Execute view to monitor the actions of all the Vusers and Vuser groups in the session.

The Status field of each Vuser group displays the current state of each Vuser in the group. The following table describes the possible Vuser states during a session.

| Status          | Description   |
|-----------------|---|
| DOWN            | The Vuser is down.  |
| PENDING         | The Vuser is ready to be initialized and is waiting for an available load generator, or is transferring files to the load generator. The Vuser will run when the conditions set in its scheduling attributes are met. |
| INITIALIZING    | The Vuser is being initialized on the remote machine.   |
| READY           | The Vuser already performed the init section of the script and is ready to run.   |
| RUNNING         | The Vuser is running. The script is being executed on a load generator.   |
| RENDEZVOUS      | The Vuser has arrived at the rendezvous and is waiting to be released by ProTune.   |
| DONE.PASSED     | The Vuser has finished running. The script passed.  |
| DONE.FAILED     | The Vuser has finished running. The script failed.  |
| ERROR           | A problem occurred with the Vuser. Check the Status field on the Vuser dialog box or the output window for a complete explanation of the error.   |
| GRADUAL EXITING | The Vuser is completing the iteration or action it is running (as defined in Tools > Options > Run-Time Settings) before exiting.   |
| EXITING         | The Vuser has finished running or has been stopped, and is now exiting.   |
| STOPPED         | The Vuser stopped when the Stop command was invoked.  |

You can also view a synopsis of the running session in the status window at the top of the Execute view.

|                     |          |   |
|---------------------|----------|---|
| Step Status         | Running  |   |
| Active Vusers       | 5        |   |
| Elapsed Time        | 01:44:10 |   |
| Passed Transactions | 31095    |  |
| Failed Transactions | 0        |  |
| Errors              | 0        |  |
| Alerts              | 209      |  |

| Status Summary      | Description  |
|---------------------|--|
| SESSION STATUS      | indicates whether the session is RUNNING or DOWN   |
| RUNNING VUSERS      | indicates how many Vusers are being executed on a load generator machine   |
| ELAPSED TIME        | indicates how much time has elapsed since the beginning of the session   |
| HITS/SECOND         | indicates how many hits (HTTP requests) there have been to the Web site being tested per second that each Vuser has been running |
| PASSED TRANSACTIONS | indicates how many transactions have been executed successfully  |
| FAILED TRANSACTIONS | indicates how many transactions have been executed unsuccessfully  |
| ERRORS              | indicates how many problems have occurred with the Vusers  |

---

**Note:** When running a goal oriented step, the measurement used for the goal appears following the elapsed time.

---

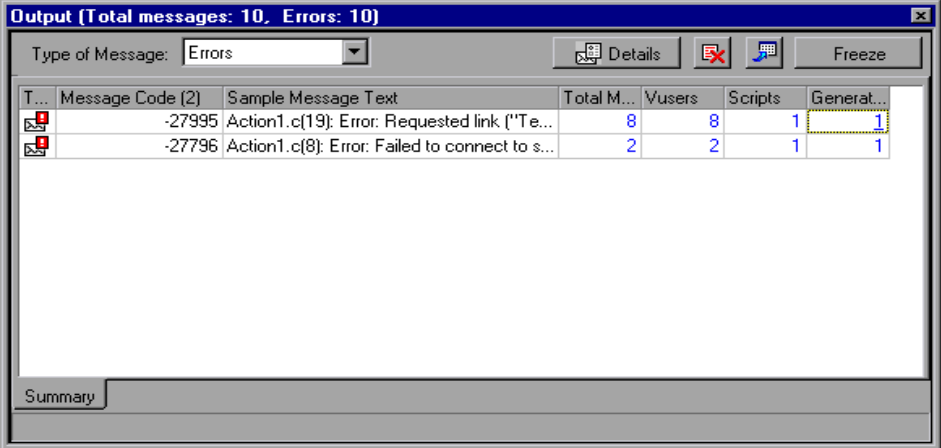
**To view details of the transactions and errors:**



- 1 Click the **Show Snapshot** button to the right of the Passed Transactions or Failed Transactions in the Session Status window. The Transactions dialog box opens.

| Name                   | TPS | Passed | Failed | Stopped |
|------------------------|-----|--------|--------|---------|
| Action1_Transaction    | 0.0 | 10     | 0      | 0       |
| vuser_end_Transaction  | 0.5 | 10     | 0      | 0       |
| vuser_init_Transaction | 0.0 | 10     | 0      | 0       |
|                        |     |        |        |         |
|                        |     |        |        |         |
|                        |     |        |        |         |
|                        |     |        |        |         |
|                        |     |        |        |         |
|                        |     |        |        |         |

The **Name** column lists the individual transactions in a script. For each transaction, the Transactions dialog box lists information concerning the number of **Transactions Per Second (TPS)**, the number of transactions that **Passed**, the number of transactions that **Failed**, and the number of transactions that **Stopped** before completion.

- 2 Choose **View > Show Output** or click the **Show Snapshot** button to the right of the Errors listing. The Output window opens, displaying a list of the Error log information.



| T...  | Message Code (2) | Sample Message Text                            | Total M... | Vusers | Scripts | Generat... |
|---|------------------|--|------------|--------|---------|------------|
|  | -27995           | Action1.c(19): Error: Requested link ("Te...   | 8          | 8      | 1       | 1          |
|  | -27796           | Action1.c(8): Error: Failed to connect to s... | 2          | 2      | 1       | 1          |

Summary

For each type of error message code, the Output window lists a sample message text, the total number of messages generated, the Vusers and load generators that generated the code, and the scripts in which the errors occurred. To view details of the log information by message, Vuser, script, or load generator, click the link in the respective column. For more information on the Output window, see the following section.

## Viewing the Output Window

While the session runs, the Vusers and load generators send error, notification, warning, debug, and batch messages to the Console. In addition, the Console generates alert messages. You can view all of these messages in the Output window.

| T... | Message Code (5) | Sample Message Text                            | Total M... | Vusers | Scripts | Generat... |
|------|------------------|--|------------|--------|---------|------------|
|      | -27995           | Action1.c(19): Error: Requested link ("Te...   | 8          | 8      | 1       | 1          |
|      | -27798           | -27798 : Action1.c(6): Warning: could no...    | 70         | 10     | 1       | 1          |
|      | -27798           | Action1.c(6): Error: could not resolve ad...   | 10         | 10     | 1       | 1          |
|      | -19890           | Action1.c(6): Error -19890 : C-interpreter ... | 10         | 10     | 1       | 1          |
|      | 0                | Error from ftp_logon_ex at Actions.c (4) : ... | 8921       | 20     | 2       | 1          |

The total number of messages received is displayed in the title bar. Note that ProTune clears the messages in the Output window at the start of each session execution, or when you reset a session.

---

**Note:** You can limit the number of messages in the Output window, and set a deletion quota for the number of messages that will be overwritten. For more information, see Appendix B, “Working in Expert Mode.”

---

The Output window provides the following information in the Summary tab:

| Column              | Description   |
|---------------------|---|
| TYPE                | the type of message sent: Alert, Error, Notify, Warning, Debug, or Batch (each represented by a different icon)<br>Note: Debug messages will only be sent if you enable the debugging feature in Tools > Options > Debug Information (Expert Mode). Batch messages will be sent instead of message boxes appearing in the Console, if you are using automation. |
| MESSAGE CODE        | the code assigned to all similar messages. The number in parentheses indicates the number of different codes displayed in the Output window.  |
| SAMPLE MESSAGE TEXT | an example of the text of a message with the specified code   |
| TOTAL MESSAGES      | the total number of sent messages with the specified code   |
| VUSERS              | the number of Vusers that generated messages with the specified code  |
| SCRIPTS             | the number of scripts whose execution caused messages with the specified code to be generated   |
| GENERATORS          | the number of load generators from which messages with the specified code were generated  |

You can view and manipulate the log information in the Summary tab using the following utilities:

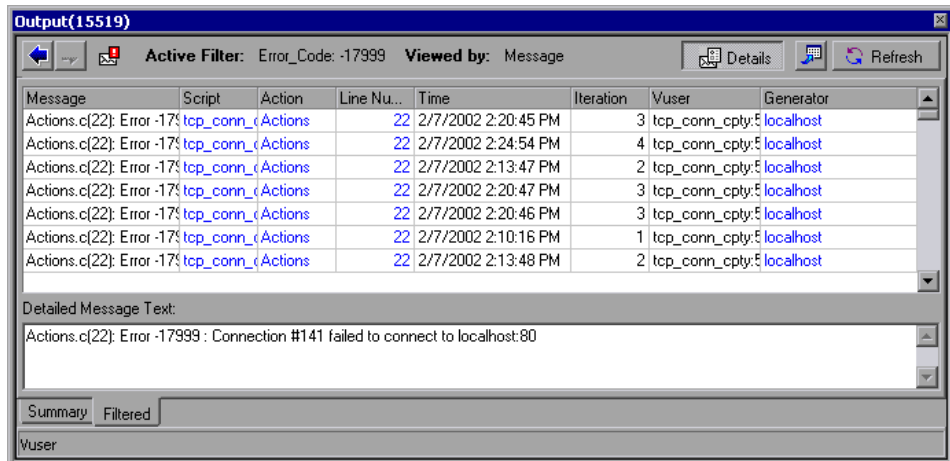
- To show (or hide) the Output window, choose **View > Show Output**.
- To sort the log information, click the appropriate column header. The messages are sorted in descending/ascending order.
- To filter the Output window to display only certain message types, select the type of message you want to view from the Type of Message box. By default, all types of output messages are displayed, unless you click the Show Snapshot button in the Session Status window.

- ▶ To view a message in detail, select the message and click the **Details** button. The Detailed Message Text box opens in the Output window displaying the complete message text.
- ▶ To save the Output window view to a file, click the **Export the view** button.
- ▶ To clear all log information from the Output window, click the **Remove all messages** button.
- ▶ To halt the updating of the Output window, click the **Freeze** button. To instruct ProTune to resume updating the Output window with messages, click the **Resume** button. Note that newly updated log information is displayed in a red frame.

## Viewing Log Information Details

You can view details of each message, Vuser, script, and load generator associated with an error code by clicking the blue link in the respective column. The Output window displays a drilled down view by message, Vuser, script, or load generator in the Filtered tab.

For example, if you drill down on the Vusers column, the Output window displays all the messages with the code you selected, grouped by the Vusers that sent the messages.



Note that the message type, the message code, and the column that you selected to drill down on, are displayed above the grid.



You can drill down further on the entries displayed in blue. Note that when you drill down on a Vuser, the Vuser log opens. When you drill down on a load generator, the Load Generators dialog box opens, displaying the load generator you selected. When you drill down on a script (or Action or Line Number), VuGen opens, displaying the script you selected.

---

**Note:** To limit the number of rows displayed when you drill down, open the *wlrun7.ini* file in any text editor, and located the following line:

MaxOutputUIRowsToShow=0

Change the 0 (no limit) to the number of rows you want to view.

---

When new messages arrive in the Output window, the Refresh button is enabled. Click **Refresh** to add the new log information to the Detailed tab view.



To move between the various drill down levels, click the **Previous view** and **Next view** buttons in the upper left-hand corner of the Output window.

## Viewing the Script Log

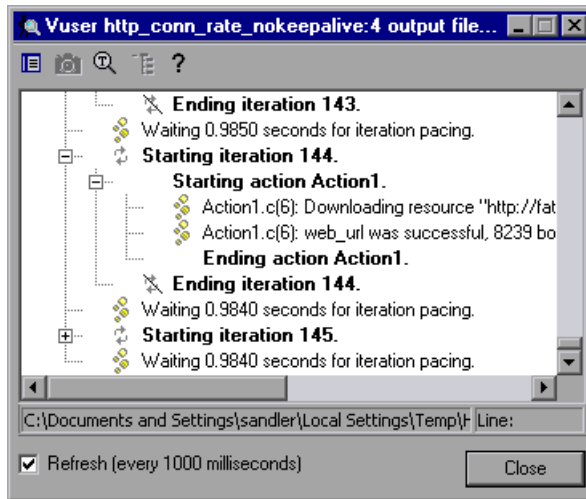
During session execution, you can view a log containing run-time information about each running Vuser.

**To view the script log for a particular Vuser:**

- 1 In the Vusers dialog box, select the Vuser whose log you want to view, and click the **Show Vuser Log** button, or right-click the Vuser and select **Show Vuser Log**.



The script log opens, displaying run-time information about the Vuser that is refreshed, by default, every 1000 milliseconds.



To change the default refresh settings, see “Options - Output Settings” on page 650.

---

**Note:** If you disabled the logging feature in the Run-Time Settings’ Log tab, the script log will be empty. If you selected the **Send messages only when an error occurs** option in the Log tab, the script log will contain output only if there are script errors.

---

- To disable the refreshing of this log, clear the **Refresh** check box.
- To view the information in text format, click the **Show Text View** button. To revert to the tree view, click the button again.
- If you are running a Web Vuser, and want to view a snapshot of the Web page where an error occurred, highlight the error in the Vuser log and click the **Display** button. Note that this option is only available for Vusers running on Windows load generators.

---

**Note:** In order to view a snapshot of the Web page where an error occurred, you must select the **Activate snapshot on error** option in the Internet Protocol > Preferences tab of the Run-Time Settings dialog box before running the session.

---



- To search the Vuser log for specific text, click the **Find Text** button, and enter the text you want to search for in the text box.

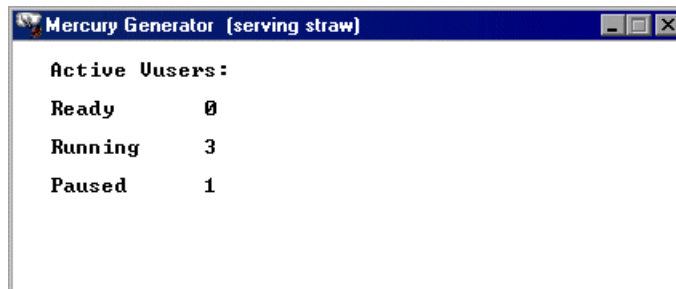


- To collapse the tree view, click the **Collapse Node** button. To revert to the expanded tree view, click the same button again.

- 2 Click **Close** to close the script log.

## Viewing the Agent Summary

When you run a session with non-GUI Vusers, the machine running the Vusers invokes an agent that controls the Vuser execution on that load generator. To view the Agent window during session execution, double-click the Mercury Generator Agent icon on the task bar. ProTune displays a summary of the Ready, Running, and Paused Vusers.





# 11

---

## Viewing the Session Summary

The Session Summary shows the activity that took place during a session, and lets you specify which information should appear in the session report.

This chapter describes:

- The Session Summary Window
- Viewing Step Run Information
- Viewing Tuning Information
- Adding Session Notes
- Deleting Session Summary Entries

### About Viewing the Session Summary

The Session Summary window displays the following information:

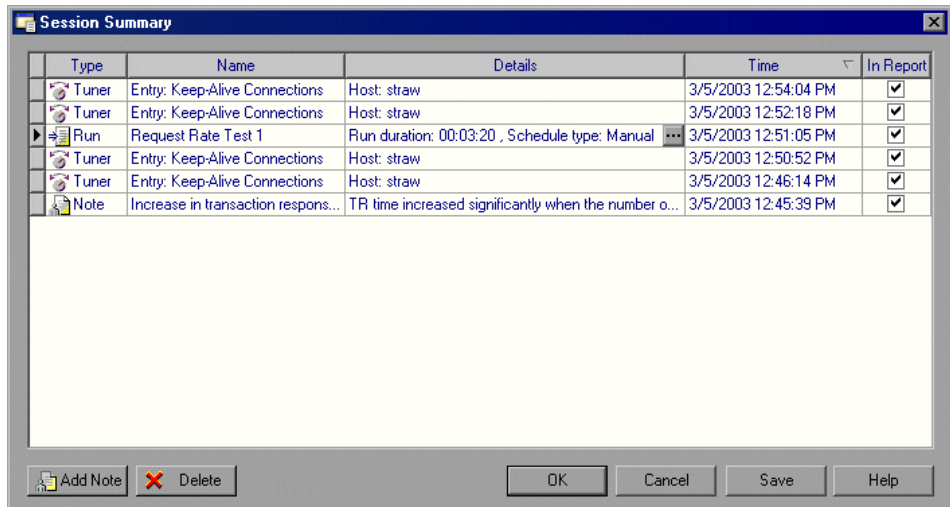
- Details of each step execution performed in a session. You can see when the step started, when it ended, and other information about the step. You can also enter information about the step run.
- Details of each tuning action. You can view the values of a tuned property, before and after tuning.
- Notes entered by the user during the session.

## The Session Summary Window



To view the Session Summary window:

- 1 Click Session Summary, or choose **Session > Summary Information**. The Session Summary window is displayed.



Each row in the table contains an entry of one of the following types:

**Run**—A step that was executed

**Tuner**—A tuning action

**Note**—A note entered by the user

## Viewing Step Run Information

The Session Summary window contains the following information about each step run:

**Name**—The step result's name.

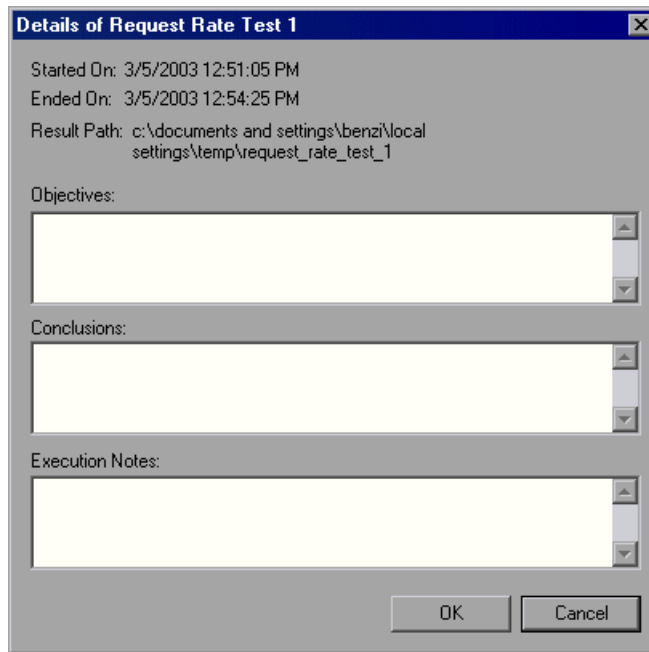
**Details**—Duration of the step run and the schedule type (Manual or Goal-Oriented).

**Time**—When the step began or the property was tuned.

**In Report**—Checking this box causes the step information to be included in the session report.

**To view additional information about a step's execution:**

- 1 Click the step's row to select it.
- 2 Click the browse button in the Details column. The Details dialog box is displayed.



This dialog box contains the following information:

- ▶ **Started On**—When the step started executing
- ▶ **Ended On**—When the step execution ended
- ▶ **Result path**—Path to the directory where the step's results are stored
- ▶ **Objectives, Conclusions and Execution Notes:** You use these fields to enter relevant information about the step's execution.

- 3 Make any necessary changes and click **OK** to save your changes, or **Cancel** to exit without saving.

## Viewing Tuning Information

The Session Summary window contains the following columns of information about tuning actions:

**Name**—Name of the tuned (changed) property

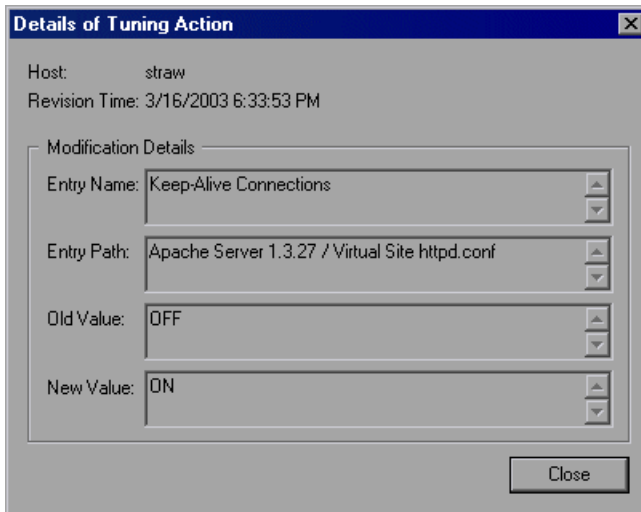
**Details**—Host name or IP address of the tuned host

**In Report**—Checking this box causes the tuning information to be included in the session report.

**Time**—When the host was tuned

**To view additional information about a tuning action:**

- 1 Click the tuning action's row to select it.
- 2 Click the row's browse button. The following window is displayed.





This window displays the host name and revision time (when the property was tuned), as well as the Modification Details section which contains the following information:


- **Entry Name**—Name of the tuned (changed) property
- **Entry Path**—Logical path to the entry (as seen in the Server Configurations tree in the Tune tab)
- **Old Value**—The property's value prior to the change
- **New Value**—The property's value after the change

3 Click **Close** to exit the window.

## Adding Session Notes

You can add session notes in the Session Summary dialog box. These notes can then be included in the Session Report.

**To add a session note:**



- 1 Click **Add Note**. The User Note dialog box appears.
- 2 Enter the subject and content of the note in the relevant fields.
- 3 Click **OK** to save the note. The note appears as an entry in the Session Summary dialog box.

To subsequently view the note, select the entry in the Session Summary dialog box and click the Browse button. ProTune displays the note, along with its creation and modification dates.

## Deleting Session Summary Entries

**To delete an entry in the Session Summary dialog box:**



- Select the entry and click **Delete**.



# 12

---

## Generating a Session Report

ProTune allows you to generate a report of the activity that took place during a session. The report is in Microsoft Word format.

The Session Report includes the information that you define in the Session Summary (see “Viewing the Session Summary,” on page 147).

This chapter describes:

- ▶ Session Summary Section
- ▶ Step Information
- ▶ Generating the Report

### About Generating a Session Report

The Session Report includes the following sections:

- ▶ Session Summary
- ▶ Step information

### Session Summary Section

The Session Summary section of the report contains information relating to the entire session, including the user-specified objectives and the system topology.

## Step Information

Since a ProTune session can contain multiple steps, the report includes a separate section for each step. The section includes the following subsections:

- ▶ Objectives
- ▶ Conclusions
- ▶ Execution Notes
- ▶ Step Summary
- ▶ Scheduler Information
- ▶ Scripts
- ▶ Run Time Settings
- ▶ Transaction Response Times
- ▶ Transaction Response Times Under Load
- ▶ Hardware Utilization - Windows Resources

## Generating the Report

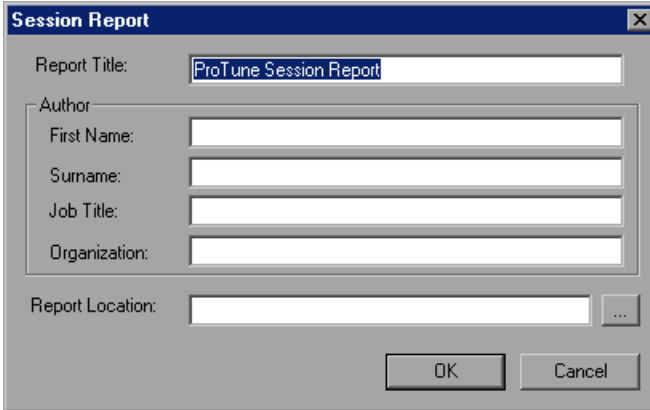
After you've executed all the steps that you want to include in the session, you generate the Session Report as follows:

- 1** In the Session Summary dialog box, choose the items that you want to include in the report, and enter any additional information (objectives, conclusions, and so on). For details, see "Viewing the Session Summary," on page 147.



- 2** Click **Session Report**.

The Session Report dialog box appears.



The image shows a Windows-style dialog box titled "Session Report". The "Report Title" field contains the text "ProTune Session Report". Below it is a section labeled "Author" which contains four sub-fields: "First Name", "Surname", "Job Title", and "Organization", all of which are currently empty. At the bottom of the dialog is a "Report Location" field with a browse button (three dots) to its right. At the very bottom are "OK" and "Cancel" buttons.

- 3** In the Report Title field, enter a name for the report, or use the default name.
- 4** Enter the relevant information in the First Name, Surname, Job Title and Organization fields. (Optional)
- 5** Click the **Browse** button and specify the report's directory and filename in the Report Location field.
- 6** Click **OK**. ProTune generates the report in Word format.

To view the report, locate it in the directory that you specified in the Report Location field, and open it with Microsoft Word.



# 13

---

## Working with Firewalls

You can run Vusers and monitor your servers, while the Console is outside of the firewall.

This chapter describes:

- ▶ Overview of Running or Monitoring over the Firewall
- ▶ Configuring the ProTune Agents in LAN1
- ▶ Configuring the Firewall to Allow Agent Access
- ▶ Installing and Configuring the MI Listener in LAN2
- ▶ Configuring Console to Run or Monitor Vusers over the Firewall

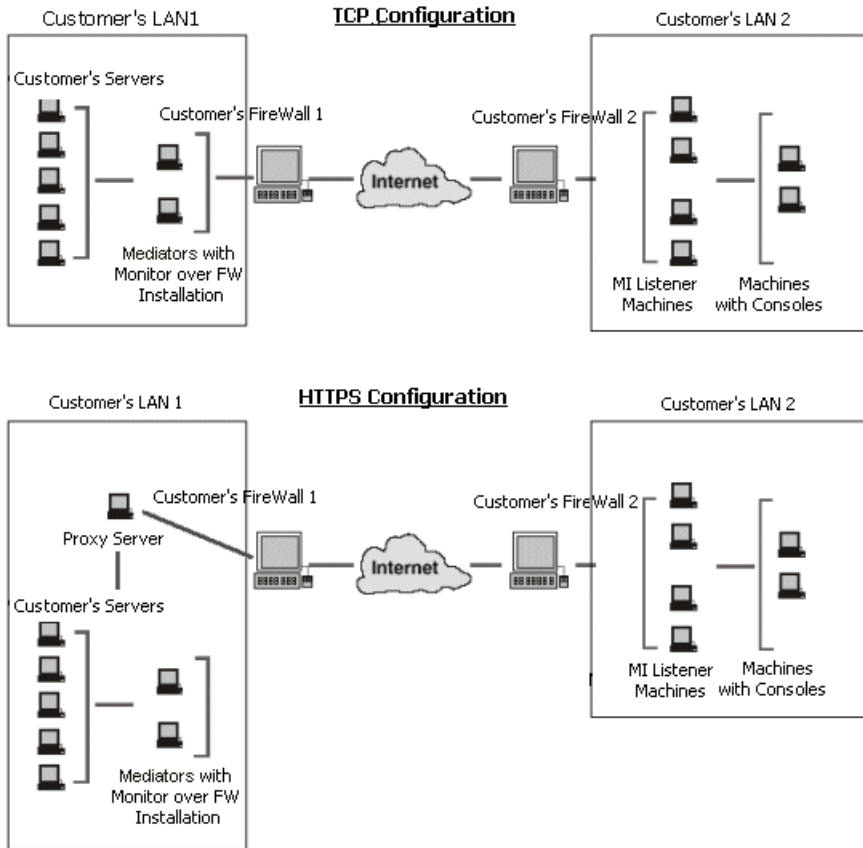
### About Using Firewalls in ProTune

Working with a firewall means that you can prevent access to the outside world and from the outside world, on specific port numbers.

For example, you can specify that there is no access to any port from the outside world, with the exception of the mail port (23), or you can specify that there is no outside connection to any ports except for the mail port and WEB port (80). The port settings are configured by the system administrator.

In a regular ProTune session (not over the firewall), the Console has direct access to the ProTune agents running on remote machines. This enables the Console to connect directly to those machines. However, when running Vusers or monitoring servers over the firewall, this direct connection is blocked due to the firewall. The connection cannot be initiated by the Console, because it does not have permissions to make an opening in the firewall.

Configure your system according to one of the following configurations to run Vusers or monitor servers over the firewall. Note that these configurations contain a firewall on each LAN. There may also be configurations where there is a firewall only for LAN1:



During installation, the ProTune agent is added either as a Windows service or as an executable run from the Startup folder. The MI Listener component serves as a router between the Console and the ProTune agent.

In the above configuration, the MI Listener is on a different machine than the Console. Every ProTune agent can behave as a MI Listener, so you can use the Console machine as the MI Listener also, and you do not need a separate installation.



### **TCP Configuration**

The TCP configuration requires every ProTune agent machine behind the FireWall 1 to be allowed to open a port in the firewall for outgoing communication. If this is the firewall configuration at hand, use the TCP configuration.

### **HTTPS Configuration**

In the HTTPS configuration, only one machine (the proxy server) is allowed to open a port in the firewall. Therefore it is necessary to tunnel all outgoing communications through the proxy server.

After installation, you configure the ProTune agent to operate over the firewall. You also modify firewall settings to enable communication between the agent machine(s) inside the firewall and machines outside the firewall. In addition, you prepare the Console to work over the firewall.

See Chapter 15, “Monitoring over a Firewall” for additional information about configuring ProTune to monitor servers from outside the firewall.

## **Overview of Running or Monitoring over the Firewall**

To prepare for running Vusers or monitoring servers over the firewall, perform the following steps:

- 1 Make sure that the ProTune agent is installed on the machines running Vusers, or on the servers to be monitored behind the firewall.**

The agents can run on Windows or Unix machines. See the diagram, “About Using Firewalls in ProTune,” on page 157.

- 2 Configure the ProTune agent to operate over the firewall.**

Configure the ProTune agent on the machines running Vusers, or agents acting as mediators for the servers to be monitored. See “Configuring the ProTune Agents in LAN1,” on page 160 for instructions.

- 3 Configure the firewall(s).**

Configure the firewall, to allow communication between the agents inside the firewall, and the machines outside the firewall. See “Configuring the Firewall to Allow Agent Access” on page 168.

#### **4 Install the Monitoring over Firewall Component.**

To monitor a server over the firewall, install this component on the machine which sits inside the firewall, and acts as a mediator between the Console, and the monitored server. See the diagram, "About Using Firewalls in ProTune," on page 157 for information about where to install the Monitoring over the Firewall component, and refer to the *ProTune Installation Guide* for installation instructions.

#### **5 Install the MI Listener on a machine outside the firewall.**

See the diagram, "About Using Firewalls in ProTune," on page 157 for information about where to install the MI Listener, and refer to the *ProTune Installation Guide* for installation instructions.

#### **6 Configure the MI Listener machines.**

Configure the security attributes on each MI Listener machine. See "Installing and Configuring the MI Listener in LAN2," on page 169.

#### **7 Configure the Console machine.**

Configure the Console machine to recognize the agent and MI Listener machines. See "Configuring Console to Run or Monitor Vusers over the Firewall," on page 170.

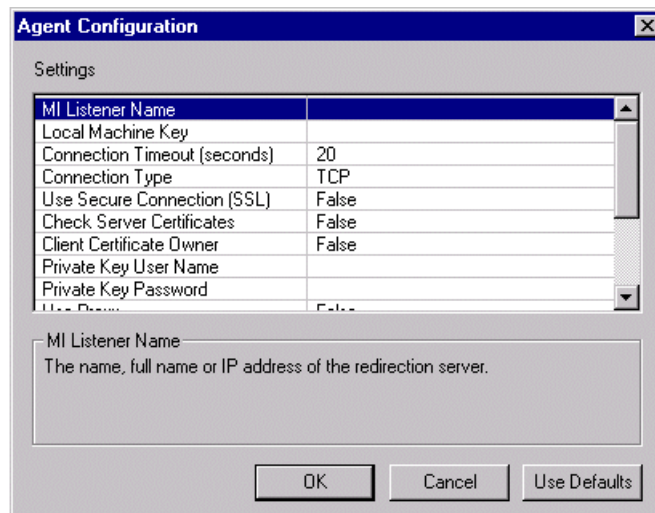
## **Configuring the ProTune Agents in LAN1**

The machines within LAN1 can either be Load Generator machines running Vusers, or mediator machines connected to the servers to be monitored by the Console. You configure the ProTune agents in LAN1 to operate over firewall. The Console machine resides outside the firewall, and LAN1 is inside the firewall.

## Configuring and Running the Windows ProTune Agent

To configure the ProTune agents on Windows machines:

- 1 Stop the ProTune agent by right-clicking its icon in the system tray and selecting **Close**.
- 2 Select **Agent Settings** from Start > Programs > ProTune > Advanced Settings (or open `<ProTune root>\launch_service\dat\br_Inch_server.cfg` in a text editor).
- 3 In the Firewall section set `FireWallServiceActive` to 1, and save your changes.
- 4 Run **Agent Configuration** from Start > Programs > ProTune > Advanced Settings, or run `<ProTune root>\launch_service\bin\AgentConfig.exe`.



- 5 Set each option as described in “Agent Configuration Settings” on page 166.
- 6 Click **OK** to save your changes, **Cancel** to cancel them, or **Use Defaults**.
- 7 Restart the ProTune agent by double-clicking the shortcut on the desktop, or from Start > Programs > ProTune > ProTune Agent Service/Process.

## Configuring and Running the UNIX ProTune Agent

To configure the ProTune agents on UNIX machines:

- 1 Open `<ProTune root folder>/dat/br_lrch_server.cfg` in a text editor.
- 2 In the Firewall section, set `FireWallServiceActive` to 1 and save your changes.
- 3 Run `agent_config` from the `<ProTune root folder>/bin` directory to display the following menu:

```
Menu:
1. Show current settings.
2. Change a setting.
3. Save changes and exit.
4. Exit without saving.
5. Use default values.
```

- 4 Enter 1 to display the current settings:

```
Settings:
-----
1. MI Listener Name =
2. Local Machine Key =
3. Connection Timeout (seconds) = 20
4. Connection Type = TCP
5. Use Secure Connection (SSL) = False
6. Check Server Certificates = False
7. Client Certificate Owner = False
8. Private Key User Name =
9. Private Key Password =
10. Proxy Name =
11. Proxy Port =
12. Proxy User Name =
13. Proxy Password =
14. Proxy Domain =

Menu:
1. Show current settings.
2. Change a setting.
3. Save changes and exit.
4. Exit without saving.
5. Use default values.
```

- 5 To change a setting, enter 2 to display the settings menu:

```
Settings:
-----
1. MI Listener Name =
2. Local Machine Key =
3. Connection Timeout (seconds) = 20
4. Connection Type = TCP
5. Use Secure Connection (SSL) = False
6. Check Server Certificates = False
7. Client Certificate Owner = False
8. Private Key User Name =
9. Private Key Password =
10. Proxy Name =
11. Proxy Port =
12. Proxy User Name =
13. Proxy Password =
14. Proxy Domain =

Enter number of setting to change or 0 to go back to menu.
```

Enter the setting and continue according to the menu instructions. Set each option according to the “Agent Configuration Settings” on page 166.

## Examples of Changing Agent Settings in Unix

To change the MI Listener Name:

- 1 Enter 1 in the Settings menu to display the following screen:

```
MI Listener Name - The name, full name or IP address of the redirection server.
Old value =
Enter new MI Listener Name.
```

Line one is a description of the setting. Line two shows the current value of the setting.

- 2 Enter the new value, (For example, 'bunji') to display the following:

```
MI Listener Name - The name, full name or IP address of the redirection server.
Old value =
Enter new MI Listener Name.
bunji
Change MI Listener Name from "" to "bunji"? 1.OK 2.CANCEL 3.FIX
```

- 3 To keep the new value and return to the menu, enter '1'.  
To discard the new value and return to the menu, enter '2'.  
To discard the new value and change the setting once more, enter '3'.

**To change the Connection Type:**

- 1 Enter 4 in the Settings menu to display the following screen:

```
· xterm
Connection Type - The connection type: TCP or HTTP.
Old value = TCP
Enter number for new Connection Type: 1.TCP 2.HTTP 3.CANCEL
```

Line one is a description of the setting. Line two shows the current value of the setting.

- 2 Enter 1 to set the connection type to TCP, or enter 2 to set it to HTTP and display the following:


```
· xterm
Connection Type - The connection type: TCP or HTTP.
Old value = TCP
Enter number for new Connection Type: 1.TCP 2.HTTP 3.CANCEL
2
Change Connection Type from "TCP" to "HTTP"? 1.OK 2.CANCEL
```

- 3 To keep the new value and return to the menu, enter '1'.  
To discard the new value and return to the menu, enter '2'.

## Viewing the Settings and Restarting the Agent

To view the current settings:

- 1 Return to the main menu by entering 1.
- 2 Enter 1 to display the settings. The following example includes the new settings for MI Listener Name and Connection Type:



```

xterm
Settings:
-----
1. MI Listener Name = bunji
2. Local Machine Key = gumbi
3. Connection Timeout (seconds) = 20
4. Connection Type = HTTP
5. Use Secure Connection (SSL) = False
6. Check Server Certificates = False
7. Client Certificate Owner = False
8. Private Key User Name =
9. Private Key Password =
10. Proxy Name =
11. Proxy Port =
12. Proxy User Name =
13. Proxy Password =
14. Proxy Domain =

Menu:
1. Show current settings.
2. Change a setting.
3. Save changes and exit.
4. Exit without saving.
5. Use default values.

```

- 3 To save your changes, enter 3 from the main menu.

To cancel your changes, enter 4.

To use the default values supplied by ProTune (as described in “Agent Configuration Settings,” on page 166), enter 5.

**To start or remove the ProTune agent:**

- 1** To start the ProTune agent, run the command 'm\_daemon\_setup -install' from the <ProTune root folder>/bin directory.
- 2** To remove the ProTune agent, run the command 'm\_daemon\_setup -remove' from the <ProTune root folder>/bin directory.

For more information about running the ProTune agent, see "UNIX Shell" in Appendix A, "Troubleshooting the Console."

**Agent Configuration Settings**

| Option                              | Default Value | Description   |
|-------------------------------------|---------------|---|
| <i>MI Listener name</i>             | none          | The name, full name or IP address of the Mercury Interactive listener machine, MI Listener.   |
| <i>Local Machine Key</i>            | none          | A string identifier used to establish a unique connection between the Console host and the agent machine, via the MI Listener machine.                                    |
| <i>Connection Timeout (seconds)</i> | 20 seconds    | The length of time you want the agent to wait before retrying to connect to the MI Listener machine. If zero, the connection is kept open from the time the agent is run. |
| <i>Connection Type</i>              | TCP           | Choose either TCP (default) or HTTP, depending on the configuration you are using.  |
| <i>Use Secure Connection (SSL)</i>  | False         | Choose True to connect using the Secure Sockets Layer protocol.   |



| Option                           | Default Value                  | Description  |
|----------------------------------|--------------------------------|--|
| <i>Check Server Certificates</i> | False                          | Choose True to authenticate SSL certificates that are sent by servers. This option is relevant only if the <b>Use Secure Connection</b> option is set to <b>True</b> .   |
| <i>Client Certificate Owner</i>  | False                          | Choose True to load the SSL certificate. In some cases the server requests a certificate to allow the connection to be made. This option is relevant only if the <b>Use Secure Connection</b> option is set to <b>True</b> . |
| <i>Private Key User Name</i>     | None                           | The user name that may be required during the SSL certificate authentication process. This option is relevant only if the <b>Client Certificate Owner</b> option is set to <b>True</b> .                                     |
| <i>Private Key Password</i>      | None                           | The password that may be required during the SSL certificate authentication process. This option is relevant only if the <b>Client Certificate Owner</b> option is set to <b>True</b> .                                      |
| <i>Proxy Name</i>                | <IE proxy server name> or None | The name of the proxy server. This option is mandatory if the <b>Connection Type</b> option is set to <b>HTTP</b> .  |
| <i>Proxy Port</i>                | <IE proxy server port> or None | The proxy server connection port. This option is mandatory if the <b>Connection Type</b> option is set to <b>HTTP</b> .  |
| <i>Proxy User Name</i>           | None                           | The username of a user with connection rights to the proxy server.   |

| Option                | Default Value | Description   |
|-----------------------|---------------|---|
| <i>Proxy Password</i> | None          | The user's password.  |
| <i>Proxy Domain</i>   | None          | The user's domain if defined in the proxy server configuration. This option is required only if NTLM is used. |

## Configuring the Firewall to Allow Agent Access

You modify your firewall settings to enable communication between the machine(s) inside the firewall and machines outside the firewall.

### TCP configuration:

The ProTune agent tries to establish a connection with MI Listener using port 443 at an interval of seconds specified in the Connection Timeout field in the agent configuration. To enable this connection, allow an outgoing connection for HTTPS service on the FireWall for port 443. As a result, the agent connects to MI Listener and MI Listener connects back to the agent. From this point on, the agent listens to commands from the MI Listener.

### HTTPS configuration:

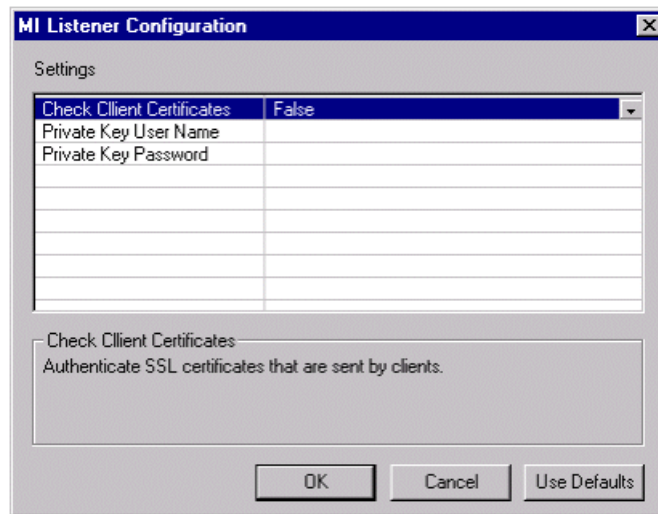
The ProTune agent tries to establish a connection with MI Listener using the proxy port specified in the Proxy Port field and at an interval of seconds specified in the Connection Timeout field in the agent configuration. On successful connection, the proxy server connects to MI Listener. To enable this connection, allow an outgoing connection for HTTPS service on the FireWall for port 443. As a result, the proxy server connects to MI Listener and MI Listener connects back to the agent through the proxy server. From this point on, the agent listens to commands from the MI Listener.

## Installing and Configuring the MI Listener in LAN2

To enable running Vusers or monitoring over a firewall, you need to install MI Listener on one or more machines in your LAN2. For instructions, refer to the *ProTune Installation Guide*.

**To configure the MI Listener:**

- 1 Open incoming HTTPS service for port 443.
- 2 Stop the ProTune agent by right-clicking its icon in the system tray and selecting **Close** from the popup menu.
- 3 Run **MI Listener Configuration** from Start > Programs > ProTune > Advanced Settings, or run `<ProTune_root_dir>\launch_service\bin\MILsnConfig.exe`.



- 4 Set each option as described in “MI Listener Configuration Settings,” on page 170.
- 5 Click **OK** to save your changes, **Cancel** to cancel them, or **Use Defaults**.
- 6 Restart the ProTune agent by double-clicking the shortcut on the desktop, or running it from Start > Programs > ProTune.
- 7 Make sure that port 433 is free on the MI Listener machine.

## MI Listener Configuration Settings

| Option                           | Default Value | Description  |
|----------------------------------|---------------|--|
| <i>Check Client Certificates</i> | False         | Choose True to request that the client send an SSL certificate when connecting, and to authenticate the certificate. |
| <i>Private Key User Name</i>     | None          | The user name that may be required during the SSL certificate authentication process.                                |
| <i>Private Key Password</i>      | None          | The password that may be required during the SSL certificate authentication process.                                 |

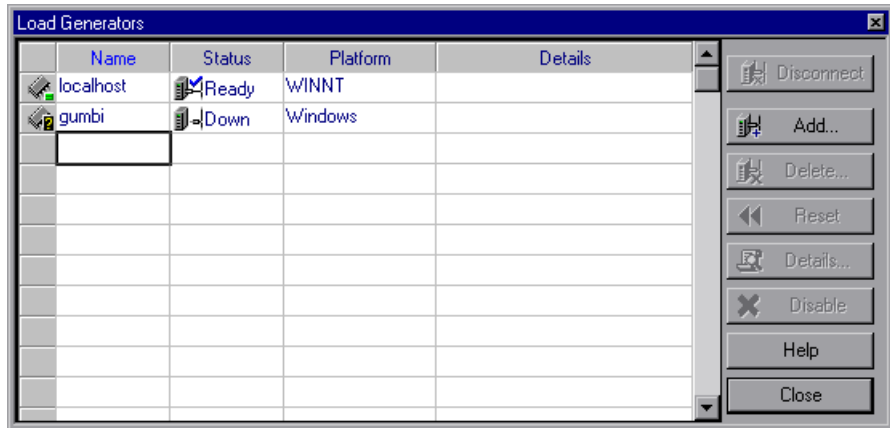
## Configuring Console to Run or Monitor Vusers over the Firewall

In order to obtain information for the monitors configured inside the firewall, or to run Vusers inside the firewall, you create a unique connection between the Console and the agent machine, via the Mercury Interactive listener machine, MI Listener. You establish this connection by defining the agent machine as a load generator.

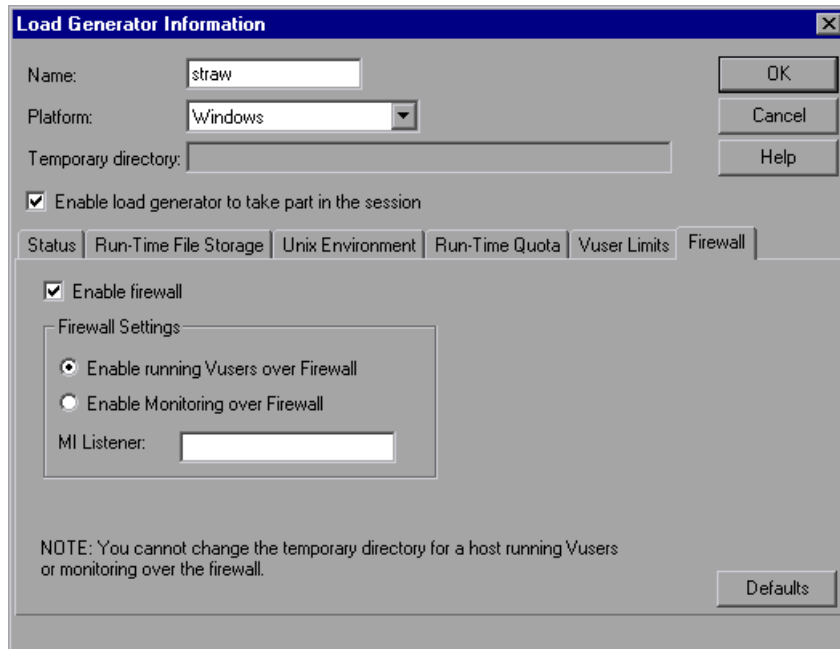
**To configure the Console for running vusers or monitoring over the firewall:**

- 1 Run the Console from Start > Programs > ProTune and create a new session, or load an existing one.
- 2 Click **Generators** to display the Load Generators window. In the Name field, enter the symbolic name of the server. This is the same name that you entered in the Local Machine Key setting in the Agent Configuration. In the example below, the server name is *gumbi*.

If the server is a UNIX server, change the Platform field to *UNIX*.



- 3 Select the Load Generator, and click **Details** to display the Load Generator Information.



- 4** In the Firewall tab, enter the MI Listener machine's name in the **MI Listener** field. This is the same name that you entered in the Agent Configuration, in the MI Listener Name setting.
- 5** To run Vusers over the Firewall, select the **Enable running Vusers over Firewall** option.
- 6** To monitor over the firewall, select the **Enable Monitoring over Firewall** option.
- 7** Click **OK** to return to the Load Generators dialog box.
- 8** Select the Load Generator and click **Connect**.

---

**Note:** Remember that you cannot change the temporary directory on the host running Vusers over the firewall or monitoring over the firewall.

---

# Part V

---

## Monitoring a Session





# 14

---

## Online Monitoring

You can monitor session step execution using the ProTune online run-time, transaction, Web resource, system resource, network delay, firewall server resource, Web server resource, Web application server resource, database server resource, streaming media resource, ERP server resource, Java performance, application deployment solutions, middleware performance, and application traffic management monitors.

The specific monitors are discussed in the next few chapters. This chapter describes the online monitor user interface:

- Choosing Monitors and Measurements
- Viewing the Monitors
- Opening Online Monitor Graphs
- Customizing the Graph Display View
- Configuring Online Monitors
- Setting Monitor Options
- Configuring Online Graphs
- Merging Graphs
- Understanding Online Monitor Graphs
- Configuring Online Measurements
- Exporting Online Monitor Graphs
- Viewing Data Offline

## About Online Monitoring

ProTune provides the following online monitors:

The **Run-Time** monitor displays the number and status of Vusers participating in the session step, as well as the number and types of errors that the Vusers generate. It also provides the User-Defined Data Point graph that displays the real-time values for user-defined points in a Vuser script.

The **Transaction** monitor displays the transaction rate and response time during session step execution. For more information, see Chapter 16, "Run-Time and Transaction Monitoring."

The **Web Resource** monitor measures statistics at the Web server(s) during session step runs. It provides information about the number of Web connections, throughput volume, HTTP responses, server retries, and downloaded pages during the session step. For more information on the Web Resource monitor, see Chapter 17, "Web Resource Monitoring."

The **System Resource** monitors gauge the Windows, UNIX, TUXEDO, SNMP, Antara FlameThrower, and SiteScope resources used during a session step. To activate the System Resource monitors, you must set the monitor options before you run your session step. For information on setting these options, see Chapter 18, "System Resource Monitoring."

The **Network Delay** monitor displays information about the network delays on your system. To activate the Network Delay monitor, you must set up the network paths to monitor before you run your session step. For more information see Chapter 19, "Network Monitoring."

The **Firewall** monitor measures statistics at the firewall servers during the session step. To activate the Firewall monitor, you must set up a list of resources to monitor before you run your session step. For more information, see Chapter 20, "Firewall Server Performance Monitoring."

The **Web Server Resource** monitors measure statistics at the Apache, Microsoft IIS, iPlanet (SNMP) and iPlanet/Netscape Web servers during the session step. To activate the Web Server Resource monitors, you must set up a list of resources to monitor before you run your session step. For more information, see Chapter 21, "Web Server Resource Monitoring."

The **Web Application Server Resource** monitors measure statistics at the Web application server(s) during the session step. To activate the Web Application Server Resource monitors, you must set up a list of resources to monitor before you run your session step. For more information, see Chapter 22, “Web Application Server Resource Monitoring.”

The **Database Server Resource** monitors measure statistics related to the SQL server, Oracle, Sybase, and DB2 databases. To activate the Database Server Resource monitors, you must set up a list of measurements to monitor before you run your session step. For more information, see Chapter 23, “Database Resource Monitoring.”

The **Streaming Media** monitors measure statistics at the Windows Media Server and RealPlayer audio/video servers, as well as the RealPlayer client. To activate the Streaming Media monitors, you must set up a list of resources to monitor before you run your session step. For more information, see Chapter 24, “Streaming Media Monitoring.”

The **ERP Server Resource** monitor measures statistics at the ERP servers during the session step. To activate the ERP Server Resource monitor, you must set up a list of resources to monitor before you run your session step. For more information, see Chapter 25, “ERP/CRM Server Resource Monitoring.”

The **Java Performance** monitors measure statistics of Java 2 Platform, Enterprise Edition (J2EE) objects, Enterprise Java Bean (EJB) objects and Java-based applications, using J2EE, EJB, JProbe, and Sitraka JMonitor machines. To activate the Java Performance monitors, you must set up lists of resources to monitor before you run your session step. For more information, see Chapter 26, “Java Performance Monitoring”, and Chapter 27, “J2EE Performance Monitoring.”

The **Application Deployment Solutions** monitor measures statistics of the Citrix MetaFrame XP and 1.8 servers during a session step run. To activate the Application Deployment Solutions monitor, you must set the monitor options before you run your session step. For information on setting these options, see Chapter 28, “Application Deployment Solution Monitoring.”

The **Middleware Performance** monitors measure statistics of the TUXEDO and IBM WebSphere MQ servers during a session step run. To activate the Middleware Performance monitors, you must set the monitor options before you run your session step. For information on setting these options, see Chapter 29, “Middleware Performance Monitoring.”

The **Application Traffic Management** monitor measures statistics of the F5 BIG-IP server during a session step run. To activate the Application Traffic Management monitor, you must set the monitor options before you run your session step. For information on setting these options, see Chapter 30, “Application Traffic Management.”

All of the monitors allow you to view a summary of the collected data at the conclusion of the session step. Using ProTune Analysis, you can generate a graph for any of the monitors. For more information, refer to the *ProTune Analysis User's Guide*.

---

**Note:** For a detailed list of ProTune's monitors, see Mercury Interactive's Web site ([http://www-heva.mercuryinteractive.com/resources/library/technical/loadtesting\\_monitors/supported.html](http://www-heva.mercuryinteractive.com/resources/library/technical/loadtesting_monitors/supported.html)).

---

## Choosing Monitors and Measurements

You can select the measurements to monitor for each of your servers via the Monitors button on the main toolbar or the Element Monitors tab in the System Topology window.

Note that you select the measurements to monitor the topology elements—not to monitor the physical hosts. If you map the same physical host to more than one topology element, you will typically be interested in monitoring those measurements that are relevant to each element. For example, if the physical host is mapped to a Web server element and a database server element, you can monitor the Web-related measurements on the Web server, and the database-related ones on the database server.

The relevant measurements will appear on the graphs belonging to the individual topology elements.

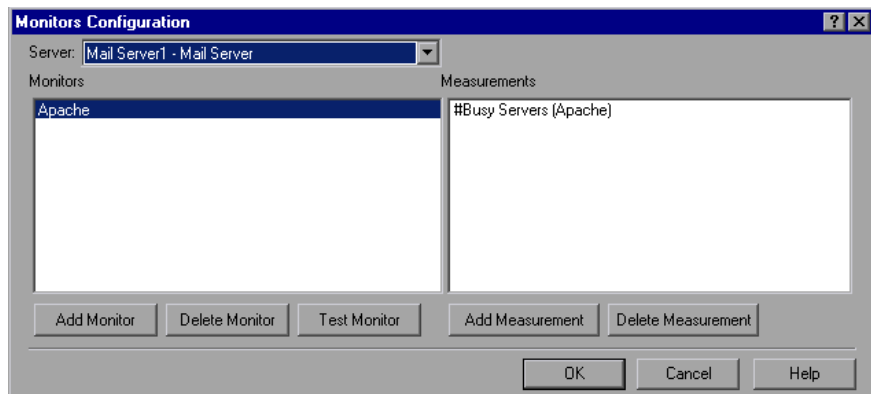
**To select measurements to monitor via the System Topology window:**

- Follow the directions in “Selecting Monitors,” on page 18.

**To select measurements to monitor via the Monitors button on the main toolbar:**



- 1 Click the **Monitors** button. The **Monitors Configuration** dialog box is displayed:



- 2 Select the server whose monitors you want to configure from the list box. The monitors that are currently assigned to monitor the specified server are displayed in the **Monitors** pane. When you click on a monitor in the **Monitors** pane, the measurements that have been specified to be monitored by that monitor are listed in the **Measurements** pane.
- 3 To add measurements to monitor, click **Add Monitor**. The **Select Measurements to Monitor** dialog box is displayed. Choose the monitor and the measurements for the specific server.

---

**Note:** For detailed instructions on using the **Select Measurements to Monitor** dialog box, see “Selecting Monitors,” on page 18.

---

- 4** To delete a monitor from the **Monitors** pane, select the monitor and click **Delete Monitor**.
- 5** To add a measurement to the list in the **Measurements** pane, click **Add Measurement**.
- 6** To delete a measurement from the list in the **Measurements** pane, select the measurement and click **Delete Measurement**.
- 7** For some monitors, a **Test Monitor** button appears in the **Monitors Configuration** dialog box. Click this button to test whether you can access the monitor.
- 8** Click **OK** to save your configuration.

## Viewing the Monitors

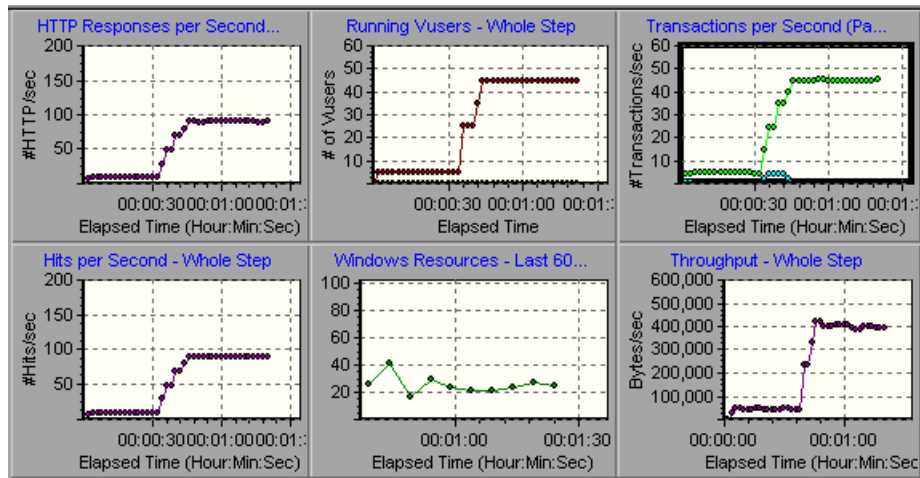
You use the online monitors to monitor Vuser status, errors, transactions, system resources, Web resources, network delay, firewall server resources, Web server resources, Web application server resources, database server resources, streaming media resources, ERP server resources, and Java performance.

The online monitors start operating as soon as you assign them to an element, even before you start a session step. The measurements from each monitor are displayed on a separate graph. **Note:** The Runtime, Transaction and Web Resource graph groups are session step dependent: they start displaying measurements only after you start a session step.

**To view the online monitors and session step dependent graphs:**

- 1** Start the session step by clicking the **Execute** button. Alternatively, you can choose **Session > Start Time > Session Step Start**, specify the start time, and click **OK**.

- 2 Click the **Execute** tab. The default graphs are displayed below the Session Step Groups window.



- 3 Double-click a graph to maximize it. Repeat the operation to restore the tiled view.
- 4 If the graph tree is not displayed, select **View > Show Available Graphs**. Click the "+" in the left pane to expand the graph tree. To hide the graph tree view, select **View > Hide Available Graphs**, or click the X button in the right-hand corner of the Available Graphs list.
- 5 Select a graph from the tree and drag it into the right pane. You can also drag graphs between panes.

---

**Note:** The Transaction Monitor graphs will not contain any data unless transactions are being executed. In addition, the System Resource, Network, Firewall, Web Server, Web Application Server, Database, Streaming Media, ERP Resource, and Java Performance graphs will not contain any data unless you set up a list of resources to monitor before running your session step.

---

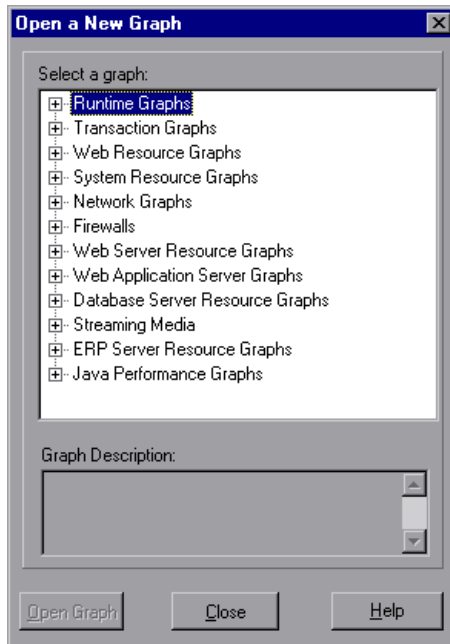
## Opening Online Monitor Graphs

By default, ProTune displays four graphs in the Execute view: Running Users, Transaction Response Time, Hits per Second, and Windows Resources. You can display the other graphs by using one of the following methods:

- ▶ Clicking graphs in the graph tree view and dragging them to the graph view area
- ▶ Opening a new graph using the Open a New Graph dialog box
- ▶ Using the Select Online Graphs dialog box

**To open a new graph using the Open a New Graph dialog box:**


- 1** Select **Monitors > Online Graphs > Add New Graph**, or right-click a graph and select **Open a New Graph**. The Open a New Graph dialog box opens.



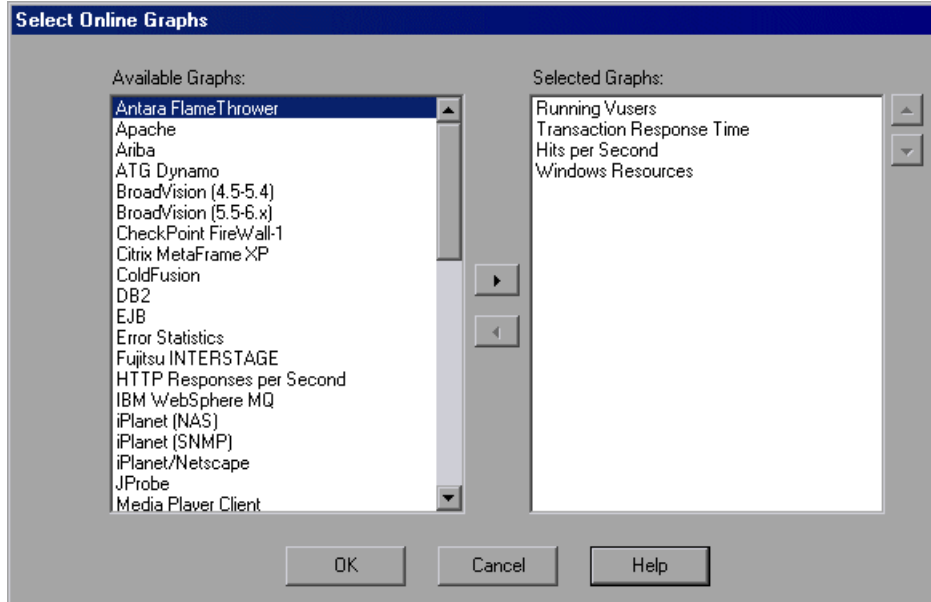
- 2** Click the "+" in the left pane to expand the graph tree, and select a graph. You can view a description of the graph in the **Graph Description** box.
- 3** Click **Open Graph**. The graph appears in the graph view area.



### To specify the graphs via the Select Online Graphs dialog box:



- 1 Click the **Graphs** button in the toolbar. The Select Online Graphs dialog box opens.



The Selected Graphs pane shows the graphs that are currently selected to be displayed in the Execute view. If you have not specified any graphs, the pane lists the four default graphs.

- 2 To add a graph to the list of graphs to display, select the graph in the Available Graphs pane and click the right-arrow. The graph is added to the Selected Graphs pane.
- 3 To remove a graph from the list of graphs to display, select the graph in the Selected Graphs pane and click the left-arrow. The graph name is removed from the list in the Selected Graphs pane.
- 4 You can also determine the position of a graph in the Execute view. The graph at the top of the list in the Selected Graphs pane will be displayed in the top row in the leftmost position. The graph at the bottom of the list will be displayed in the bottom row in the rightmost position.

To change a graph's position in the Selected Graphs pane, use the arrows to the right of the pane.

- 5 Click **OK**. In the **Execute** tab, the selected graphs appear in the order that you specified.

## Customizing the Graph Display View

ProTune lets you display up to 16 online monitor graphs simultaneously.

**To customize your online graph display:**

Click **View Graphs** and select the number of graphs you want to view. You can choose from **Show One Graph**, **Show Two Graphs**, **Show Four Graphs**, **Show Eight Graphs**, or **Custom Number**. If you select **Custom Number**, enter the number of graphs you want to view in the View Graphs dialog box, and click **OK**. The number of graphs selected open in the graph view area.

To display only one graph, double-click the graph pane. To return to the previous view, double-click the graph again.

## Configuring Online Monitors

ProTune lets you configure the settings for your online monitors. You can set graph measurements and properties, such as the sampling time, the colors of the lines, and the scale of the graph.

**Monitor options:** global sampling rate, error handling, debugging, and the frequency settings. For more information, see "Setting Monitor Options" on page 185.

**Graph properties:** refresh rate, display type, graph time for the x-axis, and the y-axis scale. For more information, see "Configuring Online Graphs" on page 187.

**Measurement settings:** line color, scale of the y-axis, and whether to show or hide the line. For more information, see "Configuring Online Measurements" on page 194.

When you save a session step, the online monitor configuration settings are saved as well.

## Setting Monitor Options

Before running your session step, you can set monitor options in the following areas:

- ▶ **Sampling Rate:** The sampling rate is the period of time (in seconds) between consecutive samples. By default, the online monitor samples the data at intervals of three seconds. If you increase the sampling rate, the data is monitored less frequently. This setting applies to all graphs. To set a sampling rate for a specific graph, see “Configuring Online Graphs” on page 187.

The sampling rate you set is applied to all server monitors that you subsequently activate. It is not applied to server monitors that have already been activated. To apply the new sampling rate to activated server monitors, save your session step and reopen it.

---

**Note:** Each monitor has a different minimum sampling rate. If the default sampling rate, or the rate set in the Options Monitors tab is less than a monitor’s minimum sampling rate, the monitor will sample data at its minimum sampling rate. For example, the minimum sampling rate for the Oracle Monitor is 10 seconds. If the sampling rate in the Options Monitors tab is set at less than 10 seconds, the Oracle Monitor will continue to monitor data at 10 second intervals.

---

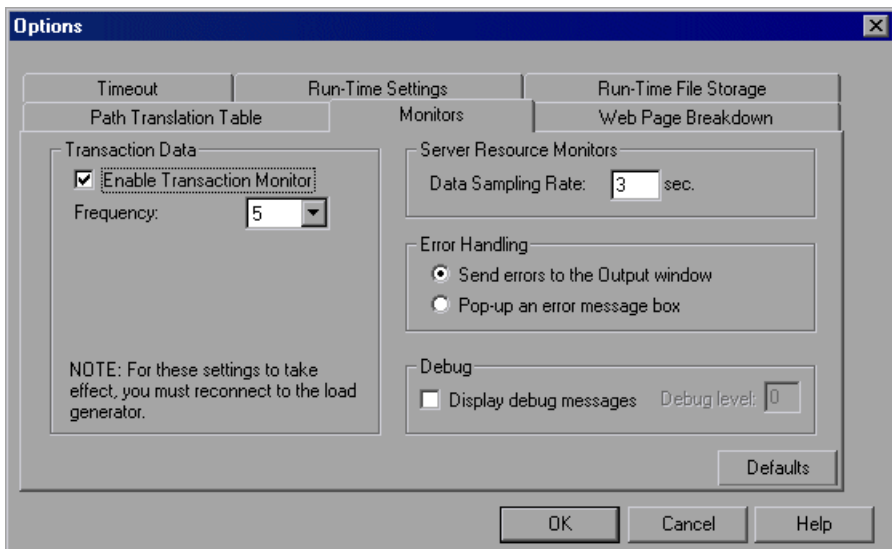
- ▶ **Error Handling:** You indicate how ProTune should behave when a monitor error occurs—issue a popup message box or send error messages to the Output window (default).
- ▶ **Debug:** The online monitor provides debugging capabilities. You can display the debug messages in the output log. For the Network monitor, you can indicate the debug (detail) level of messages sent to the log, ranging from 1-9.

- **Frequency:** You set the frequency at which the monitor sends updates to the Console for the Transaction, Data Point, and Web Resource graphs. The data is averaged for the frequency period defined, and only one value is sent to the Console.

For information on enabling and disabling the Transaction monitor and Web page breakdown, see Chapter 16, “Run-Time and Transaction Monitoring.”

**To set monitor options:**

- 1 Select **Tools > Options** and select the **Monitors** tab.



- 2 Specify the frequency at which the monitor should send updates to the Console for the Transaction, Data Point, and Web Resource graphs. The default value is 5 seconds. For a small session step, it is recommended that you use a frequency of 1. For a large session step, it is recommended that you use a frequency of 3-5. The higher the frequency, the less network traffic there will be.

---

**Note:** You cannot modify these settings during session step execution; you must stop the session step before disabling the monitor or changing its frequency.

---

- 3** Enter a sampling rate.
- 4** Set the desired **Error Handling** option.
- 5** To display debug messages in the Output window, select the **Display Debug Messages** check box. For the Network monitor, specify a **Debug level** from 1-9.
- 6** Click **OK** to save your settings and close the Options dialog box.

You can configure an additional monitor setting while working in Expert mode. For information on working in Expert mode, see Appendix B, “Working in Expert Mode.”

## Configuring Online Graphs

You can customize your graph in the following areas:

- Refresh Rate
- X-Axis Style
- Graph Time
- Display Type
- Y-Axis Style
- Network Delay View

Note that these settings can be set globally—to apply to all graphs—or per graph.

## Refresh Rate

The refresh rate is the interval in which the graph is refreshed with new data. By default, the graph is refreshed every five seconds. If you increase the refresh rate, the data is refreshed less frequently.


---

**Note:** In a large load test, it is recommended to use a refresh rate of three to five seconds. This enables you to avoid problems with CPU resource usage.


---

## X-Axis Style

You can specify how the graph displays the x-axis time: *Don't Show*, *Clock Time*, or *Relative to Session Start*. The *Don't Show* setting instructs ProTune not to display values for the x-axis. The *Clock Time* setting displays the absolute time, based on the system clock. The *Relative to Session Start* setting displays the time relative to the beginning of the session step. In the following example, the graph is shown with the *Don't Show* and *Clock Time* options:



*Don't Show*



*Clock Time*

## Graph Time

The Graph Time settings indicate the scale for a graph's x-axis when it is time-based. A graph can show 60 or 3600 seconds of activity. To see the graph in greater detail, decrease the graph time. To view the performance over a longer period of time, increase the graph time. The available graph times are: *Whole Session*, *60*, *180*, *600*, and *3600* seconds.

## Display Type

You can specify whether ProTune displays the Network Delay Time graph as a line, pie, or area graph. By default, the graph is displayed as a line graph. Note that all other graphs can only be displayed as line graphs.

## Y-Axis Style

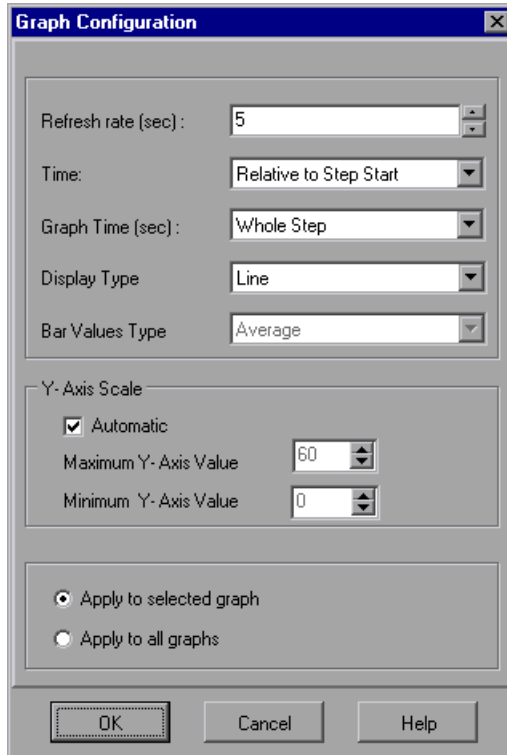
You can instruct ProTune to display graphs using the default y-axis scale, or you can specify a different y-axis scale. Click **Automatic** if you want ProTune to use the default y-axis values. Specify a maximum or minimum value for the y-axis if you want to modify the y-axis scale.

## Network Delay View

This option only appears when you configure the Network Delay Time graph. Click **SubPaths** to view the delay measurements from the source machine to each of the nodes along the network path. Click **DNS name** to view the DNS names of the measurements displayed in the legend.

**To customize your graphs:**

- 1 Select the online graph you want to configure (in either the right or left pane) and choose **Monitors > Online Graphs > Configure**. Alternatively, right-click a graph and select **Configure**. The Graph Configuration dialog box opens.





- 2** To apply the dialog box settings to all graphs, select **Apply to all graphs**.
- 3** Enter the desired refresh rate—the time between graph updates—in the Refresh Rate box.
- 4** Select a style for the x-axis from the Time box.
- 5** Select a value from the Graph Time box. The graph time is the time in seconds displayed by the x-axis.
- 6** For the Network Delay Time graph, select a graph style—Line, Pie, or Area—from the Display Type box.
- 7** If the selected display type is Bar, choose a value from the Bar Values Type box. This determines the type of value that will be displayed in the bar graph. You can choose between **Average**, **Last Value**, **Minimum** and **Maximum**.
- 8** Select a maximum or minimum value for the y-axis, or choose **Automatic** to view graphs using the default y-axis scale.
- 9** Click **OK** to save your settings and close the Graph Configuration dialog box.

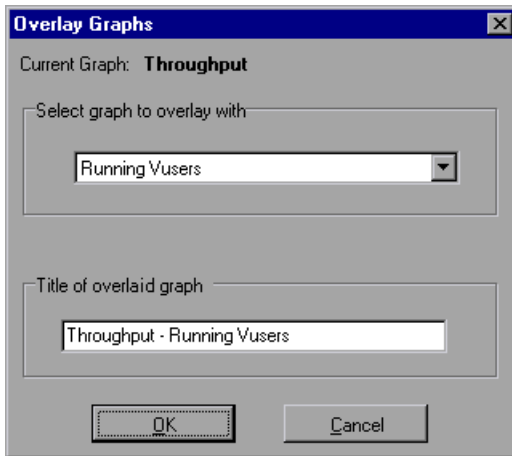
## Merging Graphs

ProTune lets you merge the results of two graphs from the same session step into a single graph. The merging allows you to compare several different measurements at once. For example, you can make a merged graph to display the Web Throughput and Hits per Second, as a function of the elapsed time. Note that in order to merge graphs, their x-axis must be the same measurement.

When you overlay the contents of two graphs that share a common x-axis, the left y-axis on the merged graph shows the current graph's values. The right y-axis shows the values of the graph that was merged.

**To overlay two graphs:**


- 1 Right-click one of the graphs you want to overlay, and select **Overlay Graphs**. The Overlay Graphs dialog box opens.



- 2 Select a graph with which you want to overlay the current graph. The drop-down list only shows the active graphs that have a common x-axis with the current graph.
- 3 Enter a title for the overlaid graph.
- 4 Click **OK**. The merged graph appears in the graph view area.

## Understanding Online Monitor Graphs

Online monitor graphs display information about the measurements listed below the graph. Each value is represented by a colored line. A legend beneath the graph indicates the color and measurement.



By default, the online monitor displays each measurement in a session step in the legend below the graphs. The legend displays the measurements for the selected graph.

---

**Note:** In a goal-oriented session step, the goal you defined is also displayed in the appropriate graph.

---

To get additional information about a measurement, right-click the measurement and choose **Description**.

To focus on a particular line, you can:

- **Highlight a measurement:** To highlight a specific measurement, select it in the legend. The corresponding line in the graph is displayed in blue.
- **Hide a measurement:** To hide a measurement, right-click the measurement and choose **Hide**.

To show a hidden measurement, right-click the measurement and choose **Show**.

- **Pause the monitor:** To pause a specific graph during session step execution, select the graph and choose **Monitors > Online Graph > Freeze**, or right-click the graph and select **Freeze**. To resume, repeat one of the above actions. When you resume, the graph displays the data for the paused period.

## Configuring Online Measurements

You can configure the following online measurement settings:

- Changing Line Colors
- Setting the Scale of the Measurement
- Showing and Hiding Transactions

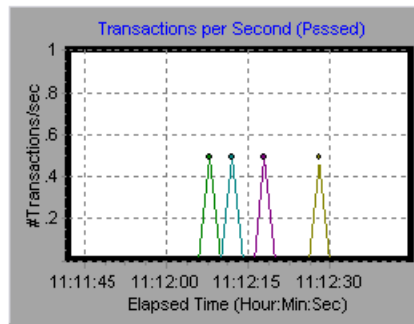
### Changing Line Colors

ProTune assigns a unique color to each measurement. You can modify the color using the configuration interface.

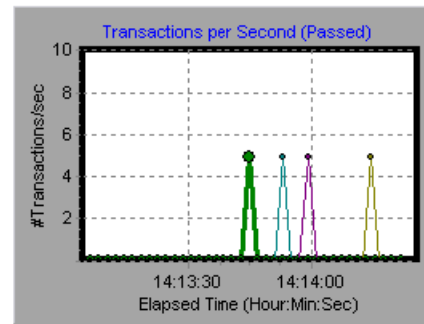
### Setting the Scale of the Measurement

You can modify the scale of a measurement—the relationship between the y-axis and the graph's actual value. For example, a scale set at 1 indicates that the measurement's value is the value of the y-axis. If you choose a scale of 10, you must divide the y-axis value by 10 to obtain the true value of the measurement.

In the following example, the same graph is displayed with a scale of 1 and 10.



*scale = 1*



*scale = 10*

The actual graph values range from 0-1, as shown in the left graph. You can view the information more accurately using a larger scale for the display, as shown in the right graph. However, to obtain the actual values, you need to divide the displayed value by the scale. In the example above, the highest value shown in the graph is 5. Since the scale is 10, the actual value is 0.5.

The legend below the graph indicates the scale factor.

| Cobr | Scale | Measurement                       | Machine | Max         | Min        | Avg        | Std        | Last      |
|------|-------|-----------------------------------|---------|-------------|------------|------------|------------|-----------|
|      | 10    | Processor Queue Length (System)   | zeus    | 3           |            | 1.823529.. | 0.705882.. | 1         |
|      | 1     | File Data Operations/sec (System) | zeus    | 127.1463... | 6.64241... | 43.36583.. | 24.31799.. | 49.928041 |

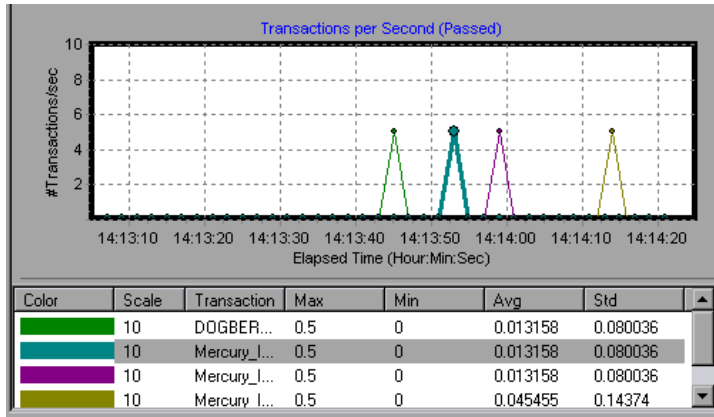
*scale factor*

By default, ProTune uses the *autoscale* option, which automatically scales the measurements by calculating the best ratio for displaying the graph.

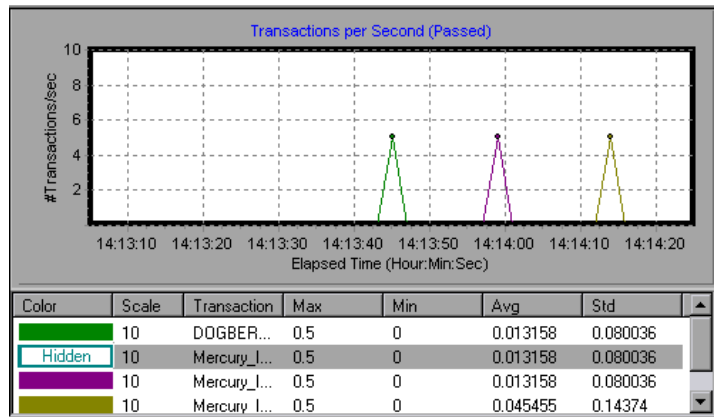
## Showing and Hiding Transactions

By default, the Transaction Monitor displays a line for each item in the transaction list. You can hide the line for any of the monitored transactions in order to focus on a specific measurement.

In the following example, a line is shown for each measurement

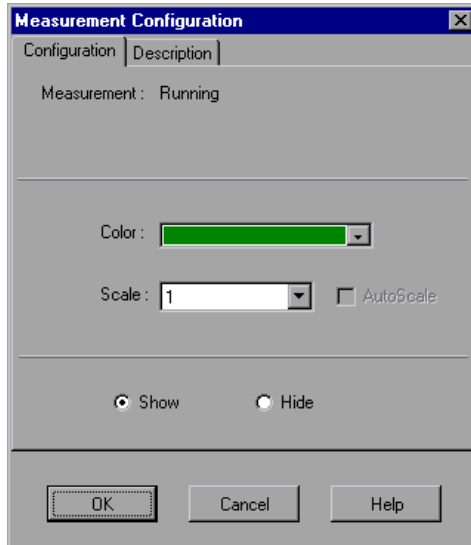


In this example, the second item in the legend is hidden.






**To configure a measurement:**

- 1 In the legend below the graphs, select the measurement you want to configure. Right-click and choose **Configure**. The Measurement Configuration dialog box opens.



- 2 To change the color of the line, select a color from the Color list.
- 3 To change the scale, clear the **Autoscale** check box and select the desired ratio from the Scale list.
- 4 To hide a measurement, click **Hide**. To show a hidden resource, click **Show**.  
Note that you can also show and hide measurements without opening the Measurement Configuration dialog box, by right-clicking a measurement in the legend and selecting **Show/Hide**.
- 5 Click **OK** to accept the settings and close the dialog box.

The specified changes are reflected in the graph and in the legend beneath the graph. The color is displayed in the first column of the legend. Hidden transactions are displayed as unfilled boxes. The scale is displayed in the legend's second column.

| Color   | Scale | Transaction  | Max | Min | Avg      | Std      |
|---|-------|--------------|-----|-----|----------|----------|
|  | 10    | DOGBER...    | 0.5 | 0   | 0.013158 | 0.080036 |
| Hidden  | 10    | Mercury_I... | 0.5 | 0   | 0.013158 | 0.080036 |
|  | 10    | Mercury_I... | 0.5 | 0   | 0.013158 | 0.080036 |
|  | 10    | Mercury_I... | 0.5 | 0   | 0.045455 | 0.14374  |

## Exporting Online Monitor Graphs

ProTune allows you to export the online graph to HTML for viewing at a later stage. When you export to HTML, the legend is also displayed with the graph. You can export all graphs or only the selected one.

**To export online graphs to HTML:**

- 1 To export a specific graph, select the graph you want to export and choose **Monitors > Online Graphs > Export to HTML**. The Select Filename and Path dialog box opens.
- 2 To export all graphs in the Online Monitor view, choose **Monitors > Export Online Graphs to HTML**. The Select Filename and Path dialog box opens.
- 3 Specify a filename and path and click **Save**.

## Viewing Data Offline

After monitoring resources during a session step run, you can view a graph of the data that was gathered using the ProTune Analysis. When you run the Analysis utility, it processes the data and generates a graph for each measurement that was monitored.

To view a graph, choose **Graph > Add Graph** in the Analysis window. For more information about working with the ProTune Analysis at the conclusion of the session step, refer to the *ProTune Analysis User's Guide*.



# 15

---

## Monitoring over a Firewall

To enable monitoring of your servers from outside the firewall, *Monitors over Firewall* is installed on designated machines inside the firewall. The installation sets up the Server Monitor mediator (referred to as the “mediator” in this chapter) as well as the Server Monitor configuration tool. You then configure the servers to monitor, and define the specific measurements that ProTune collects for each monitored server.

This chapter describes:

- Installing Monitors over Firewall
- Installing MI Listener
- Preparing for Data Collection
- Configuring Server Monitor Properties
- Adding and Removing Measurements
- Configuring Measurement Frequency
- Configuring the Network Delay Monitor over a Firewall

## About Monitoring over the Firewall

Once you set up your environment, as described in Chapter 13, “Working with Firewalls,” and install the Monitoring over Firewall component, as described in “Installing Monitors over Firewall,” on page 200, continue with the steps below:

### 1 Prepare for data collection.

Check that you can obtain information for the monitors configured inside the firewall. Refer to “Preparing for Data Collection,” on page 205.

### 2 Configure server monitor properties.

Refer to “Configuring Server Monitor Properties,” on page 205.

### 3 Add and remove measurements.

Add measurements to monitor for each server. If ProTune added default measurements, you can edit them as required. Refer to “Adding and Removing Measurements,” on page 208.

### 4 Configure measurement frequencies.

Set a measurement schedule for each measurement to be reported. Refer to “Configuring Measurement Frequency,” on page 209.

## Installing Monitors over Firewall

*Monitors over Firewall* may have been installed during ProTune installation. To check whether it was installed, click **Start > Programs > ProTune > Advanced Settings**. If the **Monitor Configuration** option appears on the list of ProTune options, then *Monitors over Firewall* was already installed, and you can proceed to “Installing MI Listener” on page 205.

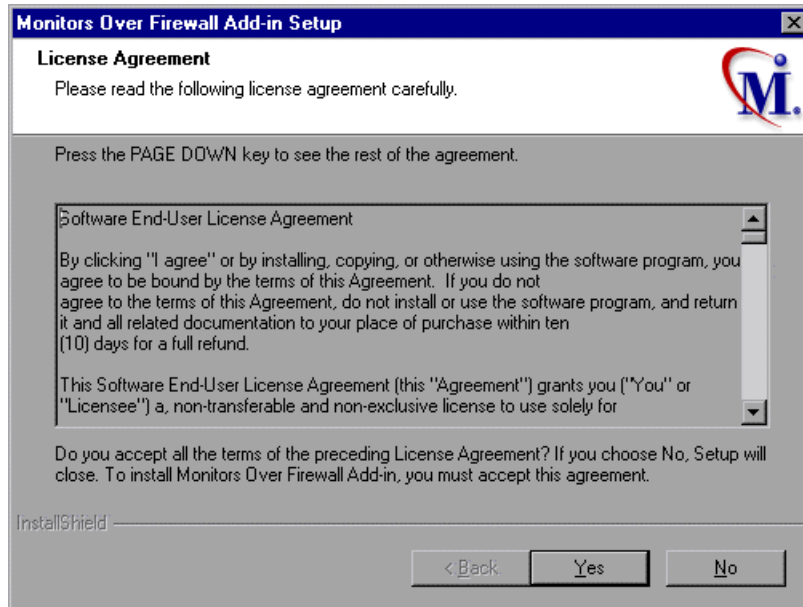
If Monitors over Firewall was not yet installed, you need to install it using one of the following:

- ▶ Perform a custom installation of ProTune from the ProTune CD, choosing only the Monitors over Firewall option. For instructions on performing a custom installation of ProTune, refer to the *ProTune Installation Guide*.

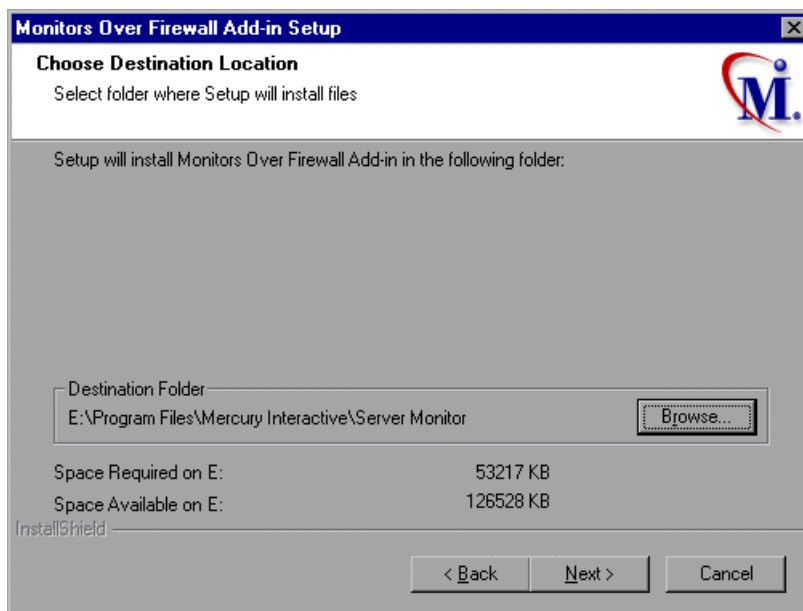
- Obtain the *Monitors over Firewall* file from the Mercury Interactive Customer Support Web site (<http://support.mercuryinteractive.com>). *Monitors over Firewall* is a stand-alone downloadable installation. It comes as a self-extracting installer file.

**To install Monitors over Firewall from the Mercury Interactive Customer Support Web site:**

- 1 Copy the self-extracting installer file to the mediator machine.
- 2 Double-click the installer file to begin installation. The software license agreement appears. Read the agreement, and click **Yes** to accept it. If you click **No**, Setup closes.

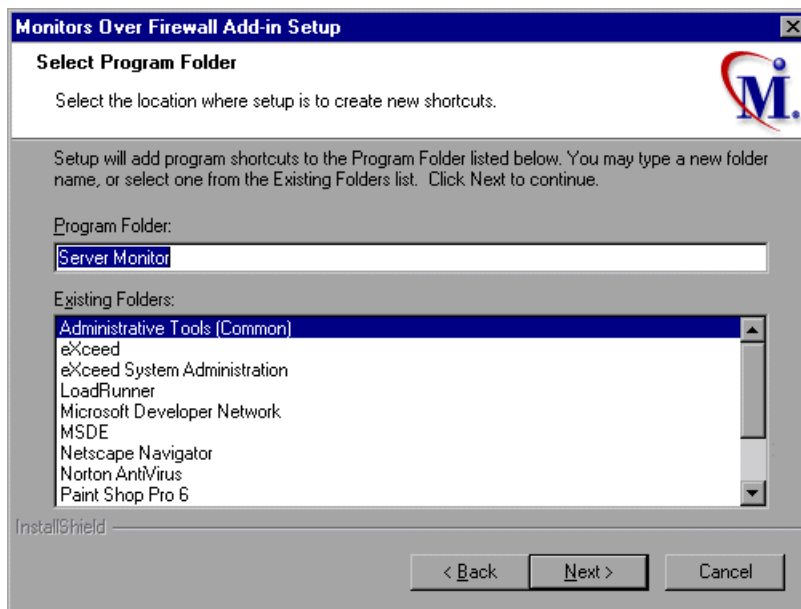


- 3 In the Choose Destination Location screen, specify the folder in which to install the add-in. To select a different location, click **Browse**, choose a folder, and click **OK**.




Click **Next**.

- 4 In the Select Program Folder screen, specify a program folder, or accept the default folder, *Server Monitor*.



Click **Next**.

- 5 In the Start Copying Files screen, review your settings. To make changes, click **Back**.



Click **Next**.

- 6 The installation process begins. To quit the installation, click **Cancel**.
- 7 Setup completes the installation process. The Setup Complete screen prompts you to restart your computer. You can delay restarting your computer until a later point, however, you must restart your computer before you use ProTune Server Monitors.

Click **Finish** to complete the setup process.

## Installing MI Listener

To enable monitoring over a firewall, you need to install MI Listener on one or more machines in the same LAN as the Console machine. Note that the Console installation automatically includes the MI Listener, so you can designate the Console as the MI Listener machine. For instructions, refer to the *ProTune Installation Guide*.

## Preparing for Data Collection

In order to obtain information for the monitors configured inside the firewall, you must create a unique connection between the Console and the mediator machine, via the Mercury Interactive listener machine, MI Listener. You establish this connection by defining the mediator machine as a load generator.

### To configure the Console for data collection:

- 1** You should already have configured the ProTune agent and the Console to operate over the firewall, as described in Chapter 13, “Working with Firewalls.”
- 2** Remember that in the Firewall tab of the Load Generator Information dialog box, you should enter the IP address of the MI Listener machine, and check **Enable Monitoring over Firewall**.
- 3** Connect to the load generator. Make sure that you obtain information for the monitors configured inside the firewall.

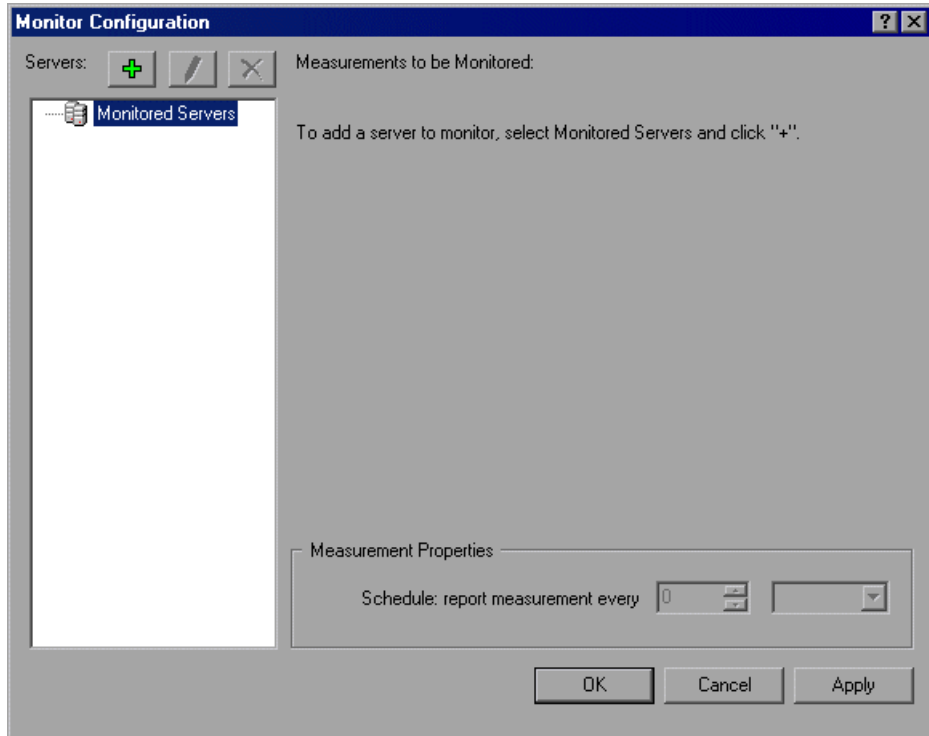
## Configuring Server Monitor Properties

The next step is to add the server monitors. You configure server monitor properties (select the server whose resources you want to monitor, and the type of monitors to run), add the measurements to monitor for each server, and specify the frequency with which you want the monitored measurements to be reported.

To enable monitoring over the firewall, you need to configure server monitor properties.

**To configure server monitor properties:**


- 1** Select **Start > Programs > ProTune > Advanced Settings > Monitor Configuration**. For machines without the complete ProTune installation, select **Start > Programs > Server Monitor > Monitor Configuration**. The Monitor Configuration dialog box opens.







- 2 Click the **Add Server** button. The New Monitored Server Properties dialog box opens.



- 3 In the Monitored Server box, type the name or IP address of the server whose resources you want to monitor.


---

**Note:** To add several servers simultaneously, separate the server names or IP ranges with commas. For example: 255.255.255.0-255.255.255.5, server1, server2.

---

- 4 From the Available Monitors list, select the monitors appropriate for the server being monitored.

- 5 Click **OK** to close the New Monitored Server Properties dialog box to display the Monitored Servers list.




Note that, for certain monitors, ProTune displays default measurements in the right pane. For details on selecting measurements, see “Adding and Removing Measurements” on page 208.

- 6 To add additional monitored servers to the list, repeat steps 1-5.
- 7 Click **Apply** to save your settings.

## Adding and Removing Measurements


After you configure one or more server machines to monitor, you add measurements to monitor for each server. If ProTune added default measurements, you can edit them as required.

### To add a measurement to monitor:

- 1 Select a server from the Monitored Servers list.
- 2  Click the **Add Measurement** button. Select the appropriate monitor. A dialog box opens, enabling you to choose measurements for the monitor you selected.
- 3 Select the measurements that you want to monitor, and click **OK**.
- 4 Click **Apply** to save your settings.

For information on configuring measurements for each server monitor, see the relevant chapter.


### To remove a measurement from the measurements list:

- 1  Select the measurement, and click the **Delete** button.
- 2 Click **Apply** to save your settings.

## Configuring Measurement Frequency

Once you have configured monitor measurements, you configure measurement frequency.

In the Measurement Properties section, you set a measurement schedule for each measurement to be reported.



**To set a measurement schedule for a measurement:**

- 1** Select the configured server measurement you want to schedule.
- 2** Specify the frequency at which you want ProTune to report the measurement.
- 3** Click **Apply** to save your settings.

## Configuring the Network Delay Monitor over a Firewall

To run the Network Delay Monitor when there are firewalls between the Console machine and the source machine, you must configure the Network Delay Monitor (see “Configuring the Network Monitor,” on page 268), and add the following to step 3 (on page 269):

In the **Monitor the network delay from machine** section, enter the server name or IP address of the source machine according to the following format:  
*<MI Listener machine>:<source machine local key>*.

where source machine local key is the unique key that you chose when configuring the ProTune Agent on the source machine.

For example: 12.12.12.3:vds



# 16

---

## Run-Time and Transaction Monitoring

While running a session step, you can use ProTune's Run-Time and Transaction monitors to view graphs of run-time status and transaction performance.

This chapter describes:

- ▶ Run-Time Graphs
- ▶ User-Defined Data Points Graph
- ▶ Transaction Monitor Graphs
- ▶ Enabling the Transaction Monitor
- ▶ Adding Transactions to a Script
- ▶ Enabling Web Page Breakdown




### About Run-Time and Transaction Graphs

The *Run-Time* monitor provides information about the status of the Vusers participating in the session step, and the number and types of errors that the Vusers generate. In addition, the Run-Time monitor provides the User-Defined Data Points graph, which displays the real time values for user-defined points in a Vuser script.

The *Transaction* monitor displays the transaction rate and response time during session step execution. For more information about transactions, see "Adding Transactions to a Script" on page 216.

## Run-Time Graphs

The monitor's **Running Vusers** graph provides information about the status of the Vusers running in the current session step on all load generator machines. The graph shows the number of running Vusers, while the information in the legend indicates the number of Vusers in each state.

| Color   | Scale | Status   | Max | Min | Avg         | Std         | Last |
|---|-------|----------|-----|-----|-------------|-------------|------|
|  | 1     | Running  | 14  | 2   | 7.632653... | 3.783389... | 14   |
|  | 1     | Error    | 0   | 0   | 0           | 0           | 0    |
|  | 1     | Finished | 0   | 0   | 0           | 0           | 0    |

The Status field of each Vuser displays the current status of the Vuser. The following table describes each Vuser status.

| Status   | Description   |
|----------|---|
| RUNNING  | The total number of Vusers currently running on all load generators.  |
| READY    | The number of Vusers that completed the initialization section of the script and are ready to run.  |
| FINISHED | The number of Vusers that have finished running. This includes both Vusers that passed and failed.  |
| ERROR    | The number of Vusers whose execution generated an error. Check the Status field in the Vuser view or the Output window for a complete explanation of the error. |

The monitor's **Error Statistics** graph provides details about the number of errors that accrue during each second of the session step run. The errors are grouped by error source—for example, the location in the script or the load generator name.

The **Vusers with Error Statistics** graph provides details about the number of Vusers that generate errors during session step execution. The errors are grouped by error source.

## User-Defined Data Points Graph

The **User-Defined Data Points** graph displays the real-time values of user-defined data points. You define a data point in your Vuser script by inserting an `lr_user_data_point` function at the appropriate place (`user_data_point` for GUI Vusers and `lr.user_data_point` for Java Vusers).

```

Action1()
{
    lr_think_time(1);
    lr_user_data_point ("data_point_1",1);
    lr_user_data_point ("data_point_2",2);
    return 0;
}

```

For Vuser protocols that support the graphical script representations such as Web and Oracle NCA, you insert a data point as a User Defined step. Data point information is gathered each time the script executes the function or step. For more information about data points, refer to the *Online Function Reference*.

By default, ProTune displays all of the data points in a single graph. The legend provides information about each data point. If desired, you can hide specific data points using the legend below the graphs.

You can also view data points offline, after the completion of the session step. For more information, refer to the *ProTune Analysis User's Guide*.

## Transaction Monitor Graphs

The *Transaction* monitor provides the following graphs:

- ▶ Transaction Response Time
- ▶ Transactions per Second (Passed)
- ▶ Transactions per Second (Failed, Stopped)
- ▶ Total Transactions per Second (Passed)

The **Transaction Response Time** graph shows the average response time of transactions in seconds (y-axis) as a function of the elapsed time in the session step (x-axis).

The **Transactions per Second (Passed)** graph shows the number of successful transactions performed per second (y-axis) as a function of the elapsed time in the session step (x-axis).

The **Transactions per Second (Failed, Stopped)** graph shows the number of failed and stopped transactions per second (y-axis) as a function of the elapsed time in the session step (x-axis).

The **Total Transactions per Second (Passed)** graph shows the total number of completed, successful transactions per second (y-axis) as a function of the elapsed time in the session step (x-axis).




## Enabling the Transaction Monitor

The Transaction monitor is enabled by default—it automatically begins monitoring Vuser transactions at the start of a session step. You can disable the Transaction monitor in order to conserve resources.

To enable the Transaction monitor:

- 1 Choose **Tools > Options** and select the **Monitors** tab.



- 2 Enable transaction monitoring by selecting the **Enable Transaction Monitor** check box. To disable transaction monitoring, clear the **Enable Transaction Monitor** check box.

## Adding Transactions to a Script

If there are no transactions defined in your Vuser script, no data will be displayed in the online graphs. To add transactions to an existing script, edit it using the appropriate tool. The following table shows the script generation tools for each script type:

| Script type     | Editing tool            |
|-----------------|-------------------------|
| GUI Windows     | WinRunner               |
| non-GUI Windows | VuGen (Vuser Generator) |
| SAP             | QuickTest for SAP       |

### To add a transaction to a script:

- 1 Click the **Design** tab to view the list of Vuser groups and scripts.
- 2 To edit a script for a Vuser group, select the group and click the **View Script** button to the right of the Session Groups window. The script generation tool opens.

To edit a script for an individual Vuser, click **Vusers**. Right-click the Vuser whose script you want to edit, and select **View Script** to open the script generation tool.

- 3 Insert Start and End Transaction functions or markers throughout your script.

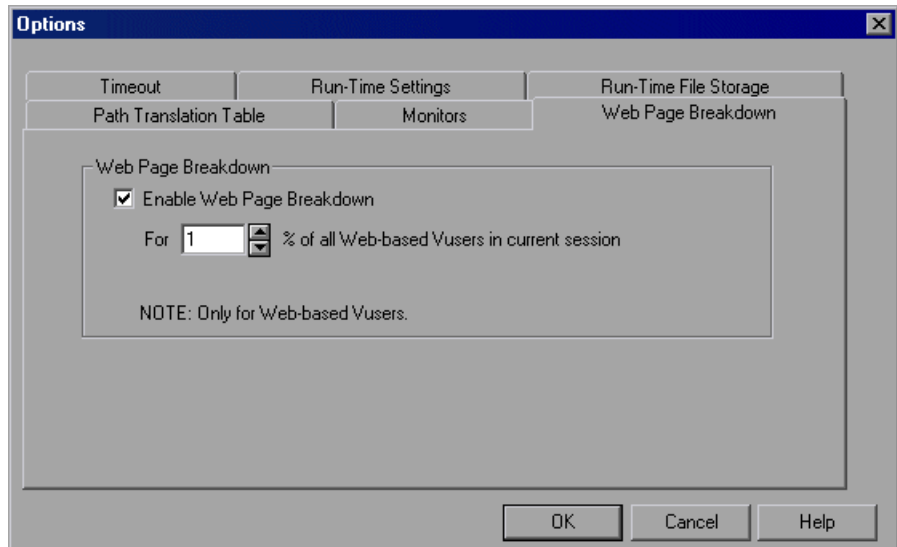
For more information, refer to the appropriate user's guide as described in the *Welcome* chapter.

## Enabling Web Page Breakdown

In order for the Analysis to generate Web Page Breakdown graphs, which provide you with performance information for each transaction and sub-transaction defined in your script, you must enable the Web page breakdown feature in the Console before running your session step.

To enable Web page breakdown:

- 1 Choose **Tools > Options** and select the **Web Page Breakdown** tab.



- 2 Select **Enable Web Page Breakdown**, and specify the percentage of Web Vusers for which you want Web page breakdown to be performed.

For more information about Web Page Breakdown graphs, refer to the *ProTune Analysis User's Guide*.



# 17

---

## Web Resource Monitoring

You can obtain information about the performance of your Web server using ProTune's Web Resource monitor.

This chapter describes:

- ▶ Hits per Second Graph
- ▶ Throughput Graph
- ▶ HTTP Responses per Second Graph
- ▶ Pages Downloaded per Second Graph
- ▶ Retries per Second Graph
- ▶ Connections Graph
- ▶ Connections per Second Graph
- ▶ SSL Connections per Second Graph

### About Web Resource Monitoring

The Web Resource monitor enables you to analyze the throughput on the Web server, the number of hits per second that occurred during the session, the number of HTTP responses per second, the HTTP status codes (which indicate the status of HTTP requests, for example, "the request was successful," "the page was not found") returned from the Web server, the number of downloaded pages per second, and the number of server retries per second.


## Hits per Second Graph

The **Hits Per Second** graph shows the number of hits (HTTP requests) to the Web server (y-axis) as a function of the elapsed time in the session step (x-axis). This graph can display the whole step, or the last 60, 180, 600, or 3600 seconds. You can compare this graph to the Transaction Response Time graph to see how the number of hits affects transaction performance.

## Throughput Graph

The **Throughput** graph shows the amount of throughput on the Web server (y-axis) during each second of the session step run (x-axis). Throughput is measured in bytes and represents the amount of data that the Vusers received from the server at any given second. You can compare this graph to the Transaction Response Time graph to see how the throughput affects transaction performance.

In the following example, the Transaction Response time graph is compared with the Throughput graph. It is apparent from the graph that as the throughput decreases, the transaction response time also decreases. The peak throughput occurred at approximately 1 minute into the step. The highest response time also occurred at this time.



## HTTP Responses per Second Graph

The **HTTP Responses per Second** graph shows the number of HTTP status codes—which indicate the status of HTTP requests, for example, “the request was successful,” “the page was not found”—(y-axis) returned from the Web server during each second of the session step run (x-axis), grouped by status code. You can group the results shown in this graph by script (using the "Group By" function) to locate scripts which generated error codes.

The following table displays a list of HTTP status codes:

| Code | Description                   |
|------|-------------------------------|
| 200  | OK                            |
| 201  | Created                       |
| 202  | Accepted                      |
| 203  | Non-Authoritative Information |
| 204  | No Content                    |
| 205  | Reset Content                 |
| 206  | Partial Content               |
| 300  | Multiple Choices              |
| 301  | Moved Permanently             |
| 302  | Found                         |
| 303  | See Other                     |
| 304  | Not Modified                  |
| 305  | Use Proxy                     |
| 307  | Temporary Redirect            |
| 400  | Bad Request                   |
| 401  | Unauthorized                  |

| Code | Description                     |
|------|---------------------------------|
| 402  | Payment Required                |
| 403  | Forbidden                       |
| 404  | Not Found                       |
| 405  | Method Not Allowed              |
| 406  | Not Acceptable                  |
| 407  | Proxy Authentication Required   |
| 408  | Request Timeout                 |
| 409  | Conflict                        |
| 410  | Gone                            |
| 411  | Length Required                 |
| 412  | Precondition Failed             |
| 413  | Request Entity Too Large        |
| 414  | Request - URI Too Large         |
| 415  | Unsupported Media Type          |
| 416  | Requested range not satisfiable |
| 417  | Expectation Failed              |
| 500  | Internal Server Error           |
| 501  | Not Implemented                 |
| 502  | Bad Gateway                     |
| 503  | Service Unavailable             |
| 504  | Gateway Timeout                 |
| 505  | HTTP Version not supported      |

For more information on the above status codes and their descriptions, see <http://www.w3.org/Protocols/rfc2616/rfc2616-sec10.html#sec10>.



## Pages Downloaded per Second Graph

The **Pages Downloaded per Second** graph shows the number of Web pages (y-axis) downloaded from the server during each second of the session step run (x-axis). This graph helps you evaluate the amount of load Vusers generate, in terms of the number of pages downloaded.

---

**Note:** In order to view the Pages Downloaded per Second graph, you must select **Pages per second (HTML Mode only)** from the script's run-time settings Preferences tab before running your session.


---

Like throughput, downloaded pages per second is a representation of the amount of data that the Vusers received from the server at any given second.

- ▶ The Throughput graph takes into account each resource and its size (for example, the size of each *.gif* file, the size of each Web page).
- ▶ The Pages Downloaded per Second graph takes into account simply the number of pages.

In the following example, the Throughput graph is compared with the Pages Downloaded per Second graph. It is apparent from the graph that throughput is not proportional to the number of pages downloaded per second.

For example, between 15 and 16 seconds into the session run, the throughput decreased while the number of pages downloaded per second increased.



## Retries per Second Graph

The **Retries Per Second** graph shows the number of attempted Web server connections (y-axis) as a function of the elapsed time in the session step (x-axis). A server connection is retried when the initial connection was unauthorized, when proxy authentication is required, when the initial connection was closed by the server, when the initial connection to the server could not be made, or when the server was initially unable to resolve the load generator's IP address.

## Connections Graph

The **Connections** graph shows the number of open TCP/IP connections (y-axis) at each point in time of the session step (x-axis). Note that one HTML page may cause the browser to open several connections, when links on the page go to different Web addresses. Two connections are opened for each Web server.

This graph is useful in indicating when additional connections are needed. For example, if the number of connections reaches a plateau, and the transaction response time increases sharply, adding connections would probably cause a dramatic improvement in performance (reduction in the transaction response time).

## Connections per Second Graph

The **Connections Per Second** graph shows the number of new TCP/IP connections (y-axis) opened each second of the session step (x-axis). This number should be a small fraction of the number of hits per second, because new TCP/IP connections are very expensive in terms of server, router and network resource consumption. Ideally, many HTTP requests should use the same connection, instead of opening a new connection for each request.

## SSL Connections per Second Graph

The **SSL Connections per Second** graph shows the number of new and reused SSL Connections (y-axis) opened in each second of the session step (x-axis). An SSL connection is opened by the browser after a TCP/IP connection has been opened to a secure server.

Because creating a new SSL connection entails heavy resource consumption, you should try to open as few new SSL connections as possible; once you've established an SSL connection, you should reuse it. There should be no more than one new SSL connection per Vuser. If you've configured ProTune to simulate a new Vuser at each iteration (via the Browser Emulation tab in the Run-Time Settings menu), you should have no more than one new SSL connection per Vuser per iteration. Ideally, you should have very few new TCP/IP and SSL connections each second.

# 18

---

## System Resource Monitoring

You can monitor a machine's system resource usage during a session step run using ProTune's System Resource monitors.

This chapter describes:

- ▶ Configuring the Windows Resources Monitor
- ▶ Configuring the UNIX Resources Monitor
- ▶ Configuring an rstatd Daemon on UNIX
- ▶ Configuring the SNMP Resources Monitor
- ▶ Configuring the TUXEDO Monitor
- ▶ Configuring the Antara FlameThrower Monitor
- ▶ Configuring the SiteScope Monitor

### About System Resource Monitoring

A primary factor in a transaction's response time is its system resource usage. Using the ProTune resource monitors, you can monitor the Windows, UNIX, SNMP, Antara FlameThrower, and SiteScope resources on a machine during a session step run, and determine why a bottleneck occurred on a particular machine.

The Windows measurements correspond to the built-in counters available from the Windows Performance Monitor.

The UNIX measurements include those available by the *rstatd* daemon: average load, collision rate, context switch rate, CPU utilization, incoming packets error rate, incoming packets rate, interrupt rate, outgoing packets error rate, outgoing packets rate, page-in rate, page-out rate, paging rate, swap-in rate, swap-out rate, system mode CPU utilization, and user mode CPU utilization.

---

**Note:** You must configure an *rstatd* daemon on all UNIX machines being monitored. For information on how to configure an *rstatd* daemon, refer to the UNIX *man* pages, or see "Configuring an rstatd Daemon on UNIX," on page 236.

---

The TUXEDO monitor can monitor the server, load generator machine, workstation handler, and queue in a TUXEDO system. Note that in order to run the TUXEDO monitor, you must install the TUXEDO client libraries on the machine you want to monitor. For information on configuring the TUXEDO monitor, see "Configuring the TUXEDO Monitor," on page 241.

The SNMP monitor is available for monitoring machines using the Simple Network Management Protocol (SNMP). SNMP monitoring is platform independent.

The Antara FlameThrower monitor can measure the following performance counters: Layer, TCP, HTTP, SSL/HTTPS, Sticky SLB, FTP, SMPT, POP3, DNS, and Attacks.

The SiteScope monitor can measure server, network, and processor performance counters. For detailed information on the performance counters that SiteScope can monitor, refer to the relevant SiteScope documentation.

The resource monitors are automatically enabled when you execute a session step. However, you must specify the machine you want to monitor and which resources to monitor for each machine. You can also add or remove machines and resources during the session step run.

## Configuring the Windows Resources Monitor

Windows NT and Windows 2000 measurements correspond to the built-in counters available from the Windows Performance Monitor.

---

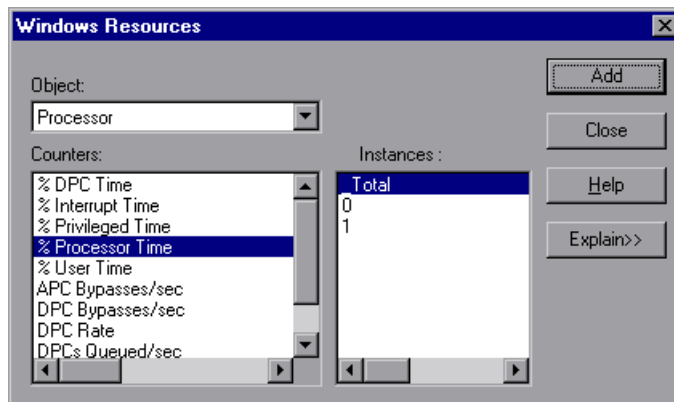
**Note:** To monitor a Windows NT or 2000 machine through a firewall, use TCP, port 139.

---

### To configure the Windows Resources monitor:



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select Windows Resources and then click **Add**. The Windows Resources dialog box is displayed.



- 5 Select an object, a counter, and an instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.

The following default measurements are available for Windows machines:

| Object    | Measurement                            | Description   |
|-----------|--|---|
| System    | <b>% Total Processor Time</b>          | The average percentage of time that all the processors on the system are busy executing non-idle threads. On a multi-processor system, if all processors are always busy, this is 100%, if all processors are 50% busy this is 50% and if 1/4th of the processors are 100% busy this is 25%. It can be viewed as the fraction of the time spent doing useful work. Each processor is assigned an Idle thread in the Idle process which consumes those unproductive processor cycles not used by any other threads.  |
| System    | <b>File Data Operations/sec</b>        | The rate at which the computer issues read and write operations to file system devices. This does not include File Control Operations.  |
| Processor | <b>% Processor Time (Windows 2000)</b> | The percentage of time that the processor is executing a non-idle thread. This counter was designed as a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the idle process in each sample interval, and subtracting that value from 100%. (Each processor has an idle thread which consumes cycles when no other threads are ready to run). It can be viewed as the percentage of the sample interval spent doing useful work. This counter displays the average percentage of busy time observed during the sample interval. It is calculated by monitoring the time the service was inactive, and then subtracting that value from 100%. |



| Object       | Measurement                   | Description  |
|--------------|-------------------------------|--|
| System       | <b>Processor Queue Length</b> | The instantaneous length of the processor queue in units of threads. This counter is always 0 unless you are also monitoring a thread counter. All processors use a single queue in which threads wait for processor cycles. This length does not include the threads that are currently executing. A sustained processor queue length greater than two generally indicates processor congestion. This is an instantaneous count, not an average over the time interval. |
| Memory       | <b>Page Faults/sec</b>        | This is a count of the page faults in the processor. A page fault occurs when a process refers to a virtual memory page that is not in its Working Set in the main memory. A page fault will not cause the page to be fetched from disk if that page is on the standby list (and hence already in main memory), or if it is in use by another process with which the page is shared.   |
| PhysicalDisk | <b>% Disk Time</b>            | The percentage of elapsed time that the selected disk drive is busy servicing read or write requests.  |
| Memory       | <b>Pool Nonpaged Bytes</b>    | The number of bytes in the nonpaged pool, a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. Nonpaged pool pages cannot be paged out to the paging file. They remain in main memory as long as they are allocated.  |

| Object  | Measurement          | Description  |
|---------|----------------------|--|
| Memory  | Pages/sec            | The number of pages read from the disk or written to the disk to resolve memory references to pages that were not in memory at the time of the reference. This is the sum of Pages Input/sec and Pages Output/sec. This counter includes paging traffic on behalf of the system cache to access file data for applications. This value also includes the pages to/from non-cached mapped memory files. This is the primary counter to observe if you are concerned about excessive memory pressure (that is, thrashing), and the excessive paging that may result. |
| System  | Total Interrupts/sec | The rate at which the computer is receiving and servicing hardware interrupts. The devices that can generate interrupts are the system timer, the mouse, data communication lines, network interface cards, and other peripheral devices. This counter provides an indication of how busy these devices are on a computer-wide basis. See also Processor:Interrupts/sec.   |
| Objects | Threads              | The number of threads in the computer at the time of data collection. Notice that this is an instantaneous count, not an average over the time interval. A thread is the basic executable entity that can execute instructions in a processor.   |
| Process | Private Bytes        | The current number of bytes that the process has allocated that cannot be shared with other processes.   |

**Note:** To change the default counters for the Windows machine monitor, see “Changing a Monitor’s Default Counters,” on page 661.

If you are monitoring a Win2000 machine, some of the NT machine default counters may not be available (such as % Total CPU usage and Interrupts/sec).

---

- 6** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.
  - 7** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.
- 

**Note:** If you want to monitor a remote Windows machine that does not use Windows domain security, you must authenticate the Console machine on the remote Windows machine. To authenticate the Console machine, create an account, or change the password of the account used to log on to the Console so that it matches the password and user name used to log on to the remote monitored Windows machine. When the remote Windows machine requests another machine’s resources, it sends the logged-in user name and password of the machine requesting the resources.

---

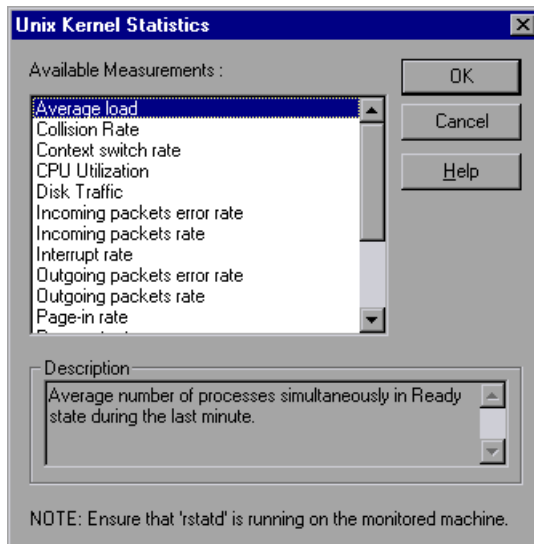
## Configuring the UNIX Resources Monitor

The UNIX kernel statistics measurements include those available by the *rstatd* daemon: average load, collision rate, context switch rate, CPU utilization, incoming packets error rate, incoming packets rate, interrupt rate, outgoing packets error rate, outgoing packets rate, page-in rate, page-out rate, paging rate, swap-in rate, swap-out rate, system mode CPU utilization, and user mode CPU utilization.

**To configure the UNIX Resources monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select UNIX Resources and then click **Add**. The UNIX Kernel Statistics dialog box opens, displaying the available measurements and server properties.



The following default measurements are available for the UNIX machine:

| <b>Measurement</b>                  | <b>Description</b>  |
|-------------------------------------|---|
| <b>Average load</b>                 | Average number of processes simultaneously in READY state during the last minute    |
| <b>Collision rate</b>               | Collisions per second detected on the Ethernet                                      |
| <b>Context switches rate</b>        | Number of switches between processes or threads, per second                         |
| <b>CPU utilization</b>              | Percent of time that the CPU is utilized  |
| <b>Disk rate</b>                    | Rate of disk transfers  |
| <b>Incoming packets error rate</b>  | Errors per second while receiving Ethernet packets                                  |
| <b>Incoming packets rate</b>        | Incoming Ethernet packets per second  |
| <b>Interrupt rate</b>               | Number of device interrupts per second  |
| <b>Outgoing packets errors rate</b> | Errors per second while sending Ethernet packets                                    |
| <b>Outgoing packets rate</b>        | Outgoing Ethernet packets per second  |
| <b>Page-in rate</b>                 | Number of pages read to physical memory, per second                                 |
| <b>Page-out rate</b>                | Number of pages written to pagefile(s) and removed from physical memory, per second |
| <b>Paging rate</b>                  | Number of pages read to physical memory or written to pagefile(s), per second       |
| <b>Swap-in rate</b>                 | Number of processes being swapped   |
| <b>Swap-out rate</b>                | Number of processes being swapped   |
| <b>System mode CPU utilization</b>  | Percent of time that the CPU is utilized in system mode                             |
| <b>User mode CPU utilization</b>    | Percent of time that the CPU is utilized in user mode                               |

---

**Note:** To change the default counters for the UNIX monitor, see "Changing a Monitor's Default Counters," on page 661.

---

- 5 To add UNIX measurements to the monitor list, select the desired measurements, and click **OK**.
- 6 Click **OK** in the UNIX Kernel Statistics dialog box to activate the UNIX monitor.

---

**Note:** Ensure that the `rstatd` daemon is correctly configured and running on the monitored UNIX machine. For more information, see "Configuring an `rstatd` Daemon on UNIX," on page 236.

---

## Configuring an `rstatd` Daemon on UNIX

To monitor UNIX resources, you must configure the `rstatd` daemon. Note that the `rstatd` daemon might already be configured, because when a machine receives an `rstatd` request, the `inetd` on that machine activates the `rstatd` automatically.

### To verify whether the `rstatd` daemon is already configured:

The `rup` command reports various machine statistics, including `rstatd` configuration. Run the following command to view the machine statistics:

```
>rup host
```

You can also use `lr_host_monitor` and see if it returns any relevant statistics.

If the command returns meaningful statistics, the `rstatd` daemon is already configured and activated. If not, or if you receive an error message, the `rstatd` daemon is not configured.

**To configure the rstatd daemon:**

- 1** Run the command: *su root*
- 2** Go to */etc/inetd.conf* and look for the rstatd row (it begins with the word rstatd). If it is commented out (with a #), remove the comment directive, and save the file.
- 3** From the command line, run:

```
kill -1 inet_pid
```

where *inet\_pid* is the pid of the inetd process. This instructs the inetd to rescan the */etc/inetd.conf* file and register all daemons which are uncommented, including the rstatd daemon.

- 4** Run *rup* again.

If the command still does not indicate that the rstatd daemon is configured, contact your system administrator.

---

**Note:** To monitor a UNIX machine through a firewall, you must run a UNIX utility called *rpcinfo* and identify the rstatd's port number. By running *rpcinfo -p <hostname>*, you will receive a list of all RPC servers registered in the host's portmapper, along with the port number. This list will not change until rstatd is stopped and rerun.

Some firewalls allow you to open an RPC program number instead of a port. In such cases, open program 100001. If are prompted to include a version number, specify versions 3 and 4.

---

## Configuring the SNMP Resources Monitor

The SNMP Resources monitor is available for monitoring any machine that runs an SNMP agent, using the Simple Network Management Protocol (SNMP).

---

**Note:** You can specify a port number in the *snmp.cfg* file. If you do not specify a port, ProTune connects to default SNMP port 161. You can also specify a machine name in the following format:  
*<server name>:<port number>*

To monitor SNMP resources through a firewall, use ports 161 or 162.

---

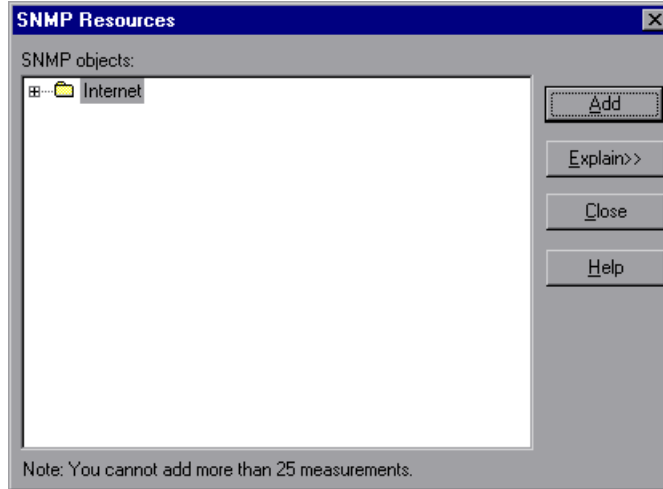
### To configure the SNMP Resources monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select SNMP (in the System Resources category) and then click **Add**.



The SNMP Resources dialog box opens.



- 5 Browse the SNMP Object tree.
- 6 To measure an object, select it, and click **Add**. For a description of each resource, click **Explain>>** to expand the dialog box. Add all the desired resources to the list, and click **Close**.

---

**Note:** The SNMP monitor can only monitor up to 25 measurements.


---

- 7 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

You can modify the list of resources that you want to monitor at any point during the session step. Note that a session step does not have to be active in order for you to monitor the resources on a remote machine.

**Note:** You can improve the level of measurement information for the SNMP monitor by enabling measurements with string values to be listed (in addition to measurements with numeric values), and by enabling the name modifier (which displays the string value as an identifying part of the measurement name).

In the following example of a measurement using the name modifier, the string value of `ProcessName` (`sched`) is displayed in addition to its instance ID (0):



To enable this feature, add the following line to the *<ProTune root folder>\dat\monitors\snmp.cfg* file:  
`SNMP_show_string_nodes=1`

**Usage Notes:** You can select more than one name modifier, but the first in the hierarchy will be used. Each time the SNMP Add Measurements dialog box opens, the information is reread from the *snmp.cfg* file. You cannot add the same measurement twice (once with a name modifier and once without it). If you do so, an error message is issued.

---

## Configuring the TUXEDO Monitor

The TUXEDO monitor allows you to measure and view your TUXEDO client's performance.

---

**Note:** If TUXEDO 7.1 or higher is installed on the Console machine, more than one TUXEDO application server can be monitored at a time. However, if TUXEDO 6.5 or below is installed on the Console machine, only one TUXEDO application server can be monitored at a time. Use a TUXEDO 6.x client if a TUXEDO 6.x server is used, and TUXEDO 7.1 or above if a TUXEDO 7.1 or above server is used.

---

### Before you set up the monitor:

- 1** Ensure that a TUXEDO workstation client (not a native client) is installed on the Console machine.
- 

**Note:** A TUXEDO workstation client communicates with the application server over the network, and is not required to run the TUXEDO application server on the same machine. A native client can only communicate with the TUXEDO application server if it is part of the relevant TUXEDO domain.

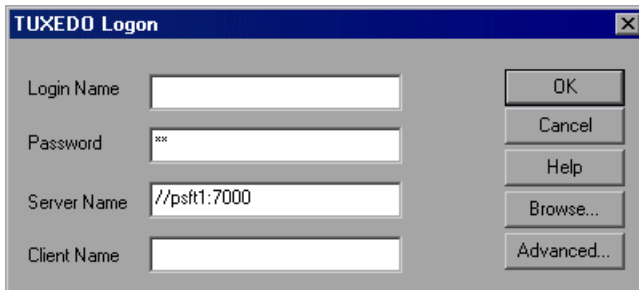
---

- 2** Define the TUXEDO environment variables on the Console machine—set the TUXDIR variable to the TUXEDO installation directory, and add the TUXEDO bin directory to the PATH variable.
- 3** Configure the TUXEDO application server so that the workstation listener (WSL) process is running. This enables the application server to accept requests from workstation clients. Note that the address and port number used to connect to the application server must match those dedicated to the WSL process.

**To configure the TUXEDO monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select TUXEDO (in the ERP-CRM Server category) and then click **Add**. The TUXEDO Logon dialog box is displayed.



- 5** Enter the TUXEDO server Login Name, Password, Server Name and Client Name.

---

**Note:** This information is located in the Logon section of the *tpinit.ini* file in the recorded script's directory. It is recommended that you use the Browse button and select the *tpinit.ini* file from a recorded script, rather than enter the values manually.

---

To obtain the correct settings for the TUXEDO monitor using the *tpinit.ini* file, click the **Browse** button and navigate to the *tpinit.ini* file of that ProTune script. You can also determine the client name from the **lrt\_tpinitialize** statement in the recorded script.

In the following example of a *tpinit.ini* file, the TUXEDO monitor was configured for a server named URANUS using port 65535, and a client named bankapp. The logon user name was Smith and the password was mypasswd.

```
[Logon]
LogonServername=//URANUS:65535
LogonUsrName=Smith
LogonCltName=bankapp
LogonGrpName=
LogonPasswd=myspasswd
LogonData=
```

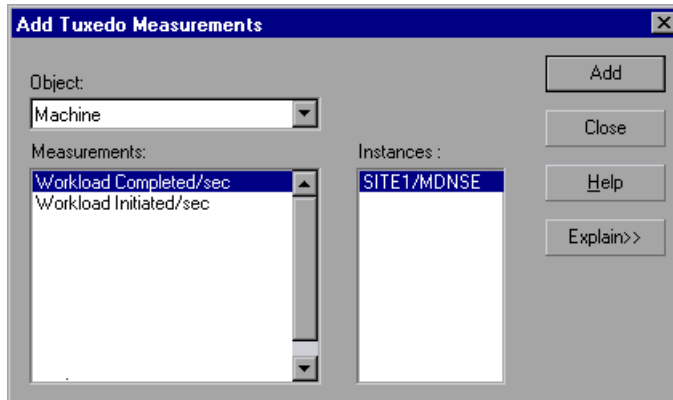
If you already know the required values, you can manually type them into the dialog box. The format of the server name is *//<machine name>:<port number>*. Alternatively, you can specify the IP address instead of the machine name. The hexadecimal format used by old versions of TUXEDO is also supported. Note that quotation marks should not be used.

---

**Note:** If you are using TUXEDO 6.5 or below, the monitor can only connect to one application server during a Console session. Once it connects to an application server, that server is the only one used by the monitor until the Console is closed. This applies even when all of the counters are deleted from the monitor.

---

- Click **OK**. The Add TUXEDO Measurements dialog box opens.



- Select a TUXEDO object from the Object list. Select the measurements and instances you want to monitor. The following table lists the available TUXEDO monitor measurements:

| Monitor | Measurements  |
|---------|---|
| Server  | <b>Requests per second</b> - How many server requests were handled per second   |
|         | <b>Workload per second</b> -The workload is a weighted measure of the server requests. Some requests could have a different weight than others. By default, the workload is always 50 times the number of requests. |

| Monitor        | Measurements   |
|----------------|--|
| <b>Machine</b> | <b>Workload completed per second</b> - The total workload on all the servers for the machine that was completed, per unit time   |
|                | <b>Workload initiated per second</b> - The total workload on all the servers for the machine that was initiated, per unit time   |
|                | <b>Current Accessers</b> - Number of clients and servers currently accessing the application either directly on this machine or through a workstation handler on this machine. |
|                | <b>Current Clients</b> - Number of clients, both native and workstation, currently logged in to this machine.  |
|                | <b>Current Transactions</b> - Number of in use transaction table entries on this machine.  |
| <b>Queue</b>   | <b>Bytes on queue</b> - The total number of bytes for all the messages waiting in the queue  |
|                | <b>Messages on queue</b> - The total number of requests that are waiting on queue. By default this is 0.   |

| Monitor                          | Measurements   |
|----------------------------------|--|
| <b>Workstation Handler (WSH)</b> | <b>Bytes received per second</b> - The total number of bytes received by the workstation handler, per unit time  |
|                                  | <b>Bytes sent per second</b> - The total number of bytes sent back to the clients by the workstation handler, per unit time  |
|                                  | <b>Messages received per second</b> - The number of messages received by the workstation handler, per unit time  |
|                                  | <b>Messages sent per second</b> - The number of messages sent back to the clients by the workstation handler, per unit time  |
|                                  | <b>Number of queue blocks per second</b> - The number of times the queue for the workstation handler blocked, per unit time. This gives an idea of how often the workstation handler was overloaded. |

- 8 Click **Add** to place the selected object on the resource list. Add all the desired objects to the list, and click **Close**.
- 9 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.



## Configuring the Antara FlameThrower Monitor

You select measurements to monitor the Antara FlameThrower server using the Antara FlameThrower Monitor Configuration dialog box.

**To configure the Antara FlameThrower monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Antara FlameThrower (in the System Resources category) and then click **Add**.

The Antara FlameThrower Monitor Configuration dialog box opens.

- 5** Browse the Measured Components tree.
- 6** Check the required performance counters in the Antara FlameThrower Monitor Configuration window's right pane.

The following tables describe the counters that can be monitored:

### Layer Performance Counters

| Measurement              | Description  |
|--------------------------|--|
| <b>TxBytes</b>           | The total number of Layer 2 data bytes transmitted.      |
| <b>TxByteRate(/sec)</b>  | The number of Layer 2 data bytes transmitted per second. |
| <b>TxFrames</b>          | The total number of packets transmitted.                 |
| <b>TxFrameRate(/sec)</b> | The number of packets transmitted per second.            |
| <b>RxBytes</b>           | The total number of Layer 2 data bytes received.         |
| <b>RxByteRate(/sec)</b>  | The number of Layer 2 data bytes received per second.    |
| <b>RxFrames</b>          | The total number of packets received.                    |
| <b>RxFrameRate(/sec)</b> | The number of packets received per second.               |

## TCP Performance Counters

| Measurement                    | Description  |
|--------------------------------|--|
| ActiveTCPConns                 | Total number of currently active TCP connections.  |
| SuccTCPConns                   | Total number of SYN ACK packets received.  |
| SuccTCPConn<br>Rate(/sec)      | Number of SYN ACK packets received per second.   |
| TCPConnLatency<br>(milisec)    | Interval between transmitting a SYN packet and receiving a SYN ACK reply packet in msec. |
| MinTCPConn<br>Latency(milisec) | Minimum TCPConnectionLatency in msec.  |
| MaxTCPConn<br>Latency(milisec) | Maximum TCPConnectionLatency in msec.  |
| TCPStdConnClose                | Total number of FIN or FIN ACK packets transmitted (Client).                             |
| TCPStdConnClose                | Total number of FIN or FIN ACK packets received (Client).                                |
| TCPStdResets                   | Total number of RST packets transmitted.   |
| TCPStdResets                   | Total number of RST packets received.  |
| SYNSent                        | Total number of SYN packets transmitted.   |
| SYNSentRate(/sec)              | Number of SYN packets transmitted per second.  |
| SYNAckSent                     | Total number of SYN ACK packets transmitted.   |
| SYNAckRate(/sec)               | Number of SYN ACK packets transmitted per second.  |

## HTTP Performance Counters

| Measurement                    | Description  |
|--------------------------------|--|
| HTTPRequests                   | Total number of HTTP Request command packets transmitted.                        |
| HTTPRequestRate (/sec)         | Number of HTTP Request packets transmitted per second.                           |
| AvgHTTPData Latency(miliseecs) | The average HTTP Data Latency over the past second in msec.                      |
| HTTPData Latency(miliseecs)    | Interval between transmitting a Request packet and receiving a response in msec. |
| DataThroughput (bytes/sec)     | The number of data bytes received from the HTTP server per second.               |
| MinHTTPData Latency(miliseecs) | Minimum HTTPDataLatency in msec.   |
| MaxHTTPData Latency(miliseecs) | Maximum HTTPDataLatency in msec.   |
| MinData Throughput (bytes/sec) | Minimum HTTPDataThroughput in seconds.   |
| MaxData Throughput (bytes/sec) | Maximum HTTPDataThroughput in seconds.   |
| SucHTTPRequests                | Total number of successful HTTP Request Replies (200 OK) received.               |
| SucHTTPRequest Rate(/sec)      | Number of successful HTTP Request Replies (200 OK) received per second.          |
| UnSucHTTP Requests             | Number of unsuccessful HTTP Requests.  |

## SSL/HTTPS Performance Counters

| Measurement                         | Description   |
|-------------------------------------|---|
| <b>SSLConnections</b>               | Number of ClientHello messages sent by the Client.  |
| <b>SSLConnection Rate(/sec)</b>     | Number of ClientHello messages sent per second.   |
| <b>SuccSSL Connections</b>          | Number of successful SSL Connections. A successful connection is one in which the Client receives the Server's finished handshake message without any errors. |
| <b>SuccSSLConnection Rate(/sec)</b> | Number of successful SSL connections established per second.  |
| <b>SSLAlertErrors</b>               | Number of SSL alert messages received by the client (e.g. bad_record_mac, decryption_failed, handshake_failure, etc.)   |
| <b>SuccSSLResumed Sessions</b>      | Number of SSL Sessions that were successfully resumed.  |
| <b>FailedSSLResumed Sessions</b>    | Number of SSL Sessions that were unable to be resumed.  |

## Sticky SLB Performance Counters

| Measurement                      | Description   |
|----------------------------------|---|
| <b>Cookie AuthenticationFail</b> | The number of Cookie's that were not authenticated by the Server.     |
| <b>SuccCookie Authentication</b> | The number of Cookie's authenticated by the server.                   |
| <b>SSLClientHellos</b>           | The number of Client Hello packets sent to the server.                |
| <b>SSLServerHellos</b>           | The number of Server Hello packets sent to back to the client.        |
| <b>SSLSessionsFailed</b>         | The number of Session ID's that were not authenticated by the server. |
| <b>SSLSessions Resumed</b>       | The number of Session ID's authenticated by the server.               |

| Measurement         | Description  |
|---------------------|--|
| succSSLClientHellos | The number of Client Hello replies received by the client or packets received by the server. |
| succSSLServerHellos | The number of Server Hello's received by the client.   |

### FTP Performance Counters

| Measurement                  | Description   |
|------------------------------|---|
| TPUsers                      | Total number of Ftp User command packets transmitted.                                     |
| FTPUserRate(/sec)            | Number of Ftp User command packets transmitted per second.                                |
| FTPUserLatency (milisecs)    | Interval between transmitting a Ftp User command packet and receiving a response in msec. |
| MinFTPUserLatency (milisecs) | Minimum FTPUsersLatency in msec.  |
| MaxFTPUserLatency (milisecs) | Maximum FTPUsersLatency in msec.  |
| SuccFTPUsers                 | Total number of successful Ftp User command replies received.                             |
| SuccFTPUserRate (/sec)       | Number of successful Ftp User command replies received per second.                        |
| FTPPasses                    | Total number of FTP PASS packets transmitted.   |
| FTPPassRate(/sec)            | Number of FTP PASS packets transmitted per second.  |
| FTPPassLatency (milisecs)    | Interval between transmitting a Ftp PASS packet and receiving a response in msec.         |
| MinFTPPassLatency (milisecs) | Minimum FTPPassLatency in msec.   |
| MaxFTPPassLatency (milisecs) | Maximum FTPPassLatency in msec.   |
| SuccFTPPasses                | Total number of successful FTP PASS replies received.                                     |

| Measurement                                 | Description  |
|---|--|
| <b>SuccFTPPassRate (/sec)</b>               | Number of successful FTP PASS replies received per second.   |
| <b>FTPControl Connections</b>               | Total number of SYN packets transmitted by the FTP client.   |
| <b>FTPControl ConnectionRate (/sec)</b>     | Number of SYN packets transmitted by the FTP client per second.  |
| <b>SuccFTPControl Connections</b>           | Total number of SYN ACK packets received by the FTP client.  |
| <b>SuccFTPControl ConnectionRate (/sec)</b> | Number of SYN ACK packets received by the FTP Client per second.   |
| <b>FTPData Connections</b>                  | Number of SYN ACK packets received by the FTP client per second.   |
| <b>FTPDataConnection Rate(/sec)</b>         | Number of SYN ACK packets transmitted by the FTP Client or received by the FTP Server per second.        |
| <b>SuccFTPData Connections</b>              | Total number of SYN ACK packets transmitted by the FTP Client or received by the FTP Server.             |
| <b>SuccFTPData ConnectionRate (/sec)</b>    | Number of SYN ACK packets received by the FTP server per second.   |
| <b>FtpAuthFailed</b>                        | Total number of error replies received by the FTP client.  |
| <b>FTPGets</b>                              | Total number of client Get requests.   |
| <b>FTPPuts</b>                              | Total number of client Put requests.   |
| <b>SuccFTPGets</b>                          | Total number of successful Get requests (data has been successfully transferred from server to client).  |
| <b>SuccFTPPuts</b>                          | Total number of successful Put requests (data has been successfully transferred from client to server) . |

## SMTP Performance Counters

| Measurement                       | Description  |
|-----------------------------------|--|
| SMTPHelos                         | Total number of HELO packets transmitted.  |
| SMTPHeloRate(/sec)                | Number of HELO packets transmitted per second.                                     |
| SMTPHeloLatency (milisecs)        | Interval between transmitting a HELO packet and receiving a response in msec.      |
| MinSMTPHelo Latency(milisecs)     | Minimum SMTPHeloLatency in msec.   |
| MaxSMTPHelo Latency(milisecs)     | Maximum SMTPHeloLatency in msec.   |
| SuccSMTPHelos                     | Total number of successful HELO replies received.                                  |
| SuccSMTPHelo Rate(/sec)           | Number of successful HELO replies received per second.                             |
| SMTPMailFroms                     | Total number of Mail From packets transmitted.                                     |
| SMTPMailFromRate (/sec)           | Number of Mail From packets transmitted per second.                                |
| SMTPMailFrom Latency(milisecs)    | Interval between transmitting a Mail From packet and receiving a response in msec. |
| MinSMTPMailFrom Latency(milisecs) | Minimum SMTPMailFromLatency in msec.   |
| MaxSMTPMailFrom Latency(milisecs) | Maximum SMTPMailFromLatency in msec.   |
| SuccSMTPMail Froms                | Total number of successful Mail From replies received.                             |
| SuccSMTPMailFrom Rate(/sec)       | Number of successful Mail From replies received per second.                        |
| SMTPRcptTos                       | Total number of RcptTo packets transmitted.  |
| SMTPRcptToRate (/sec)             | Number of RcptTo packets transmitted per second.                                   |

| Measurement                            | Description   |
|--|---|
| <b>SMTPRcptTo Latency(miliseCs)</b>    | Interval between transmitting a RcptTo packet and receiving a response in msec. |
| <b>MinSMTPRcptTo Latency(miliseCs)</b> | Minimum SMTPRcptToLatency in msec.  |
| <b>MaxSMTPRcptTo Latency(miliseCs)</b> | Maximum SMTPRcptToLatency in msec.  |
| <b>SuccSMTPRcptTos</b>                 | Total number of successful RcptTo replies received.                             |
| <b>SuccSMTPRcptTo Rate(/sec)</b>       | Number of successful RcptTo replies received per second.                        |
| <b>SMTPDatas</b>                       | Total number of Data packets transmitted.                                       |
| <b>SMTPDataRate(/sec)</b>              | Number of Data packets transmitted per second.                                  |
| <b>SMTPDataLatency (miliseCs)</b>      | Interval between transmitting a Data packet and receiving a response in msec.   |
| <b>MinSMTPData Latency(miliseCs)</b>   | Minimum SMTPDataLatency in msec.  |
| <b>MaxSMTPData Latency(miliseCs)</b>   | Maximum SMTPDataLatency in msec.  |
| <b>SuccSMTPDatas</b>                   | Total number of successful Data replies received.                               |
| <b>SuccSMTPDataRate (/sec)</b>         | Number of successful Data replies received per second.                          |

### POP3 Performance Counters

| Measurement                       | Description  |
|-----------------------------------|--|
| <b>POP3Users</b>                  | Total number of Pop3 User command packets transmitted.                                     |
| <b>POP3UserRate(/sec)</b>         | Number of Pop3 User command packets transmitted per second.                                |
| <b>POP3UserLatency (miliseCs)</b> | Interval between transmitting a Pop3 User command packet and receiving a response in msec. |



| Measurement                           | Description  |
|---------------------------------------|--|
| <b>MinPOP3User Latency(miliseccs)</b> | Minimum POP3UserLatency in msec.   |
| <b>MaxPOP3User Latency(miliseccs)</b> | Maximum POP3UserLatency in msec.   |
| <b>SuccPOP3Users</b>                  | Total number of successful Pop3 User replies received.                             |
| <b>SuccPOP3UserRate (/sec)</b>        | Number of successful Pop3 User replies received per second.                        |
| <b>POP3Passes</b>                     | Total number of Pop3 Pass command packets transmitted.                             |
| <b>POP3PassRate(/sec)</b>             | Number of Pop3 Pass command packets transmitted per second.                        |
| <b>POP3PassLatency (miliseccs)</b>    | Interval between transmitting a Pop3 Pass packet and receiving a response in msec. |
| <b>MinPOP3Pass Latency(miliseccs)</b> | Minimum POP3PassLatency in msec.   |
| <b>MaxPOP3Pass Latency(miliseccs)</b> | Maximum POP3PassLatency in msec.   |
| <b>SuccPOP3Passes</b>                 | Total number of successful Pop3 Pass replies received.                             |
| <b>SuccPOP3PassRate (/sec)</b>        | Number of successful Pop3 Pass replies received per second.                        |
| <b>POP3Stats</b>                      | Total number of Pop3 Stat command packets sent.                                    |
| <b>POP3StatRate(/sec)</b>             | Number of Pop3 Stat command packets transmitted per second.                        |
| <b>POP3StatLatency (miliseccs)</b>    | Interval between transmitting a Pop3 Stat packet and receiving a response in msec. |
| <b>MinPOP3Stat Latency(miliseccs)</b> | Minimum POP3StartLatency in msec.  |
| <b>MaxPOP3Stat Latency(miliseccs)</b> | Maximum POP3StartLatency in msec.  |
| <b>SuccPOP3Stats</b>                  | Total number of successful Pop3 Stat replies received.                             |

| Measurement                           | Description  |
|---------------------------------------|--|
| <b>SuccPOP3StatRate (/sec)</b>        | Number of successful Pop3 Stat replies received per second.                        |
| <b>POP3Lists</b>                      | Total number of Pop3 List command packets transmitted.                             |
| <b>POP3ListRate(/sec)</b>             | Number of Pop3 List command packets transmitted per second.                        |
| <b>POP3ListLatency (milisechs)</b>    | Interval between transmitting a Pop3 List packet and receiving a response in msec. |
| <b>MinPOP3List Latency(milisechs)</b> | Minimum POP3ListLatency in msec.   |
| <b>MaxPOP3List Latency(milisechs)</b> | Maximum POP3ListLatency in msec.   |
| <b>SuccPOP3Lists</b>                  | Total number of successful Pop3Lists received.                                     |
| <b>SuccPOP3ListRate (/sec)</b>        | Number of successful Pop3Lists received per second.                                |
| <b>POP3Retrs</b>                      | Total number of Pop3 Retr packets transmitted.                                     |
| <b>POP3RetrRate(/sec)</b>             | Number of Pop3 Retr packets transmitted per second.                                |
| <b>POP3RetrLatency (milisechs)</b>    | Interval between transmitting a Pop3 Retr packet and receiving a response in msec. |
| <b>MinPOP3Retr Latency(milisechs)</b> | Minimum POP3RetrLatency in msec.   |
| <b>MaxPOP3Retr Latency(milisechs)</b> | Maximum POP3RetrLatency in msec.   |
| <b>SuccPOP3Retrs</b>                  | Total number of successful Pop3Retrs received.                                     |
| <b>SuccPOP3RetrRate (/sec)</b>        | Number of successful Pop3Retrs received per second.                                |

## DNS Performance Counters

| Measurement                          | Description   |
|--------------------------------------|---|
| <b>SuccPrimaryDNS Request</b>        | Total number of Successful DNS requests made to the Primary DNS server.         |
| <b>SuccSecondaryDNS Request</b>      | Total number of Successful DNS requests made to the Secondary DNS server.       |
| <b>SuccDNSData RequestRate(/sec)</b> | Number of Successful DNS Request packets transmitted per second.                |
| <b>PrimaryDNSFailure</b>             | Total number of DNS requests failures received from the Primary DNS server.     |
| <b>PrimaryDNSRequest</b>             | Total number of DNS requests made to the Primary DNS server.                    |
| <b>SecondaryDNS Failure</b>          | Total number of DNS requests failures received from the Secondary DNS server.   |
| <b>SecondaryDNS Request</b>          | Total number of DNS requests made to the Secondary DNS server.                  |
| <b>MinDNSData Latency</b>            | Minimum DNS Data Latency in msec.   |
| <b>MaxDNSData Latency</b>            | Maximum DNS Data Latency in msec.   |
| <b>CurDNSData Latency</b>            | Interval between sending a DNS request packet and receiving a response in msec. |
| <b>DNSDataRequest Rate(/sec)</b>     | Number of DNS Request packets transmitted per second.                           |
| <b>NoOf ReTransmission</b>           | Total number of DNS Request packets re  |
| <b>NoOfAnswers</b>                   | Total number of Answers to the DNS Request packets.                             |

### Attacks Performance Counters

| Measurement             | Description   |
|-------------------------|---|
| <b>Attacks</b>          | Total number of attack packets transmitted (All Attacks)                          |
| <b>AttackRate(/sec)</b> | Number of attack packets transmitted per second (ARP, Land, Ping, SYN, and Smurf) |
| <b>Havoc Flood</b>      | Number of Havoc packets generated (Stacheldraht only)                             |
| <b>Icmp Flood</b>       | Number of ICMP attack packets generated (TFN, TFN2K, & Stacheldraht)              |
| <b>Mix Flood</b>        | Number of Mix packets generated (TFN2K only)                                      |
| <b>Mstream Flood</b>    | Number of Mstream packets generated (Stacheldraht only)                           |
| <b>Null Flood</b>       | Number of Null packets generated (Stacheldraht only)                              |
| <b>Smurf Flood</b>      | Number of Smurf packets generated (TFN, TFN2K, & Stacheldraht)                    |
| <b>Syn Flood</b>        | Number of SYN packets generated (TFN, TFN2K, & Stacheldraht)                      |
| <b>Targa Flood</b>      | Number of Targa packets generated (TFN2K only)                                    |
| <b>Udp Flood</b>        | Number of UDP packets generated (All DDoS Attacks only)                           |

- Click **OK** in the Antara FlameThrower Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the Antara FlameThrower monitor.

## Configuring the SiteScope Monitor

You select measurements to poll from SiteScope using the SiteScope Monitor Configuration dialog box.

---

**Note:** SiteScope can be monitored by only one Console at a time.

---

### Before setting up the SiteScope monitor:

- 1** Make sure that SiteScope has been installed on a machine. If SiteScope is not installed on the same machine as the Console, verify that the SiteScope machine is accessible from the Console machine.
  - 2** On the machine where SiteScope is installed, configure SiteScope to monitor the required machines. When you assign a name to a monitor, include the hostname in the monitor name. This avoids confusion about the host to which the machine belongs.
- 

**Note:** SiteScope's default sampling rate is 10 minutes, and its minimum rate 15 seconds.

---

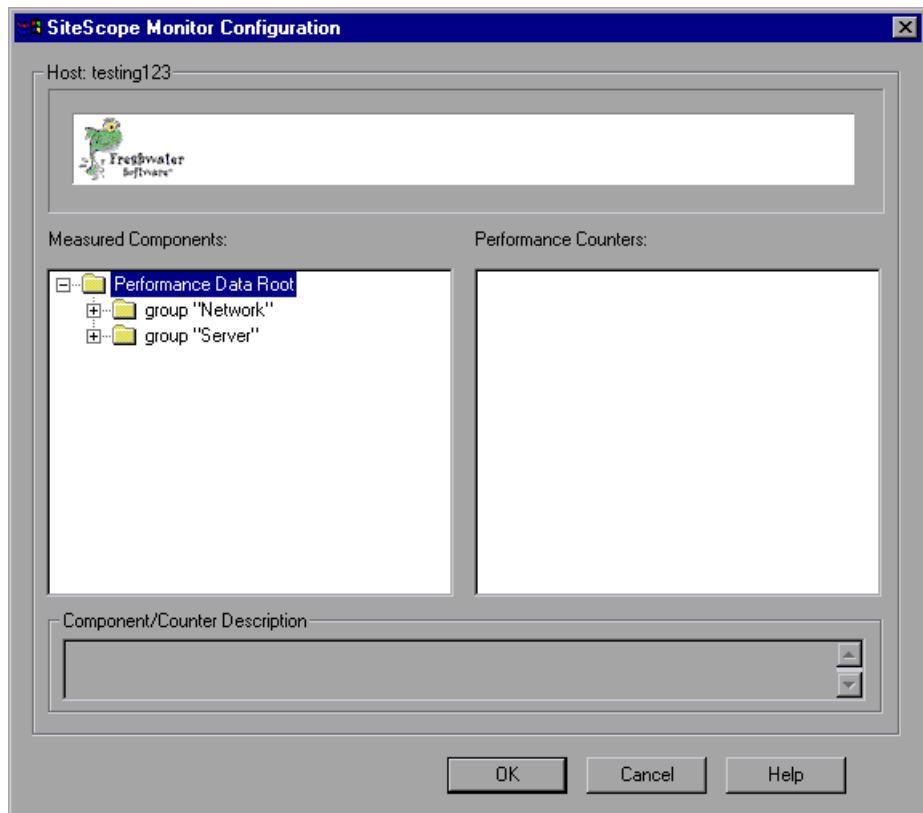
- 3** Verify that SiteScope is collecting the required data from the machines it is monitoring.
- 4** In the System Topology window, add a SiteScope Server element for the machine on which SiteScope is running.

You can configure the SiteScope monitor in one of the following ways:

- In the System Topology window
- In the **Execute** tab

**To configure the SiteScope monitor via the System Topology window:**

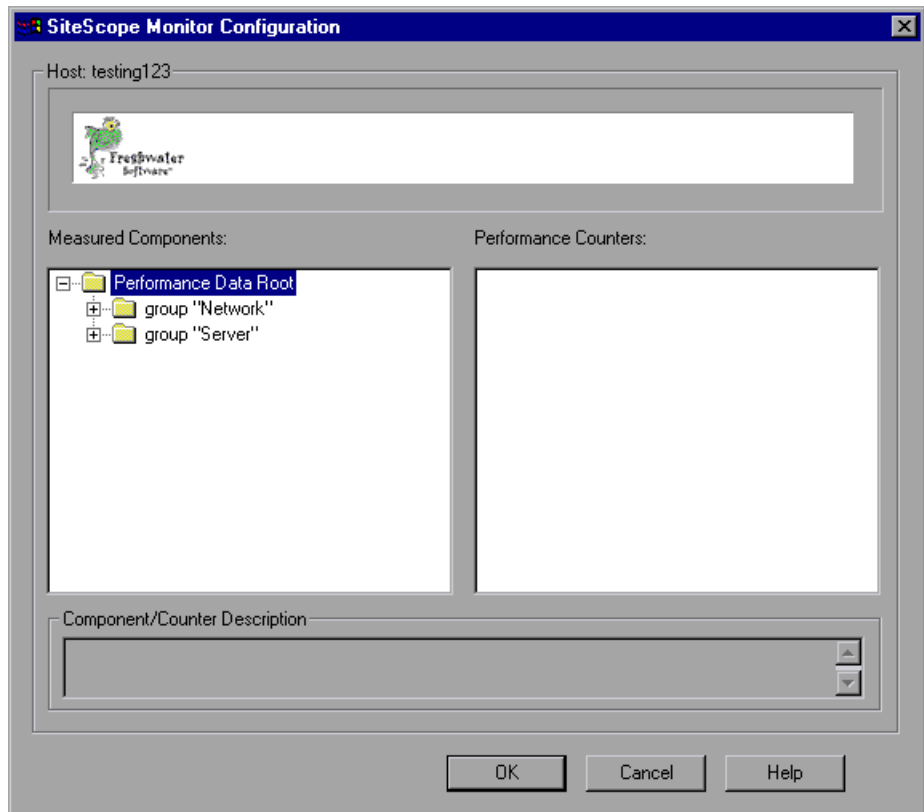
- 1** Click the SiteScope Server element to select it.
- 2** In the **Element Monitors** tab, click **Add**. The **Select Measurements to Monitor** dialog box appears.
- 3** Check the Show All Available Monitors box, expand the System Resources element, and then click SiteScope.
- 4** Click **Add**. The SiteScope Monitor Configuration dialog box is displayed.



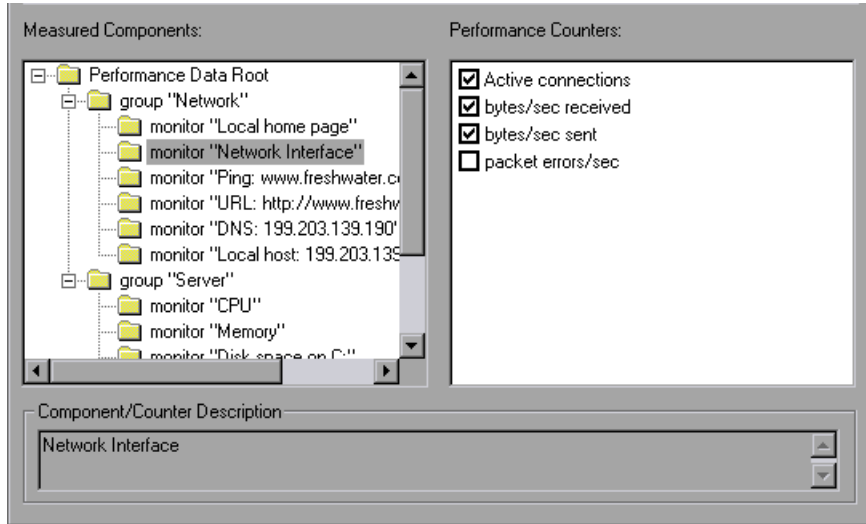
- 5** Continue with step 4 in the following section.

**To configure the SiteScope monitor via the Execute tab:**

- 1 Click the SiteScope graph in the Available Graphs tree, and drag it into the right pane of the **Execute** tab.
- 2 Right-click the SiteScope graph in the right pane, and choose **Monitors Configuration**, or click the **Monitors** button on the toolbar. The Monitors Configuration dialog box appears.
- 3 From the Server list, choose the server running SiteScope and click **Add**. The SiteScope Monitor Configuration dialog box is displayed.



- 4 In the Measured Components pane, locate the SiteScope measurement that you are monitoring and click it. The performance counters that SiteScope is monitoring on the selected component are displayed in the Performance Counters pane.



- 5 Check the required performance counters in the Performance Counters pane and click **OK**.
- 6 The **Select Measurements to Monitor** dialog box appears with the selected SiteScope measurements in the Selected Measurements pane. Click **Close** in the Select Measurements to Monitor dialog box.

---

**Note:** If you are activating the SiteScope monitor from the **Execute** tab, click **OK** in the **Monitors Configuration** dialog box to activate the monitor.

---



# 19

---

## Network Monitoring

You can use Network monitoring to determine whether your network is causing a delay in the session step. You can also determine the problematic network segment.

---

**Note:** You must have administrator privileges on the Windows source machine in order to run the Network monitor (unless you are using the ICMP protocol).

---

This chapter describes:


- ▶ Network Monitoring from a UNIX Source Machine
- ▶ Configuring the Network Monitor
- ▶ Viewing the Network Delay Time Graph

### About Network Monitoring

Network configuration is a primary factor in the performance of applications. A poorly designed network can slow client activity to unacceptable levels.

In a true Web or client/server system, there are many network segments. A single network segment with poor performance can affect the entire system.

The following diagram shows a typical network. In order to go from the server machine to the Vuser machine, data must travel over several segments.



To measure network performance, the Network monitor sends packets of data across the network. When a packet returns, the monitor calculates the time it takes for the packet to go to the requested node and return. This time is the delay which appears in the Network Delay Time graph.

Using the online Network Delay Time graph, you can locate the network-related problem so that it can be fixed.

---

**Note:** The delays from the source machine to each of the nodes are measured concurrently, yet independently. It is therefore possible that the delay from the source machine to one of the nodes could be greater than the delay for the complete path between the source and destination machines.

---

## Network Monitoring from a UNIX Source Machine

You can run the Network monitor on UNIX machines, using UDP or ICMP.

---

**Note:** If the server is running HP-UX, you need to use ICMP.

---

Before running the Network monitor from a UNIX source machine:

- ▶ configure the source machine by assigning root permissions to the *merc\_webtrace* process.
- ▶ make the necessary adjustments to either connect to the source machine through RSH, or through the agent.

### Configuring the Source Machine

To configure the source machine, where ProTune is installed locally:

To assign root permissions to the *merc\_webtrace* process, add an s-bit to *merc\_webtrace*'s permissions, as follows:

- 1** Log in to the source machine as root.
- 2** Type: `cd <ProTune>_installation/bin` to change to the *bin* directory.
- 3** Type: `chown root merc_webtrace` to make the root user the owner of the *merc\_webtrace* file.
- 4** Type: `chmod +s merc_webtrace` to add the s-bit to the file permissions.
- 5** To verify, type `ls -l merc_webtrace`. The permissions should look like: `-rwsrwsr-x`.

**To configure the source machine, where ProTune is installed on the network:**

In a ProTune network installation, the *merc\_webtrace* process is on the network, not on the source machine disk. The following procedure copies the *merc\_webtrace* file to the local disk, configures *mdrv.dat* to recognize the process, and assigns root permissions to *merc\_webtrace*:

- 1** Copy *merc\_webtrace* from *<ProTune>\_installation/bin* to anywhere on the local disk of the source machine. For example, to copy the file to the */local/ProTune* directory, type: `cp /net/tools/ProTune_installation/bin/merc_webtrace /local/ProTune`

---

**Note:** All of the source machines that use the same network installation must copy *merc\_webtrace* to the identical directory path on their local disk (for example, */local/ProTune*), since all of them use the same *mdrv.dat*.

---

- 2** Add the following line to the *ProTune\_installation/dat/mdrv.dat* file, in the [monitors\_server] section:

```
ExtCmdLine=-merc_webtrace_path /local/xxx
```

- 3** Log in to the source machine as root.
- 4** Type: `cd ProTune_installation/bin` to change to the *bin* directory.
- 5** Type: `chown root merc_webtrace` to make the root user the owner of the *merc\_webtrace* file.
- 6** Type: `chmod +s merc_webtrace` to add the s-bit to the file permissions.
- 7** To verify, type `ls -l merc_webtrace`. The permissions should look like: `-rwsrwsr-x`.

## Connecting to the Source Machine Through RSH

If the Console is connected to the source machine through RSH (default connection mode), then you don't need to activate the agent daemon. Before running the Network monitor the first time, you enter an encrypted user name and password in the Network monitor configuration file.

### To create an encrypted user name and password:

- 1 On the Console machine, type: `cd <ProTune>_installation/bin` to change to the *bin* directory.

---

**Note:** In a network or workstation installation, `<ProTune>_installation/bin` is the location on the network in which you installed the ProTune setup files.

---

- 2 Run *CryptonApp.exe*.
- 3 Type your RSH user name and password, separated by a vertical bar symbol. For example, `myname|mypw`.
- 4 Copy the encrypted string to the clipboard (highlight the string and click **ctrl+c**).
- 5 Add the following line to the `<ProTune>_installation/dat/monitors/ndm.cfg` file, in the [hosts] section:  
Host = <encrypted string copied from clipboard>
- 6 Close and open the current session step. ProTune will read the updated configuration file and recognize the source machine for monitoring.

## Connecting to the Source Machine Through the Agent

If the Console is not connected to the source machine through RSH, then make sure that the agent daemon is active on the source machine before running the Network monitor. For more information about working without RSH, refer to the section titled “UNIX Shell” in Appendix A, “Troubleshooting the Console.”

### To activate the agent daemon:

If you are not working in RSH, invoke the agent daemon on the source machine.

- 1 Type `m_daemon_setup -install` from the `<ProTune>_installation/bin` directory.
- 2 Make sure that the agent daemon is running whenever you activate the Network monitor.

---

**Tip:** To deactivate the agent daemon, type `m_daemon_setup -remove`.

---

## Configuring the Network Monitor

You configure the Network monitor from the Execute tab of the Console before you begin running a session step. Using the Network Delay Time and Add Destination Machines for Network Delay Monitoring dialog boxes, you select the network path you want to monitor.

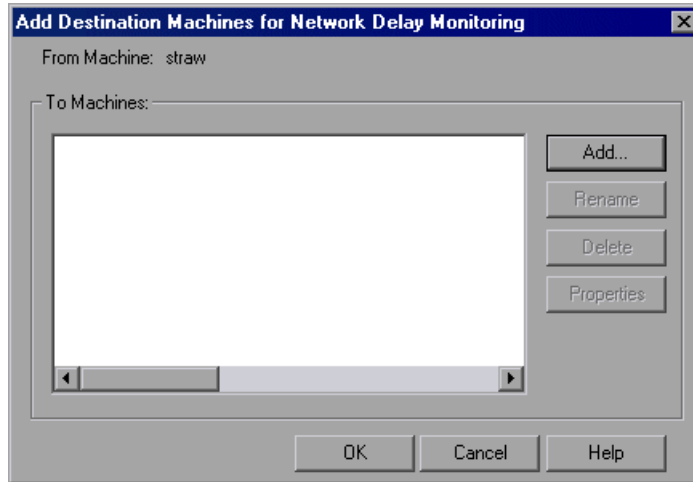
---

**Note:** To enable network monitoring, you must install the ProTune agent on the source machine. You do not have to install the ProTune agent on the destination machine.

---

**To configure the Network monitor:**

- 1** Click **Monitors** to display the Monitors Configuration dialog box, click **Add Monitor**, and check Show All Available Monitors.
- 2** In the list of monitors in the left section, expand the Load Generator category, click Network Delay, and then click **Add**. The Add Destination Machines for Network Delay Monitoring dialog box is displayed.



- 3** Click **Add**, enter the name of the destination machine, and click **OK**. The name of the machine appears in the Add Destination Machines for Network Delay Monitoring dialog box. Repeat this procedure for each path you want to monitor.

---


**Note:** If the destination machine is *localhost*, enter the local machine's name and not *localhost*.

---

To rename a machine, click **Rename**, and enter a new name for the machine.

To delete a machine, select it and click **Delete**.

- 4 Click **Properties** to configure additional network monitor settings. The Network Monitor Settings for Defined Path dialog box opens.



- 5 In the Monitor Settings box, select the protocol and enter the port number being used by the network path. The Network monitor supports three protocols: TCP, UDP, and ICMP. It is recommended that you use the default protocol. In Windows, the default is TCP, and in UNIX, the default is UDP.
- 6 Select **Enable display of network nodes by DNS names** if you want to view the DNS name of each node along the network path, in addition to its IP address. Note that selecting this option will decrease the speed of the Network monitor.
- 7 In the Monitoring Frequency box, select the number of milliseconds the monitor should wait between receiving a packet and sending out the next packet. The default value is 3000 milliseconds. If you have a long, steady session step, you can increase the interval by several seconds.



- 8 In the Monitoring Packet Retries box, select the maximum number of seconds that the monitor should wait for a packet to return before it retries to send the packet. The default value is 3 seconds. If your network is very large and loaded (an internet connection with a low capacity), you should increase the value by several seconds. If you have a small network (such as a LAN), you can decrease the value.

In addition, select the number of times the Network monitor should try resending a packet to a node if the packet is not initially returned. The default value is 0.

### **Network Monitoring over a Firewall**

If you are monitoring a network in which there are firewalls between the source and the destination machines, you must configure the firewalls to allow the network data packets to reach their destinations.

- ▶ If you are using the TCP protocol, the firewall that protects the destination machine should not block outgoing ICMP\_TIMEEXCEEDED packets (packets that are sent outside the firewall from the machine). In addition, the firewall protecting the source machine should allow ICMP\_TIMEEXCEEDED packets to enter, as well as TCP packets to exit.
- ▶ If you are using the ICMP protocol, the destination machine's firewall should not block incoming ICMP\_ECHO\_REQUEST packets, or outgoing ICMP\_ECHO\_REPLY and ICMP\_ECHO\_TIMEEXCEEDED packets. In addition, the firewall protecting the source machine should allow ICMP\_ECHO\_REPLY and ICMP\_ECHO\_TIMEEXCEEDED packets to enter, and ICMP\_ECHO\_REQUEST packets to exit.
- ▶ If you are using the UDP protocol, ensure that the UDP protocol can access the destination machine from the source machine. The destination machine's firewall should not block outgoing ICMP\_DEST\_UNREACHABLE and ICMP\_ECHO\_TIMEEXCEEDED packets. In addition, the firewall protecting the source machine should allow ICMP\_DEST\_UNREACHABLE and ICMP\_ECHO\_TIMEEXCEEDED packets to enter.

---


**Note:** To run the Network Delay Monitor when there are firewalls between the Console machine and the source machine, you must configure the ProTune agent, MI Listener, and Network monitor for monitoring over a firewall. For more information see “Configuring the ProTune Agents in LAN1,” on page 160, “Installing and Configuring the MI Listener in LAN2,” on page 169, and “Configuring the Network Delay Monitor over a Firewall,” on page 209.

---

## Viewing the Network Delay Time Graph

The **Network Delay Time** graph shows the delay for the complete path between the source and destination machines (y-axis) as a function of the elapsed session step time (x-axis).

Each path defined in the Add Destination Machines for Network Delay Monitoring dialog box is represented by a separate line with a different color in the graph.




To view the DNS names of the measurements displayed in the legend, right-click the graph and select **View as DNS Name**.

To view the delay time from the source machine to each of the nodes along the network path, right-click the graph and select **Configure**. In the Graph Configuration dialog box, click **SubPaths**.

In addition, you can view the delay time for each segment of the path.

**To view the delay time for the network segments:**

- 1 Right-click the Network Delay Time graph, and select **View Segments**. The Network Breakdown dialog box opens.



- 2 Select the path that you want to break down.
- 3 Choose whether you want to view the network segments of the graph of the graph you chose as an area graph or a pie graph.
- 4 Click **OK** to close the Network Breakdown dialog box. The delay time for the network segments of the path you chose is displayed in the graph view area.

---

**Note:** The segment delays are measured approximately, and do not add up to the network path delay which is measured exactly. The delay for each segment of the path is estimated by calculating the delay from the source machine to one node and subtracting the delay from the source machine to another node. For example, the delay for segment B to C is calculated by measuring the delay from the source machine to point C, and subtracting the delay from the source machine to point B.

---

To return to the complete path delay time view, select **Hide Segments** from the right-click menu.



# 20

---

## Firewall Server Performance Monitoring

During a session step run, you can monitor the firewall server in order to isolate server performance bottlenecks.

This chapter describes:

- ▶ Configuring the CheckPoint FireWall-1 Server Monitor

### About the Firewall Server Monitor

The Firewall server online monitor measures the performance of a Firewall server during session step execution. In order to obtain performance data, you need to activate the Firewall server monitor before executing the session step, and indicate which statistics and measurements you want to monitor.

### Configuring the CheckPoint FireWall-1 Server Monitor

To monitor the CheckPoint FireWall-1 server, you must select the counters you want the CheckPoint FireWall-1 server monitor to measure. You select these counters using the CheckPoint FireWall-1 SNMP Resources dialog box.

---

**Note:** You can specify a port number in the *snmp.cfg* file. If you do not specify a port number, ProTune connects to port 260, the default port for the CheckPoint FireWall-1 SNMP agent.

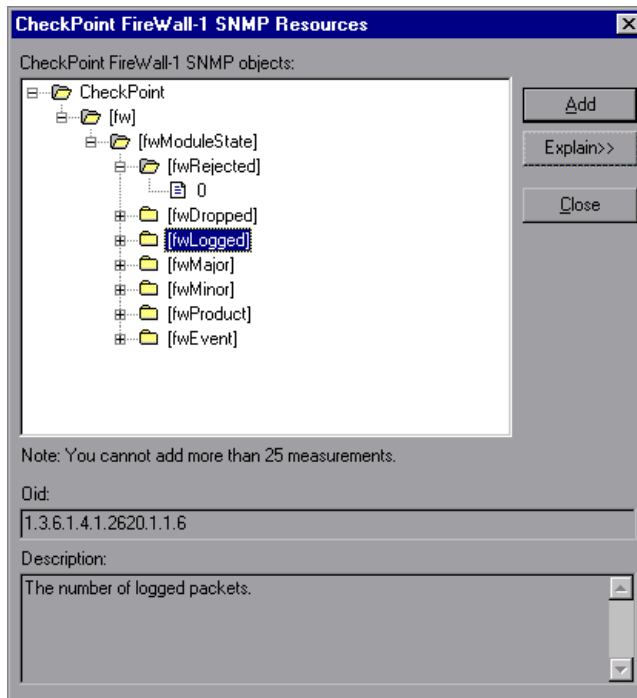
---

**To configure the CheckPoint FireWall-1 server monitor:**



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select CheckPoint FireWall-1 (in the Firewall category) and then click **Add**.

The CheckPoint FireWall-1 SNMP Resources dialog box opens.



- 5** Select the measurements you want to monitor. The following default counters can be monitored:

| Measurement | Description                     |
|-------------|---------------------------------|
| fwRejected  | The number of rejected packets. |
| fwDropped   | The number of dropped packets.  |
| fwLogged    | The number of logged packets.   |

- 6** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.

---


**Note:** The CheckPoint FireWall-1 monitor can only monitor up to 25 measurements.

---

- 7** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

**Note:** You can improve the level of measurement information for the CheckPoint FireWall-1 monitor by enabling measurements with string values to be listed (in addition to measurements with numeric values), and by enabling the name modifier (which displays the string value as an identifying part of the measurement name).

In the following example of a measurement using the name modifier, the string value of ProcessName (sched) is displayed in addition to its instance ID (0):



To enable this feature, add the following line to the *<ProTune root folder>\dat\monitors\snmp.cfg* file:  
SNMP\_show\_string\_nodes=1

Usage Notes: You can select more than one name modifier, but the first in the hierarchy will be used. Each time the CheckPoint FireWall-1 SNMP Resources dialog box opens, the information is reread from the *snmp.cfg* file. You cannot add the same measurement twice (once with a name modifier and once without it). If you do so, an error message is issued.

---



# 21

---

## Web Server Resource Monitoring

Using ProTune's Web Server Resource monitors, you can monitor the Apache, Microsoft IIS, iPlanet (SNMP), and iPlanet/Netscape servers during a session step run and isolate server performance bottlenecks.

This chapter describes:

- ▶ Configuring the Apache Monitor
- ▶ Configuring the Microsoft IIS Monitor
- ▶ Configuring the iPlanet/Netscape Monitor
- ▶ Configuring the iPlanet (SNMP) Monitor
- ▶ Monitoring Using a Proxy Server

### About Web Server Resource Monitors

Web Server Resource monitors provide you with information about the resource usage of the Apache, Microsoft IIS, iPlanet (SNMP), and iPlanet/Netscape Web servers during session step execution. In order to obtain this data, you need to activate the online monitor for the server and specify which resources you want to measure before executing the session step.

The procedures for selecting monitor measurements and configuring the monitors vary according to server type. The following sections contain specific configuration instructions for each server type.

**Note:** Certain measurements or counters are especially useful for determining server performance and isolating the cause of a bottleneck during an initial stress test on a Web server. For more information about these counters, see “Useful Counters for Stress Testing” on page 662.

---

## Configuring the Apache Monitor

To monitor an Apache server you need to know the server statistics information URL. A simple way to verify the statistics information URL is to try to view it through the browser.

The URL should be in the following format:

`http://<server name/IP address>:<port number>/server-status?auto`

For example:

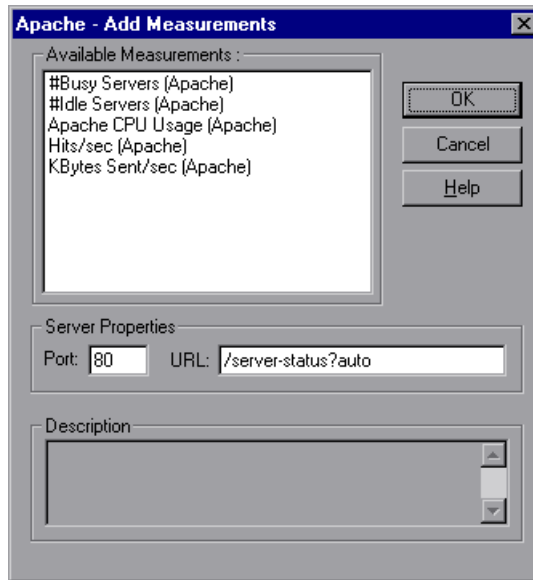
`http://stimpjy:80/server-status?auto`

### To configure the Apache monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Apache (in the Web Server category) and then click **Add**.

The Apache - Add Measurements dialog box opens, displaying the available measurements and server properties.



Select the required measurements. You can select multiple measurements using the **Ctrl** key.

The following table describes the measurements and server properties that can be monitored:

| Measurement      | Description   |
|------------------|---|
| # Busy Servers   | The number of servers in the Busy state                         |
| # Idle Servers   | The number of servers in the Idle state                         |
| Apache CPU Usage | The percentage of time the CPU is utilized by the Apache server |
| Hits/sec         | The HTTP request rate   |
| KBytes Sent/sec  | The rate at which data bytes are sent from the Web server       |

- 5 In the Server Properties section, enter the Port number and URL (without the server name), and click **OK**. The default URL is `/server-status?auto`.
- 6 Click **OK** in the Apache dialog box to activate the monitor.

---

**Note:** The default port number and URL can vary from one server to another. Please consult your Web server administrator.

---

**To change the default server properties:**

- 1 Open the `apache.cfg` file in the `<ProTune root folder>\dat\monitors\` directory.
- 2 Edit the following parameters after the `Delimiter=:` statement:

|                     |   |
|---------------------|---|
| <b>InfoURL</b>      | server statistics information URL   |
| <b>ServerPort</b>   | server port number  |
| <b>SamplingRate</b> | rate (milliseconds) at which the ProTune monitor will poll the server for the statistics information. If this value is greater than 1000, ProTune will use it as its sampling rate. Otherwise, it will use the sampling rate defined in the Monitors tab of the Options dialog box. |

---

**Note:** To monitor an Apache server through a firewall, use the Web server port (by default, port 80).

---

## Configuring the Microsoft IIS Monitor

You select measurements for the Microsoft IIS Server monitor using the MS IIS dialog box.

---

**Note:** To monitor an IIS server through a firewall, use TCP, port 139.

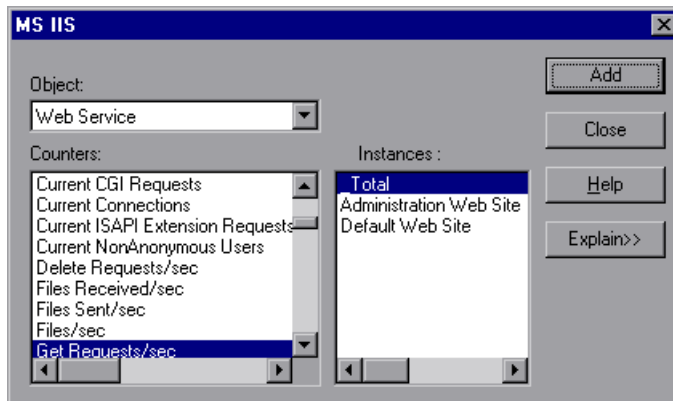
---

### To configure the IIS server monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select MS IIS (in the Web Server category) and then click **Add**.

A dialog box displaying the Web Service object, its counters, and instances opens.



The following table describes the default measurements that can be monitored:

| <b>Object</b>      | <b>Measurement</b>                | <b>Description</b>  |
|--------------------|-----------------------------------|---|
| <b>Web Service</b> | <b>Bytes Sent/sec</b>             | The rate at which the data bytes are sent by the Web service  |
| <b>Web Service</b> | <b>Bytes Received/sec</b>         | The rate at which the data bytes are received by the Web service  |
| <b>Web Service</b> | <b>Get Requests/sec</b>           | The rate at which HTTP requests using the GET method are made. Get requests are generally used for basic file retrievals or image maps, though they can be used with forms.                           |
| <b>Web Service</b> | <b>Post Requests/sec</b>          | The rate at which HTTP requests using the POST method are made. Post requests are generally used for forms or gateway requests.   |
| <b>Web Service</b> | <b>Maximum Connections</b>        | The maximum number of simultaneous connections established with the Web service   |
| <b>Web Service</b> | <b>Current Connections</b>        | The current number of connections established with the Web service  |
| <b>Web Service</b> | <b>Current NonAnonymous Users</b> | The number of users that currently have a non-anonymous connection using the Web service  |
| <b>Web Service</b> | <b>Not Found Errors/sec</b>       | The rate of errors due to requests that could not be satisfied by the server because the requested document could not be found. These are generally reported to the client as an HTTP 404 error code. |
| <b>Process</b>     | <b>Private Bytes</b>              | The current number of bytes that the process has allocated that cannot be shared with other processes.  |

---

**Note:** To change the default counters for the Microsoft IIS Server monitor, see “Changing a Monitor’s Default Counters” on page 661.

---

- 5** Select a counter and an instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.
- 6** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.
- 7** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the iPlanet/Netscape Monitor

To monitor an iPlanet/Netscape server, you need to know the administration server URL. A simple way to verify the administration server URL, is to try to view it through the browser.

The URL should be in the following format:

```
http://<admin_srv_name/IP address>:<port number>/https-<admin_srv_name/  
IP address>/bin/sitemon?doit
```

for example:

```
http://lazarus:12000/https-lazarus.mercury.co.il/bin/sitemon?doit
```

---

**Note:** In some server configurations, the URL must contain the administration server name and not the IP address.

In addition, the administration server name may differ from the iPlanet/Netscape server name.

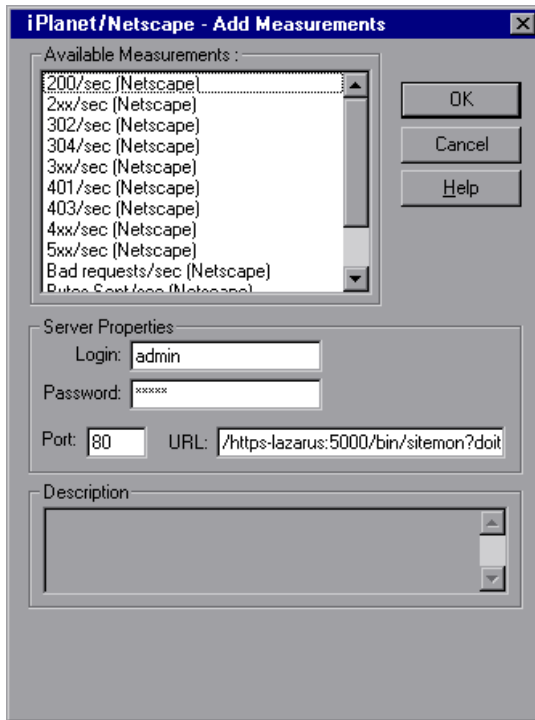
---

**To activate the iPlanet/Netscape monitor from the Console:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select iPlanet/Netscape (in the Web Server category) and then click **Add**.

The iPlanet/Netscape - Add Measurements dialog box opens, displaying the available measurements and server properties:



Select the required measurements. You can select multiple measurements using the **Ctrl** key.



The following table describes the measurements and server properties that can be monitored:

| Measurement      | Description   |
|------------------|---|
| 200/sec          | The rate of successful transactions being processed by the server   |
| 2xx/sec          | The rate at which the server handles status codes in the 200 to 299 range   |
| 302/sec          | The rate of relocated URLs being processed by the server  |
| 304/sec          | The rate of requests for which the server tells the user to use a local copy of a URL instead of retrieving a newer version from the server |
| 3xx/sec          | The rate at which the server handles status codes in the 300 to 399 range   |
| 401/sec          | The rate of unauthorized requests handled by the server   |
| 403/sec          | The rate of forbidden URL status codes handled by the server  |
| 4xx/sec          | The rate at which the server handles status codes in the 400 to 499 range   |
| 5xx/sec          | The rate at which the server handles status codes 500 and higher  |
| Bad requests/sec | The rate at which the server handles bad requests   |
| Bytes sent/sec   | The rate at which bytes of data are sent from the Web server  |
| Hits/sec         | The HTTP request rate   |
| xxx/sec          | The rate of all status codes (2xx-5xx) handled by the server, excluding timeouts and other errors that did not return an HTTP status code   |

##### 5 Fill in the Server Properties:

- Enter the user login name and password. The user must have administrator permissions on the server.
- Enter the port number and URL (without the server name), and click **OK**. The default URL is `/https-<admin_server>/bin/sitemon?doit`.

- 6 Click **OK** in the iPlanet/Netscape dialog box to activate the monitor.

---

**Note:** The default port number and URL can vary from one server to another. Please consult the Web server administrator. In some server configurations, the URL must contain the administration server name and not the IP address.

---

**To change the default server properties:**

- 1 Open the *Netscape.cfg* file in the `<ProTune root folder>\dat\monitors\` directory.
- 2 Edit the following parameters in the [Netscape] section:

|                       |   |
|-----------------------|---|
| <b>Counters</b>       | number of counters that the ProTune iPlanet/Netscape monitor will show you. This value should match the number of counters defined in the file.   |
| <b>InfoURL</b>        | server statistics information URL   |
| <b>ServerPort</b>     | server port number  |
| <b>ServerLogin</b>    | login name to the server  |
| <b>ServerPassword</b> | login password for the login name   |
| <b>SamplingRate</b>   | rate (milliseconds) at which the ProTune monitor will poll the server for the statistics information. If this value is greater than 1000, ProTune will use it as its sampling rate. Otherwise, it will use the sampling rate defined in the Monitors tab of the Options dialog box. |

---

**Note:** To monitor an iPlanet/Netscape server through a firewall, use the iPlanet/Netscape Administration server port. Configure this port during the server installation process.

---

## Configuring the iPlanet (SNMP) Monitor

The iPlanet (SNMP) monitor uses the Simple Network Management Protocol (SNMP) to retrieve iPlanet (SNMP) server statistics. You define the measurements for the iPlanet (SNMP) monitor using the iPlanet (SNMP) dialog box.

---

**Note:** To monitor a iPlanet (SNMP) server, use port 161 or 162, depending on the configuration of the agent.

---

### To configure the iPlanet (SNMP) Resources monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select iPlanet (SNMP) (in the Web Server category) and then click **Add**.

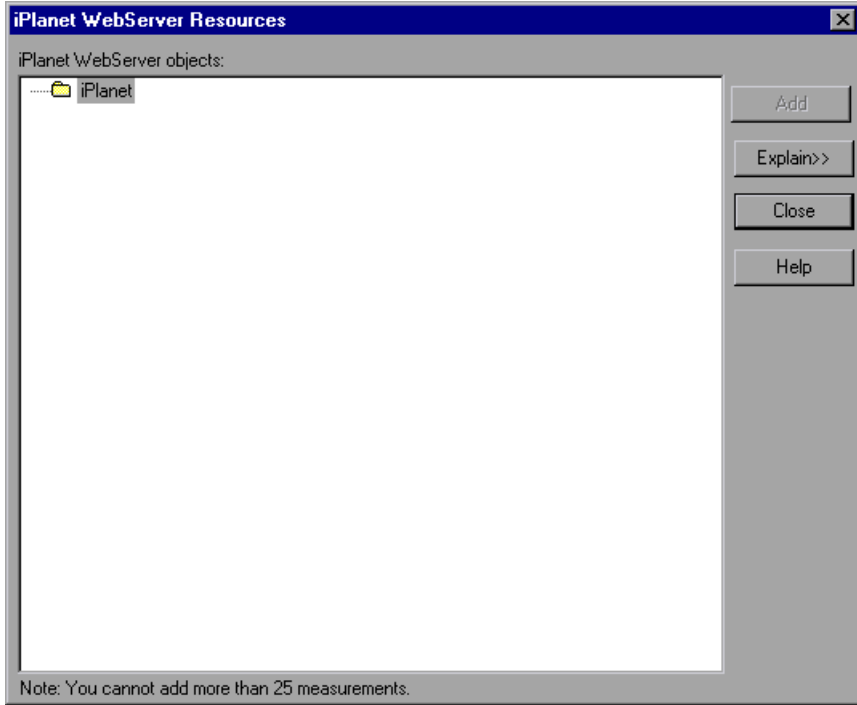
---

**Note:** You need to define the port number if the iPlanet SNMP agent is running on a different port than the default SNMP port. You can define the default port for your iPlanet server in the configuration file, *snmp.cfg*, located in `<ProTune root folder>\dat\monitors`. For example, if the port used by the SNMP agent on your iPlanet server is 8888, you should edit the *snmp.cfg* file as follows:

```
; iPlanet (WebServer)
[cm_snmp_mon_iws60]
port=8888
```

---

The iPlanet WebServer Resources dialog box opens.



**5** Browse the iPlanet WebServer Resources Object tree.

The following table describes the measurements and server properties that can be monitored:

| Measurement                   | Description                      |
|-------------------------------|----------------------------------|
| <b>iwsInstanceTable</b>       | iPlanet Web Server instances     |
| <b>iwsInstanceEntry</b>       | iPlanet Web Server instances     |
| <b>iwsInstanceIndex</b>       | Server instance index            |
| <b>iwsInstanceId</b>          | Server instance identifier       |
| <b>iwsInstanceVersion</b>     | Server instance software version |
| <b>iwsInstanceDescription</b> | Description of server instance   |

| Measurement                    | Description   |
|--------------------------------|---|
| <b>iwsInstanceOrganization</b> | Organization responsible for server instance                      |
| <b>iwsInstanceContact</b>      | Contact information for person(s) responsible for server instance |
| <b>iwsInstanceLocation</b>     | Location of server instance                                       |
| <b>iwsInstanceStatus</b>       | Server instance status  |
| <b>iwsInstanceUptime</b>       | Server instance uptime  |
| <b>iwsInstanceDeathCount</b>   | Number of times server instance processes have died               |
| <b>iwsInstanceRequests</b>     | Number of requests processed                                      |
| <b>iwsInstanceInOctets</b>     | Number of octets received   |
| <b>iwsInstanceOutOctets</b>    | Number of octets transmitted                                      |
| <b>iwsInstanceCount2xx</b>     | Number of 200-level (Successful) responses issued                 |
| <b>iwsInstanceCount3xx</b>     | Number of 300-level (Redirection) responses issued                |
| <b>iwsInstanceCount4xx</b>     | Number of 400-level (Client Error) responses issued               |
| <b>iwsInstanceCount5xx</b>     | Number of 500-level (Server Error) responses issued               |
| <b>iwsInstanceCountOther</b>   | Number of other (neither 2xx, 3xx, 4xx, nor 5xx) responses issued |
| <b>iwsInstanceCount200</b>     | Number of 200 (OK) responses issued                               |
| <b>iwsInstanceCount302</b>     | Number of 302 (Moved Temporarily) responses issued                |
| <b>iwsInstanceCount304</b>     | Number of 304 (Not Modified) responses issued                     |
| <b>iwsInstanceCount400</b>     | Number of 400 (Bad Request) responses issued                      |
| <b>iwsInstanceCount401</b>     | Number of 401 (Unauthorized) responses issued                     |
| <b>iwsInstanceCount403</b>     | Number of 403 (Forbidden) responses issued                        |

| Measurement                                | Description   |
|--|---|
| <b>iwsInstanceCount404</b>                 | Number of 404 (Not Found) responses issued                        |
| <b>iwsInstanceCount503</b>                 | Number of 503 (Unavailable) responses issued                      |
| <b>iwsInstanceLoad<br/>1MinuteAverage</b>  | System load average for 1 minute                                  |
| <b>iwsInstanceLoad<br/>5MinuteAverage</b>  | System load average for 5 minutes                                 |
| <b>iwsInstanceLoad<br/>15MinuteAverage</b> | System load average for 15 minutes                                |
| <b>iwsInstanceNetwork<br/>InOctets</b>     | Number of octets transmitted on the network per second            |
| <b>iwsInstanceNetwork<br/>OutOctets</b>    | Number of octets received on the network per second               |
| <b>iwsVsTable</b>                          | iPlanet Web Server virtual servers                                |
| <b>iwsVsEntry</b>                          | iPlanet Web Server virtual server                                 |
| <b>iwsVsIndex</b>                          | Virtual server index  |
| <b>iwsVsId</b>                             | Virtual server identifier   |
| <b>iwsVsRequests</b>                       | Number of requests processed                                      |
| <b>iwsVsInOctets</b>                       | Number of octets received   |
| <b>iwsVsOutOctets</b>                      | Number of octets transmitted                                      |
| <b>iwsVsCount2xx</b>                       | Number of 200-level (Successful) responses issued                 |
| <b>iwsVsCount3xx</b>                       | Number of 300-level (Redirection) responses issued                |
| <b>iwsVsCount4xx</b>                       | Number of 400-level (Client Error) responses issued               |
| <b>iwsVsCount5xx</b>                       | Number of 500-level (Server Error) responses issued               |
| <b>iwsVsCountOther</b>                     | Number of other (neither 2xx, 3xx, 4xx, nor 5xx) responses issued |

| Measurement                                | Description  |
|--|--|
| <b>iwsVsCount200</b>                       | Number of 200 (OK) responses issued                                |
| <b>iwsVsCount302</b>                       | Number of 302 (Moved Temporarily) responses issued                 |
| <b>iwsVsCount304</b>                       | Number of 304 (Not Modified) responses issued                      |
| <b>iwsVsCount400</b>                       | Number of 400 (Bad Request) responses issued                       |
| <b>iwsVsCount401</b>                       | Number of 401 (Unauthorized) responses issued                      |
| <b>iwsVsCount403</b>                       | Number of 403 (Forbidden) responses issued                         |
| <b>iwsVsCount404</b>                       | Number of 404 (Not Found) responses issued                         |
| <b>iwsVsCount503</b>                       | Number of 503 (Unavailable) responses issued                       |
| <b>iwsProcessTable</b>                     | iPlanet Web Server processes                                       |
| <b>iwsProcessEntry</b>                     | iPlanet Web Server process   |
| <b>iwsProcessIndex</b>                     | Process index  |
| <b>iwsProcessId</b>                        | Operating system process identifier                                |
| <b>iwsProcessThreadCount</b>               | Number of request processing threads                               |
| <b>iwsProcessThreadIdle</b>                | Number of request processing threads currently idle                |
| <b>iwsProcessConnection QueueCount</b>     | Number of connections currently in connection queue                |
| <b>iwsProcessConnection QueuePeak</b>      | Largest number of connections that have been queued simultaneously |
| <b>iwsProcessConnection QueueMax</b>       | Maximum number of connections allowed in connection queue          |
| <b>iwsProcessConnection QueueTotal</b>     | Number of connections that have been accepted                      |
| <b>iwsProcessConnection QueueOverflows</b> | Number of connections rejected due to connection queue overflow    |
| <b>iwsProcessKeepalive Count</b>           | Number of connections currently in keepalive queue                 |

| Measurement                                | Description   |
|--|---|
| <b>iwsProcessKeepaliveMax</b>              | Maximum number of connections allowed in keepalive queue        |
| <b>iwsProcessSizeVirtual</b>               | Process size in kbytes  |
| <b>iwsProcessSizeResident</b>              | Process resident size in kbytes                                 |
| <b>iwsProcessFractionSystemMemoryUsage</b> | Fraction of process memory in system memory                     |
| <b>iwsListenTable</b>                      | iPlanet Web Server listen sockets                               |
| <b>iwsListenEntry</b>                      | iPlanet Web Server listen socket                                |
| <b>iwsListenIndex</b>                      | Listen socket index   |
| <b>iwsListenId</b>                         | Listen socket identifier  |
| <b>iwsListenAddress</b>                    | Address socket is listening on                                  |
| <b>iwsListenPort</b>                       | Port socket is listening on                                     |
| <b>iwsListenSecurity</b>                   | Encryption support  |
| <b>iwsThreadPoolTable</b>                  | iPlanet Web Server thread pools                                 |
| <b>iwsThreadPoolEntry</b>                  | iPlanet Web Server thread pool                                  |
| <b>iwsThreadPoolIndex</b>                  | Thread pool index   |
| <b>iwsThreadPoolId</b>                     | Thread pool identifier  |
| <b>iwsThreadPoolCount</b>                  | Number of requests queued                                       |
| <b>iwsThreadPoolPeak</b>                   | Largest number of requests that have been queued simultaneously |
| <b>iwsThreadPoolMax</b>                    | Maximum number of requests allowed in queue                     |
| <b>iwsCpuTable</b>                         | iPlanet Web Server CPUs   |
| <b>iwsCpuEntry</b>                         | iPlanet Web Server CPU  |
| <b>iwsCpuIndex</b>                         | CPU index   |
| <b>iwsCpuId</b>                            | CPU identifier  |
| <b>iwsCpuIdleTime</b>                      | CPU Idle Time   |



| Measurement      | Description     |
|------------------|-----------------|
| iwsCpuUserTime   | CPU User Time   |
| iwsCpuKernelTime | CPU Kernel Time |

- 6 To measure an object, select it, and click **Add**. For a description of each resource, click **Explain>>** to expand the dialog box. Add all the desired resources to the list, and click **Close**.

---


**Note:** The iPlanet (SNMP) monitor can only monitor up to 25 measurements.

---

- 7 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

**Note:** You can improve the level of measurement information for the iPlanet (SNMP) monitor by enabling measurements with string values to be listed (in addition to measurements with numeric values), and by enabling the name modifier (which displays the string value as an identifying part of the measurement name).

In the following example of a measurement using the name modifier, the string value of ProcessName (sched) is displayed in addition to its instance ID (0):



To enable this feature, add the following line to the *<ProTune root folder>\dat\monitors\snmp.cfg* file:  
SNMP\_show\_string\_nodes=1

Usage Notes: You can select more than one name modifier, but the first in the hierarchy will be used. Each time the iPlanet SNMP Add Measurements dialog box opens, the information is reread from the *snmp.cfg* file. You cannot add the same measurement twice (once with a name modifier and once without it). If you do so, an error message is issued.

---

## Monitoring Using a Proxy Server

ProTune allows you to monitor using the Apache and Netscape monitors when there is a proxy server between the Console and the monitored server. To enable this, you must define settings in your configuration file: in *<LR root folder>\dat\monitors\apache.cfg* for the Apache monitor, or in *<LR root folder>\dat\monitors\Netscape.cfg* for the Netscape monitor.

Before defining settings, you need to determine whether you want ProTune to obtain proxy settings from your Internet Explorer connection configuration, or from the proxy settings in the configuration file.

### To have ProTune read proxy settings from your Internet Explorer connection:

- 1 In the Proxy Settings section of the configuration file, assign **useProxy** a value of 1.
- 2 If the proxy requires a username, password, or domain, enter these parameters on the lines **proxyUsername**, **proxyPassword**, and **proxyDomain**.

### To have ProTune read proxy settings from the configuration file:

- 1 In the Proxy Settings section of the configuration file, enter the proxy information on the **httpProxy** line. Use the format:  

```
[<protocol>=][<scheme>://]<proxy>[:<port>][[<protocol>=][<scheme>://]  
<proxy>[:<port>]]
```

For example:

```
httpProxy=http=http://my_http_proxy:8080 https=https://my_https_proxy:9000
```

- 2 If the proxy requires a username, password, or domain, enter these parameters on the lines **proxyUsername**, **proxyPassword**, and **proxyDomain**.

### To have ProTune connect directly to the server (any proxy settings are ignored):

In the Proxy Settings section of the configuration file, assign **useProxy** a value of 0.



# 22

---

## Web Application Server Resource Monitoring

You can monitor a Web application server during a session step run and isolate application server performance bottlenecks using ProTune's Web Application Server Resource monitors.

This chapter describes:

- ▶ Configuring the Ariba Monitor
- ▶ Configuring the ATG Dynamo Monitor
- ▶ Configuring the BroadVision Monitor
- ▶ Configuring the ColdFusion Monitor
- ▶ Configuring the Fujitsu INTERSTAGE Monitor
- ▶ Configuring the iPlanet (NAS) Monitor
- ▶ Configuring the Microsoft Active Server Pages Monitor
- ▶ Configuring the Oracle9iAS HTTP Monitor
- ▶ Configuring the SilverStream Monitor
- ▶ Configuring the WebLogic (SNMP) Monitor
- ▶ Configuring the WebLogic (JMX) Monitor
- ▶ Configuring the WebSphere Monitor
- ▶ Configuring the WebSphere (EPM) Monitor

## About Web Application Server Resource Monitors

Web Application Server Resource monitors provide you with information about the resource usage of the Ariba, ATG Dynamo, BroadVision, ColdFusion, Fujitsu INTERSTAGE, iPlanet (NAS), Microsoft ASP, Oracle9iAS HTTP, SilverStream, WebLogic (SNMP), WebLogic (JMX), and WebSphere application servers during session step execution. In order to obtain performance data, you need to activate the online monitor for the server and specify which resources you want to measure before executing the session step.

The procedures for selecting monitor measurements and configuring the monitors vary according to server type. The following sections contain specific configuration instructions for each server type.

### Configuring the Ariba Monitor

You select measurements to monitor the Ariba server using the Ariba Monitor Configuration dialog box.

---

**Note:** The port you use to monitor an Ariba server through a firewall depends on the configuration of your server.

---

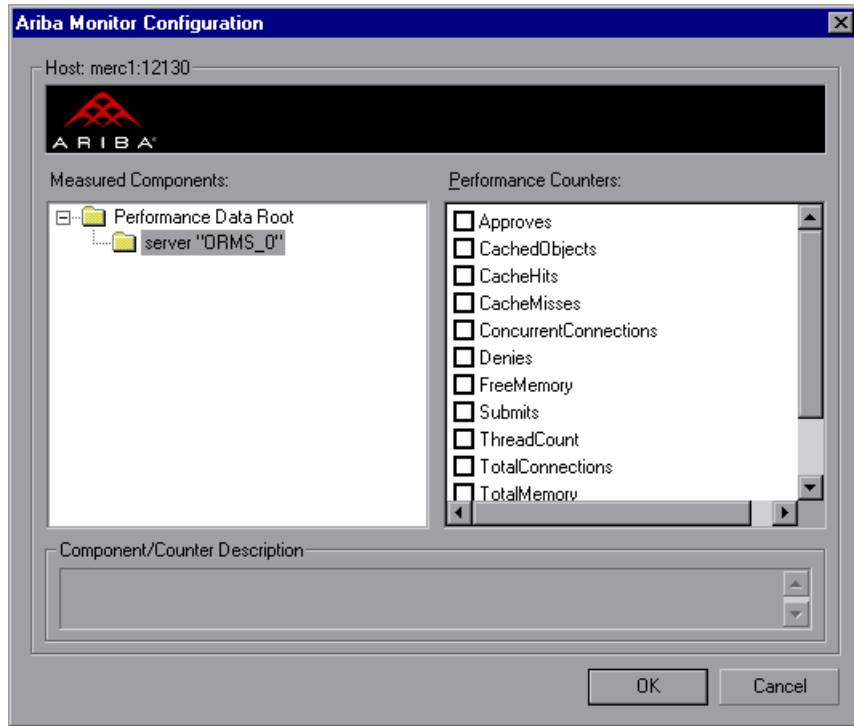
#### To configure the Ariba monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Ariba (in the Application Server category) and then click **Add**.

The Ariba Monitor Configuration dialog box opens.

5 Browse the Measured Components tree.



6 Check the required performance counters in the Ariba Monitor Configuration window's right pane.

The following tables describe the counters that can be monitored:

**Core Server Performance Counters**

| Measurement                     | Description   |
|---------------------------------|---|
| <b>Requisitions Finished</b>    | The instantaneous reading of the length of the worker queue at the moment this metric is obtained. The longer the worker queue, the more user requests are delayed for processing.        |
| <b>Worker Queue Length</b>      | The instantaneous reading of the length of the worker queue at the moment this metric is obtained. The longer the worker queue, the more user requests are delayed for processing.        |
| <b>Concurrent Connections</b>   | The instantaneous reading of the number of concurrent user connections at the moment this metric is obtained  |
| <b>Total Connections</b>        | The cumulative number of concurrent user connections since Ariba Buyer was started.   |
| <b>Total Memory</b>             | The instantaneous reading of the memory (in KB) being used by Ariba Buyer at the moment this metric is obtained   |
| <b>Free Memory</b>              | The instantaneous reading of the reserved memory (in KB) that is not currently in use at the moment this metric is obtained   |
| <b>Up Time</b>                  | The amount of time (in hours and minutes) that Ariba Buyer has been running since the previous time it was started  |
| <b>Number of Threads</b>        | The instantaneous reading of the number of server threads in existence at the moment this metric is obtained  |
| <b>Number of Cached Objects</b> | The instantaneous reading of the number of Ariba Buyer objects being held in memory at the moment this metric is obtained   |
| <b>Average Session Length</b>   | The average length of the user sessions (in seconds) of all users who logged out since previous sampling time. This value indicates on average how long a user stays connected to server. |



| Measurement           | Description   |
|-----------------------|---|
| Average Idle Time     | The average idle time (in seconds) for all the users who are active since previous sampling time. The idle time is the period of time between two consecutive user requests from the same user. |
| Approves              | The cumulative count of the number of approves that happened during the sampling period. An Approve consists of a user approving one Approvable.  |
| Submits               | The cumulative count of the number of Approvables submitted since previous sampling time  |
| Denies                | The cumulative count of the number of submitted Approvables denied since previous sampling time   |
| Object Cache Accesses | The cumulative count of accesses (both reads and writes) to the object cache since previous sampling time   |
| Object Cache Hits     | The cumulative count of accesses to the object cache that are successful (cache hits) since previous sampling time  |

### System Related Performance Counters

| Measurement                | Description   |
|----------------------------|---|
| Database Response Time     | The average response time (in seconds) to the database requests since the previous sampling time    |
| Buyer to DB server Traffic | The cumulative number of bytes that Ariba Buyer sent to DB server since the previous sampling time. |
| DB to Buyer server Traffic | The cumulative number of bytes that DB server sent to Ariba Buyer since the previous sampling time  |
| Database Query Packets     | The average number of packets that Ariba Buyer sent to DB server since the previous sampling time   |
| Database Response Packets  | The average number of packets that DB server sent to Ariba Buyer since the previous sampling time   |

- Click **OK** in the Ariba Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the Ariba monitor.

### XML Accessibility Verification

Only browsers that are XML-compatible will allow you to view the performance XML file.

#### To verify whether the XML file is accessible:

Display the XML file through the browser. The URL should be in the following format: `http://<server name>:<port number>/metrics?query=getStats`

For example: `http://merc1:12130/metrics?query=getStats`

---

**Note:** In some cases, although the browser is XML-compatible, it may still return the error: The XML page cannot be displayed. In these cases, the XML file can be accessed by the Ariba performance monitor, although it cannot be viewed by the browser.

---

## Configuring the ATG Dynamo Monitor

The ATG Dynamo monitor uses SNMP to retrieve ATG Dynamo server statistics. You define the measurements for the ATG Dynamo monitor using the ATG Dynamo Resources dialog box.

#### To configure the ATG Dynamo server monitor:



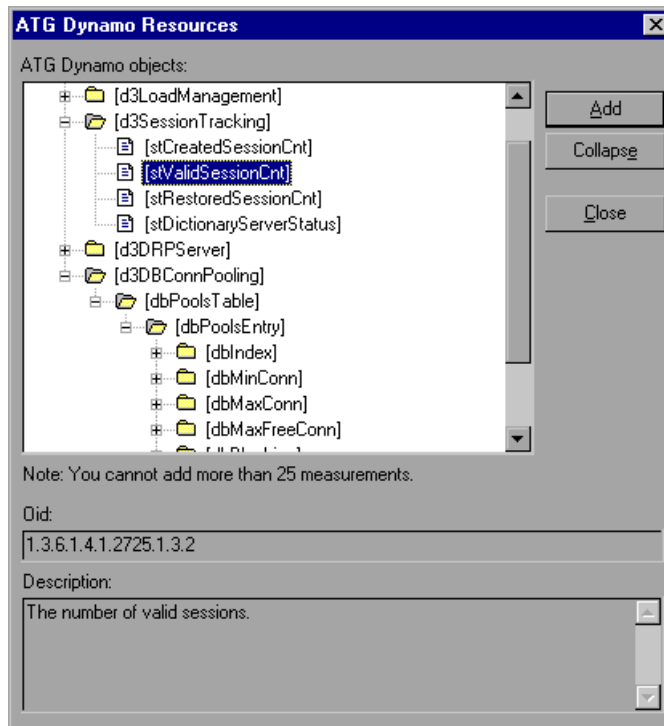
- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select ATG Dynamo (in the Application Server category) and then click **Add**.

**Note:** You need to define the port number if the ATG SNMP agent is running on a different port than the default ATG SNMP port 8870. You can define the default port for your ATG server in the configuration file, *snmp.cfg*, located in *<ProTune root folder>\dat\monitors*. For example, if the port used by the SNMP agent on your ATG system is 8888, you should edit the *snmp.cfg* file as follows:

```
; ATG Dynamo  
[cm_snmp_mon_atg]  
port=8888
```

The ATG Dynamo Resources dialog box opens.

- 5 Browse the ATG Dynamo Object tree, and select the measurements you want to monitor.



The following tables describe the measurements that can be monitored:

### d3System

| Measurement              | Description  |
|--------------------------|--|
| <b>sysTotalMem</b>       | The total amount of memory currently available for allocating objects, measured in bytes                           |
| <b>sysFreeMem</b>        | An approximation of the total amount of memory currently available for future allocated objects, measured in bytes |
| <b>sysNumInfoMsgs</b>    | The number of system global info messages written  |
| <b>sysNumWarningMsgs</b> | The number of system global warning messages written   |
| <b>sysNumErrorMsgs</b>   | The number of system global error messages written   |

### d3LoadManagement

| Measurement               | Description  |
|---------------------------|--|
| <b>lmIsManager</b>        | True if the Dynamo is running a load manager   |
| <b>lmManagerIndex</b>     | Returns the Dynamo's offset into the list of load managing entities  |
| <b>lmIsPrimaryManager</b> | True if the load manager is an acting primary manager  |
| <b>lmServicingCMs</b>     | True if the load manager has serviced any connection module requests in the amount of time set as the connection module polling interval |
| <b>lmCMLDRPPort</b>       | The port of the connection module agent  |
| <b>lmIndex</b>            | A unique value for each managed entity   |
| <b>lmSNMPPort</b>         | The port for the entry's SNMP agent  |
| <b>lmProbability</b>      | The probability that the entry will be given a new session   |

| Measurement                 | Description   |
|-----------------------------|---|
| <b>ImNewSessions</b>        | Indicates whether or not the entry is accepting new sessions, or if the load manager is allowing new sessions to be sent to the entry. This value is inclusive of any override indicated by ImNewSessionOverride. |
| <b>ImNewSessionOverride</b> | The override set for whether or not a server is accepting new sessions  |

### d3SessionTracking

| Measurement                     | Description                                   |
|---------------------------------|---|
| <b>stCreatedSessionCnt</b>      | The number of created sessions                |
| <b>stValidSessionCnt</b>        | The number of valid sessions                  |
| <b>stRestoredSessionCnt</b>     | The number of sessions migrated to the server |
| <b>StDictionaryServerStatus</b> | d3Session Tracking                            |

### d3DRPServer

| Measurement               | Description                                       |
|---------------------------|---|
| <b>drpPort</b>            | The port of the DRP server                        |
| <b>drpTotalReqsServed</b> | Total number of DRP requests serviced             |
| <b>drpTotalReqTime</b>    | Total service time in msec for all DRP requests   |
| <b>drpAvgReqTime</b>      | Average service time in msec for each DRP request |
| <b>drpNewessions</b>      | True if the Dynamo is accepting new sessions      |

### d3DBConnPooling

| Measurement         | Description  |
|---------------------|--|
| <b>dbPoolsEntry</b> | A pooling service entry containing information about the pool configuration and current status |
| <b>dbIndex</b>      | A unique value for each pooling service  |

| Measurement             | Description   |
|-------------------------|---|
| <b>dbPoolID</b>         | The name of the DB connection pool service  |
| <b>dbMinConn</b>        | The minimum number of connections pooled  |
| <b>dbMaxConn</b>        | The maximum number of connections pooled  |
| <b>dbMaxFreeConn</b>    | The maximum number of free pooled connections at a time   |
| <b>dbBlocking</b>       | Indicates whether or not the pool is to block out check outs  |
| <b>dbConnOut</b>        | Returns the number of connections checked out   |
| <b>dbFreeResources</b>  | Returns the number of free connections in the pool. This number refers to connections actually created that are not currently checked out. It does not include how many more connections are allowed to be created as set by the maximum number of connections allowed in the pool. |
| <b>dbTotalResources</b> | Returns the number of total connections in the pool. This number refers to connections actually created and is not an indication of how many more connections may be created and used in the pool.  |

- Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.

---

**Note:** The ATG Dynamo monitor can only monitor up to 25 measurements.


---

- Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

---

**Note:** You can improve the level of measurement information for the ATG Dynamo monitor by enabling measurements with string values to be listed (in addition to measurements with numeric values), and by enabling the name modifier (which displays the string value as an identifying part of the measurement name).

In the following example of a measurement using the name modifier, the string value of ProcessName (sched) is displayed in addition to its instance ID (0):



To enable this feature, add the following line to the *<ProTune root folder>\dat\monitors\snmp.cfg* file:  
SNMP\_show\_string\_nodes=1

**Usage Notes:** You can select more than one name modifier, but the first in the hierarchy will be used. Each time the ATG Dynamo Add Measurements dialog box opens, the information is reread from the *snmp.cfg* file. You cannot add the same measurement twice (once with a name modifier and once without it). If you do so, an error message is issued.

---

## Configuring the BroadVision Monitor

To monitor a BroadVision server, you must grant the client permission to invoke or launch services on the server.

---

**Note:** The port you use to monitor a BroadVision server through a firewall depends on the configuration of your server.

---

### To grant permission for a BroadVision server:

- Use the Iona Technologies (Orbix) command for setting user and access permission on a load generator machine:

```
chmodit [-h <host>] [-v] { <server> | -a <dir> }  
{i{+,-}{user,group} | l{+,-}{user,group} }
```

- If you experience problems connecting to the BroadVision monitor, you may need to redefine the permissions to "all."

To invoke permission for all, enter the following command at the BroadVision server command prompt:

```
# chmodit <server> i+all
```

To launch permission for all, enter the following command at the BroadVision server command prompt:

```
# chmodit <server> l+all
```

- Alternatively, set ORBIX\_ACL. Setting ORBIX\_ACL=i+all l+all in the BroadVision/Orbix configuration file gives permission to all.

In addition, to monitor a BroadVision server, you need to have JDK 1.2 or higher installed on the Console machine.

You can install JDK 1.2 by following the download and installation instructions at the following Web site: <http://java.sun.com/products/jdk/1.2/>

Before activating the monitor, make sure that your Java environment is configured properly.



**To configure your Java environment:**

- 1** Open the Windows Registry.
- 2** The registry should contain the correct path to the Java executable (java.exe) under the JDK 1.2 installation directory. Verify the following registry key:

HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\App Paths\java.exe

- 3** The registry should contain the correct path to the Java run-time environment (JRE) under the JRE 1.2 installation directory. Verify the following registry key:

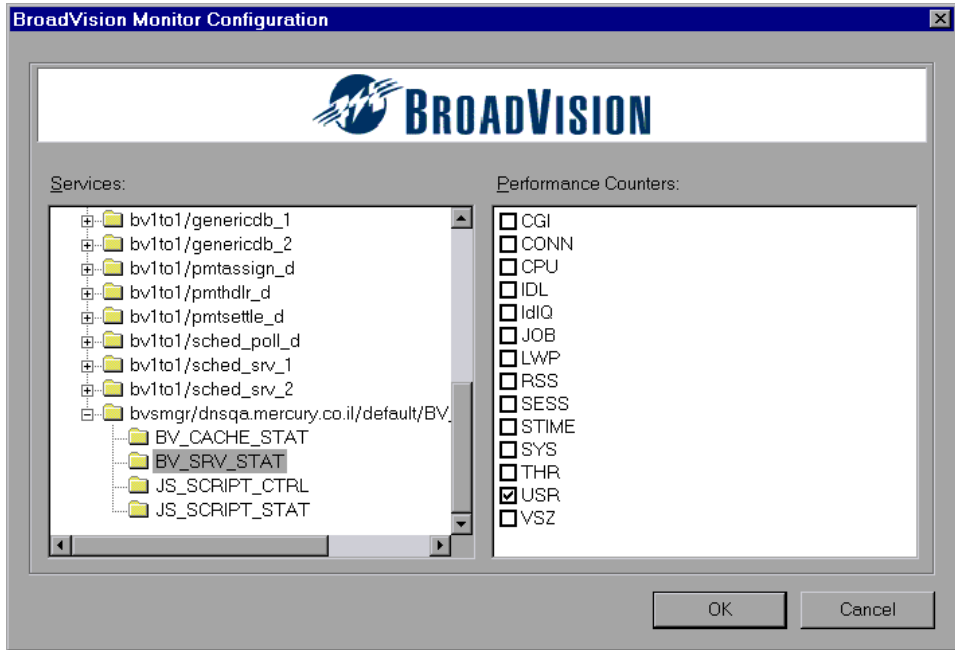
HKEY\_LOCAL\_MACHINE\SOFTWARE\JavaSoft\Java Runtime Environment\1.2\JavaHome

**To configure the BroadVision online monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select BroadVision (in the Application Server category) and then click **Add**.

The BroadVision Monitor Configuration dialog box opens, displaying the available measurements:



- 5 Browse the Services tree and check the required performance counters in the BroadVision Monitor Configuration window's right pane. For a description of the performance counters, see the information below.
- 6 Click **OK** in the BroadVision Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the BroadVision monitor.

The following table describes the servers/services that can be monitored:

| Server    | Multiple Instances | Description   |
|-----------|--------------------|---|
| adm_srv   | No                 | One-To-One user administration server. There must be one.           |
| alert_srv | No                 | Alert server handles direct IDL function calls to the Alert system. |

| Server              | Multiple Instances | Description  |
|---------------------|--------------------|--|
| <b>bvconf_srv</b>   | No                 | One-To-One configuration management server. There must be one.   |
| <b>cmsdb</b>        | Yes                | Visitor management database server.  |
| <b>cntdb</b>        | Yes                | Content database server.   |
| <b>deliv_smtp_d</b> | Yes                | Notification delivery server for e-mail type messages. Each instance of this server must have its own ID, numbered sequentially starting with "1".   |
| <b>deliv_comp_d</b> | No                 | Notification delivery completion processor.  |
| <b>extdbacc</b>     | Yes                | External database accessor. You need at least one for each external data source.   |
| <b>genericdb</b>    | No                 | Generic database accessor handles content query requests from applications, when specifically called from the application. This is also used by the One-To-One Command Center.   |
| <b>hostmgr</b>      | Yes                | Defines a host manager process for each machine that participates in One-To-One, but doesn't run any One-To-One servers. For example, you need a hostmgr on a machine that runs only servers. You don't need a separate hostmgr on a machine that already has one of the servers in this list. |
| <b>g1_ofbe_srv</b>  | No                 | Order fulfillment back-end server.   |
| <b>g1_ofdb</b>      | Yes                | Order fulfillment database server.   |
| <b>g1_om_srv</b>    | No                 | Order management server.   |
| <b>pmtassign_d</b>  | No                 | The payment archiving daemon routes payment records to the archives by periodically checking the invoices table, looking for records with completed payment transactions, and then moving those records into an archive table.   |

| Server                    | Multiple Instances | Description  |
|---------------------------|--------------------|--|
| <code>pmthdlr_d</code>    | Yes                | For each payment processing method, you need one or more authorization daemons to periodically acquire the authorization when a request is made.                                 |
| <code>pmtsettle_d</code>  | Yes                | Payment settlement daemon periodically checks the database for orders of the associated payment processing method that need to be settled, and then authorizes the transactions. |
| <code>sched_poll_d</code> | No                 | Notification schedule poller scans the database tables to determine when a notification must be run.   |
| <code>sched_srv</code>    | Yes                | Notification schedule server runs the scripts that generate the visitor notification messages.   |

### Performance Counters

Performance counters for each server/service are divided into logical groups according to the service type.

The following section describes all the available counters under each group. Note that the same group can have a different number of counters, depending on the service.

Counter groups:

- BV\_DB\_STAT
- BV\_SRV\_CTRL
- BV\_SRV\_STAT
- NS\_STAT
- BV\_CACHE\_STAT
- JS\_SCRIPT\_CTRL
- JS\_SCRIPT\_STAT

#### **BV\_DB\_STAT**

The database accessor processes have additional statistics available from the BV\_DB\_STAT memory block. These statistics provide information about database accesses, including the count of selects, updates, inserts, deletes, and stored procedure executions.

- DELETE - Count of deletes executions
- INSERT - Count of inserts executions
- SELECT - Count of selects executions
- SPROC - Count of stored procedure executions.
- UPDATE - Count of updates executions

#### **BV\_SRV\_CTRL**

- SHUTDOWN

## NS\_STAT

The NS process displays the namespace for the current One-To-One environment, and optionally can update objects in a name space.

- Bind
- List
- New
- Rebnd
- Rsvlv
- Unbnd

## BV\_SRV\_STAT

The display for Interaction Manager processes includes information about the current count of sessions, connections, idle sessions, threads in use, and count of CGI requests processed.

- **HOST** - Host machine running the process.
- **ID** - Instance of the process (of which multiple can be configured in the *bv1to1.conf* file), or engine ID of the Interaction Manager.
- **CGI** - Current count of CGI requests processed.
- **CONN** - Current count of connections.
- **CPU** - CPU percentage consumed by this process. If a process is using most of the CPU time, consider moving it to another host, or creating an additional process, possibly running on another machine. Both of these specifications are done in the *bv1to1.conf* file. The CPU % reported is against a single processor. If a server is taking up a whole CPU on a 4 processor machine, this statistic will report 100%, while the Windows Task Manager will report 25%. The value reported by this statistic is consistent with "% Processor Time" on the Windows Performance Monitor.
- **GROUP** - Process group (which is defined in the *bv1to1.conf* file), or Interaction Manager application name.

- **STIME** - Start time of server. The start times should be relatively close. Later times might be an indication that a server crashed and was automatically restarted.
- **IDL** - Total count of IDL requests received, not including those to the monitor.
- **IdIQ**
- **JOB**
- **LWP** - Number of light-weight processes (threads).
- **RSS** - Resident memory size of server process (in kilobytes).
- **STIME** - System start time.
- **SESS** - Current count of sessions.
- **SYS** - Accumulated system mode CPU time (seconds).
- **THR** - Current count of threads.
- **USR** - Accumulated user mode CPU time (seconds).
- **VSZ** - Virtual memory size of server process (in kilobytes). If a process is growing in size, it probably has a memory leak. If it is an Interaction Manager process, the culprit is most likely a component or dynamic object (though Interaction Manager servers do grow and shrink from garbage collection during normal use).

#### **BV\_CACHE\_STAT**

Monitors the request cache status.

The available counters for each request are:

- **CNT- Request\_Name-HIT** - Count of requests found in the cache.
- **CNT- Request\_Name-MAX** - Maximum size of the cache in bytes
- **CNT- Request\_Name-SWAP** - Count of items that got swapped out of the cache.
- **CNT- Request\_Name-MISS** - Count of requests that were not in the cache.
- **CNT- Request\_Name-SIZE** - Count of items currently in the cache.

## Cache Metrics

Cache metrics are available for the following items:

- ▶ **AD**
- ▶ **ALERTSCHED** - Notification schedules are defined in the BV\_ALERTSCHED and BV\_MSGSCHED tables. They are defined by the One-To-One Command Center user or by an application.
- ▶ **CATEGORY\_CONTENT**
- ▶ **DISCUSSION** - The One-To-One discussion groups provide moderated system of messages and threads of messages aligned to a particular topic. Use the Discussion group interfaces for creating, retrieving and deleting individual messages in a discussion group. To create, delete, or retrieve discussion groups, use the generic content management API. The BV\_DiscussionDB object provides access to the threads and messages in the discussion group database.
- ▶ **EXT\_FIN\_PRODUCT**
- ▶ **EDITORIAL** - Using the Editorials content module, you can point cast and community cast personalized editorial content, and sell published text on your One-To-One site. You can solicit editorial content, such as investment reports and weekly columns, from outside authors and publishers, and create your own articles, reviews, reports, and other informative media. In addition to text, you can use images, sounds, music, and video presentations as editorial content.
- ▶ **INCENTIVE** - Contains sales incentives
- ▶ **MSGSCHED** - Contains the specifications of visitor-message jobs. Notification schedules are defined in the BV\_ALERTSCHED and BV\_MSGSCHED tables. They are defined by the One-To-One Command Center user or by an application.
- ▶ **MSGSCRIPT** - Contains the descriptions of the JavaScripts that generate visitor messages and alert messages. Contains the descriptions of the JavaScripts that generate targeted messages and alert messages. Use the Command Center to add message script information to this table by selecting the Visitor Messages module in the Notifications group. For more information, see the Command Center User's Guide.



- **PRODUCT** - BV\_PRODUCT contains information about the products that a visitor can purchase.
- **QUERY** - BV\_QUERY contains queries.
- **SCRIPT** - BV\_SCRIPT contains page scripts.
- **SECURITIES**
- **TEMPLATE** - The Templates content module enables you to store in the content database any BroadVision page templates used on your One-To-One site. Combining BroadVision page templates with BroadVision dynamic objects in the One-To-One Design Center application is one way for site developers to create One-To-One Web sites. If your developers use these page templates, you can use the Command Center to enter and manage them in your content database. If your site doesn't use BroadVision page template, you will not use this content module.

#### **JS\_SCRIPT\_CTRL**

- CACHE
- DUMP
- FLUSH
- METER
- TRACE

#### **JS\_SCRIPT\_STAT**

- ALLOC
- ERROR
- FAIL
- JSPERR
- RELEASE
- STOP
- SUCC
- SYNTAX

## Configuring the ColdFusion Monitor

You select measurements to monitor the ColdFusion server using the ColdFusion dialog box.

---

**Note:** The ColdFusion monitor works via HTTP and supports UNIX platforms. If you want to monitor the ColdFusion server on Windows platforms, you can also use the Windows Resource monitor.

---

### To set up the ColdFusion monitor environment:

Copy the `<ProTune installation>\dat\monitors\perfmon.cfm` file into the `<ColdFusion Home>\cfide\administrator` directory. By default, the ColdFusion monitor checks for the `<ColdFusion Home>\cfide\administrator\perfmon.cfm` file.

---

**Note:** The port you use to monitor a ColdFusion server through a firewall depends on the configuration of your server.

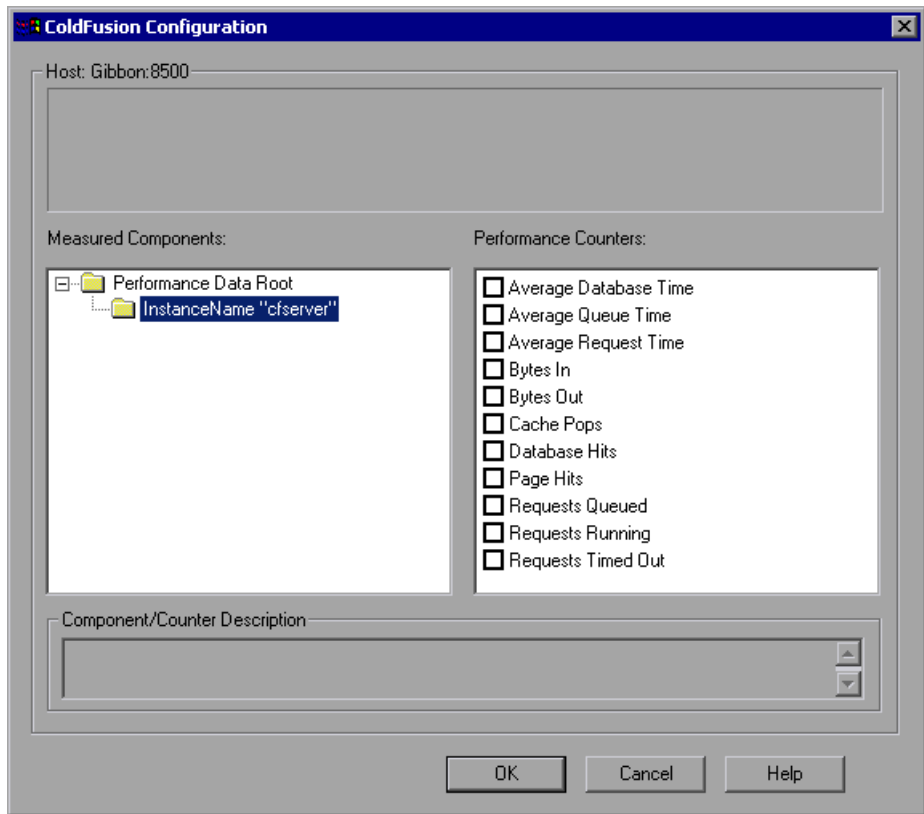
---

### To configure the ColdFusion monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select ColdFusion (in the Application Server category) and then click **Add**.
- 5** If you are prompted for a username and password, enter them and click **OK**.

The ColdFusion Configuration dialog box displays the available measurements.

**6** Browse the Measured Components tree.**7** Check the required performance counters in the ColdFusion Monitor Configuration window's right pane.

The following table describes the default counters that can be measured:

| Measurement                      | Description   |
|----------------------------------|---|
| <b>Avg. Database Time (msec)</b> | The running average of the amount of time, in milliseconds, that it takes ColdFusion to process database requests.  |
| <b>Avg. Queue Time (msec)</b>    | The running average of the amount of time, in milliseconds, that requests spent waiting in the ColdFusion input queue before ColdFusion began to process the request.   |
| <b>Avg Req Time (msec)</b>       | The running average of the total amount of time, in milliseconds, that it takes ColdFusion to process a request. In addition to general page processing time, this value includes both queue time and database processing time. |
| <b>Bytes In/sec</b>              | The number of bytes per second sent to the ColdFusion server.   |
| <b>Bytes Out/sec</b>             | The number of bytes per second returned by the ColdFusion server.   |
| <b>Cache Pops</b>                | Cache pops.   |
| <b>Database Hits/sec</b>         | This is the number of database hits generated per second by the ColdFusion server.  |
| <b>Page Hits/sec</b>             | This is the number of Web pages processed per second by the ColdFusion server.  |
| <b>Queued Requests</b>           | The number of requests currently waiting to be processed by the ColdFusion server.  |
| <b>Running Requests</b>          | The number of requests currently being actively processed by the ColdFusion server.   |
| <b>Timed Out Requests</b>        | The number of requests that timed out due to inactivity timeouts.   |

- 8 Click **OK** in the ColdFusion Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the ColdFusion monitor.

## Configuring the Fujitsu INTERSTAGE Monitor

The Fujitsu INTERSTAGE monitor uses SNMP to retrieve Fujitsu INTERSTAGE server statistics. You define the measurements for the Fujitsu INTERSTAGE monitor using the Fujitsu INTERSTAGE SNMP Resources dialog box.

**To configure the Fujitsu INTERSTAGE server monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Fujitsu INTERSTAGE (in the Application Server category) and then click **Add**.

---

**Note:** You need to define the port number if the Fujitsu INTERSTAGE SNMP agent is running on a different port than the default SNMP port 161. You can define the default port for your Fujitsu INTERSTAGE server in the configuration file, *snmp.cfg*, located in *<ProTune root folder>\dat\monitors*. For example, if the port used by the SNMP agent on your Fujitsu INTERSTAGE system is 8888, you should edit the *snmp.cfg* file as follows:

```

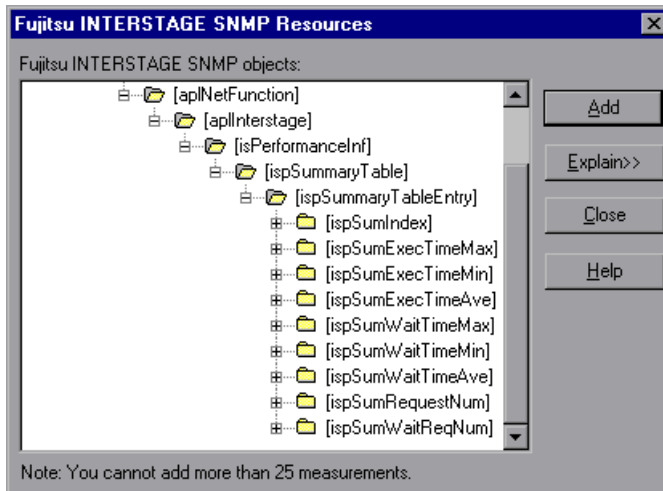
; Fujitsu INTERSTAGE
[cm_snmp_mon_isp]
port=8888

```

---

The Fujitsu INTERSTAGE SNMP Resources dialog box opens.

- Browse the Fujitsu INTERSTAGE SNMP Object tree, and select the measurements you want to monitor.



The following tables describe the measurements that can be monitored:

| Measurement              | Description  |
|--------------------------|--|
| <b>IspSumObjectName</b>  | The object name of the application for which performance information is measured                 |
| <b>IspSumExecTimeMax</b> | The maximum processing time of the application within a certain period of time                   |
| <b>IspSumExecTimeMin</b> | The minimum processing time of the application within a certain period of time                   |
| <b>IspSumExecTimeAve</b> | The average processing time of the application within a certain period of time                   |
| <b>IspSumWaitTimeMax</b> | The maximum time required for INTERSTAGE to start an application after a start request is issued |
| <b>IspSumWaitTimeMin</b> | The minimum time required for INTERSTAGE to start an application after a start request is issued |
| <b>IspSumWaitTimeAve</b> | The average time required for INTERSTAGE to start an application after a start request is issued |

| Measurement      | Description  |
|------------------|--|
| IspSumRequestNum | The number of requests to start an application         |
| IspSumWaitReqNum | The number of requests awaiting application activation |

- 6** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.

---

**Note:** The Fujitsu INTERSTAGE monitor can only monitor up to 25 measurements.

---

- 7** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the iPlanet (NAS) Monitor

The iPlanet (NAS) monitor uses the SNMP to retrieve iPlanet (NAS) server statistics. You define the measurements for the iPlanet (NAS) monitor using the iPlanet (NAS) dialog box.

### Setting up the iPlanet (NAS) Application Server

This section offers a short explanation on setting up SNMP monitoring of the iPlanet Application Server. It is intended to supplement the iPlanet documentation, not act as a replacement. For an explanation of the SNMP reporting architecture and theory, refer to the iPlanet documentation.

---

**Note:** The instructions below assume that SNMP statistics will be collected on the standard SNMP port 161.

---

### **SNMP Summary**

- Solaris has a native SNMP agent, "snmpdx", that is started automatically at boot time by the script /etc/rc3.d/S76snmpdx. This daemon communicates on the standard SNMP port 161. The port number can be changed with the -p <port> option.
- Planet Products are shipped with their own SNMP agents. The architecture is such that there is one "master agent" per host, which a network management station communicates with, and one or more "subagents" that collect data from various iPlanet products and forward statistics to the master agent. The master agent also defaults to communicating on port 161.
- To run both the Solaris SNMP agent and the iPlanet SNMP agent, a proxy must be used that makes the Sun agent look like a subagent to the iPlanet master agent.

### **Steps Overview**

- Login to the system as root.
- Change the port number for the Solaris SNMP agent.
- Configure and run the iPlanet agents "magt" and "sagt".
- Start the Solaris SNMP agent.
- Configure iPlanet Application Server for SNMP statistics.
- Start SNMP subagents for iPlanet Directory Server and iPlanet Web Server (optional).

### **To change the port number for the Solaris SNMP agent:**

- 1** Login to the system as root. (Only a root user can change the port number and run the agents).
- 2** Stop the SNMP agent by running /etc/rc2.d/K76snmpdx stop.
- 3** Edit /etc/rc3.d/S76snmpdx to run the Solaris daemon on a non-standard port number. For example, 1161:  
Replace  
    /usr/lib/snmp/snmpdx -y -c /etc/snmp/conf  
with  
    /usr/lib/snmp/snmpdx -p 1161 -y -c /etc/snmp/conf



**To configure and run the iPlanet agents "magt" and "sagt":**

The master and proxy agents and startup scripts are found in *<ias install directory>/snmp*.

- 1** In the script `S75snmpagt`, add a line to the environment variable `GX_ROOTDIR` so that it points to your iAS installation. For example, if the iPlanet Application Server is installed in `/usr/iplanet/ias6/ias`:

```
GX_ROOTDIR=/usr/iplanet/ias6/ias
exprt GX_ROOTDIR
```

- 2** Copy the script `S75snmpagt` to `/etc/rc3.d`
- 3** `chmod 755 /etc/rc3.d/S75snmpagt`
- 4** In `/etc/rc3.d/S75snmpagt` `/etc/rc2.d/K07snmpagt`
- 5** You can configure system information and traps.

In the example below, information has been added about the system owner and location, and SNMP traps have been sent to a network manager station ("mde.uk.sun.com").

```
COMMUNITY    public
ALLOW ALL OPERATIONS
INITIAL sysLocation "Under Joe Bloggs' Desk in Headquarters"
INITIAL sysContact "Joe Bloggs
Email: Joe.Bloggs@Sun.COM
Voice: +1 650 555 1212"
MANAGER mde.uk.sun.com
SEND ALL TRAPS TO PORT 162
WITH COMMUNITY public
```

---

**Note:** There is no need to edit the proxy agent's configuration file (`CONFIG_SAGT`).

---

- 6** Start the iPlanet agents by running the command:  
`/etc/rc3.d/S75snmpagt start`

**To start the Solaris SNMP agent:**

Restart the Solaris SNMP agent by running the command:  
`/etc/rc3.d/S76snmpdx start`

**To configure iPlanet Application Server for SNMP statistics:**

- 1** Start the iPlanet Application Server admin tool `ksvradmin`.
- 2** In the General View, select the instance name that you want to manage.
- 3** Click the **SNMP** tab in the management frame.
- 4** Select **Enable SNMP Administration and Monitoring** and **Enable SNMP Debug**.
- 5** Type "60" in the Connection Attempt Interval field, and exit `ksvradmin`.
- 6** Restart the iPlanet Application Server with the commands:  

```
iascontrol stop
iascontrol kill
iascontrol start
```
- 7** Check in the log file `<ias install directory>/logs/ias.log` that the application server successfully connected to the master agent. You should see the following line:  
`kas> SNMP: Connected to master agent`

**To start SNMP subagents for iPlanet Web Server:**

- 1** Use your Web browser to access the iPlanet Web Server.
- 2** Choose the Web server you wish to administer, and click the **Manage** button.
- 3** Select the **Monitor** tab, and click **SNMP Subagent Configuration** on the left side of the page.
- 4** Type in the configuration information and set the radio button **Enable SNMP Statistics Collection** to "On".
- 5** Click **SNMP Subagent Control**.
- 6** Click the **Start** button.

**To start SNMP subagents for iPlanet Directory Server:**

- 1** Use the Netscape Administration Console to manage the iPlanet Directory Server.
- 2** Select the **Configuration** tab.
- 3** Click the **SNMP** tab in the Configuration frame.
- 4** Select the **Enable statistics collection** check box.
- 5** Set "Master Host" to "localhost".
- 6** Set "Master port" to 199.
- 7** In the other fields, enter the appropriate information.
- 8** Click the **Start Subagent** button.

**Summary note**

Use your SNMP management tool to query the SNMP master agent on port 161. You should see all the information provided by the Solaris SNMP agent, as well as any iPlanet subagents that you have configured.

The next time that you boot Solaris, the Sun and iPlanet SNMP agents will be started automatically by the boot scripts which you have configured.

**Configuring the iPlanet (NAS) Monitor in the Console**

Once you have configured the iPlanet SNMP Service, you must select the counters that you want the iPlanet (NAS) monitor to measure. You select these measurements using the iPlanet (NAS) Resources dialog box.

**To configure the iPlanet (NAS) Resources monitor:**

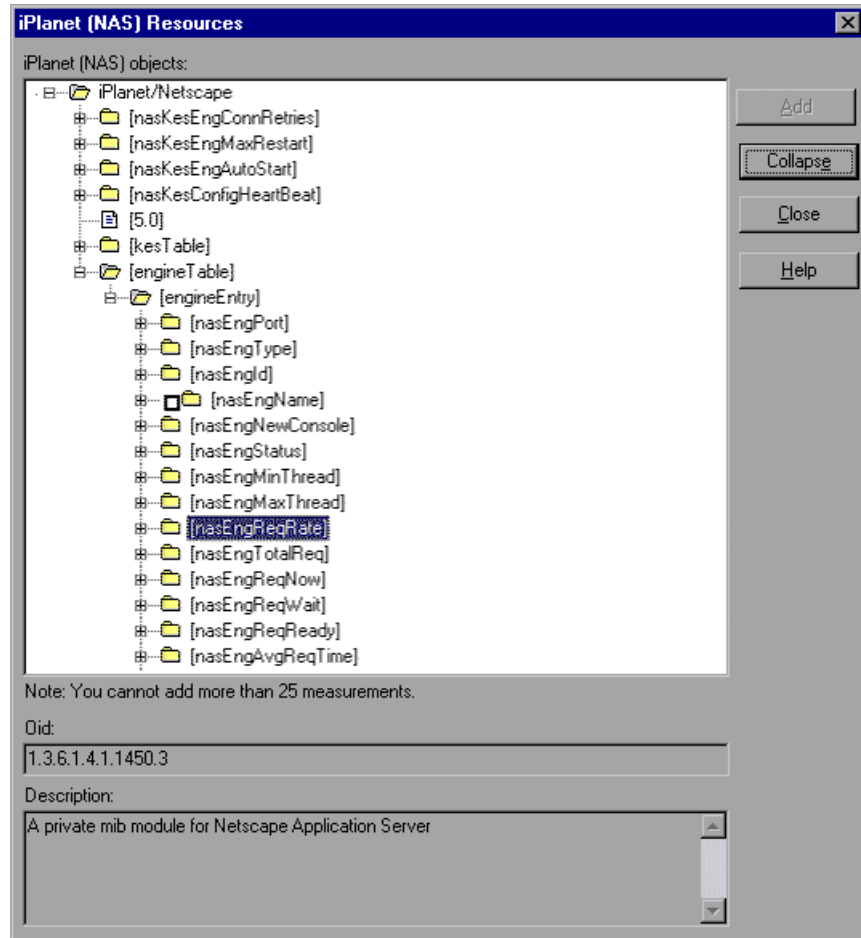
- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select iPlanet (NAS) (in the Application Server category) and then click **Add**.

**Note:** You need to define the port number if the iPlanet SNMP agent is running on a different port than the default SNMP port. You can define the default port for your iPlanet server in the configuration file, *snmp.cfg*, located in *<ProTune root folder>\dat\monitors*. For example, if the port used by the SNMP agent on your iPlanet server is 8888, you should edit the *snmp.cfg* file as follows:

```
; iPlanet (NAS)
[cm_snmp_mon_nas]
port=8888
```

---

The iPlanet (NAS) Resources dialog box opens.



**5** Browse the iPlanet (NAS) Resources Object tree.

The following tables describe the counters that can be monitored:

### Netscape Performance Counters

| Measurement                  | Description   |
|------------------------------|---|
| <b>nasKesEngConnRetries</b>  | The maximum number of times the administration server will try to connect to an engine.       |
| <b>nasKesEngMaxRestart</b>   | The maximum number of times the administration server will restart an engine after a failure. |
| <b>nasKesEngAutoStart</b>    | Start all the engines at startup of the administration server.                                |
| <b>nasKesConfigHeartBeat</b> | Heart Beat.   |

### KES Performance Counters

| Measurement                      | Description  |
|----------------------------------|--|
| <b>nasKesId</b>                  | The ID of the KES this engine belongs to.  |
| <b>nasKesMinThread</b>           | The default minimum number of threads per engine.  |
| <b>nasKesMaxThread</b>           | The default maximum number of threads per engine.  |
| <b>nasKesLoadBalancerDisable</b> | Enable or Disable the load balancer service.   |
| <b>nasKesCpuLoad</b>             | The total CPU usage on this host.  |
| <b>nasKesDiskLoad</b>            | The total disk usage on this host.   |
| <b>nasKesMemLoad</b>             | The total memory usage on this host.   |
| <b>nasKesRequestLoad</b>         | The number of requests on this NAS.  |
| <b>nasKesCpuLoadFactor</b>       | The relative importance of CPU usage in computing the server load. This number is specified as a percent. The sum of all server load factors, CPULoad, DiskLoad, MemLoad and ExecReqs must equal 100%. |

| Measurement                         | Description  |
|-------------------------------------|--|
| <b>nasKesDiskLoadFactor</b>         | The relative importance of Disk usage in computing the server load. This number is specified as a percent. The sum of all server load factors, CPUload, DiskLoad, MemLoad and ExecReqs must equal 100%.  |
| <b>nasKesMemLoadFactor</b>          | The relative importance of Memory usage in computing the server load. This number is specified as a percent. The sum of all server load factors, CPUload, DiskLoad, MemLoad and ExecReqs must equal 100%.  |
| <b>nasKesAppLogicsRunningFactor</b> | The relative importance of the number of times an AppLogic is run in computing the AppLogic execution performance. This figure is specified as a percent. The sum of all agent load factors, ResultCached, AvgExecTime, LastExecTime, and ServerLoad must equal 100%.    |
| <b>nasKesResultsCacheFactor</b>     | The relative importance of the cached results of an AppLogic in computing the AppLogic execution performance. This figure is specified as a percent. The sum of all agent load factors, ResultCached, AvgExecTime, LastExecTime, and ServerLoad must equal 100%.         |
| <b>nasKesAvgExecTimeFactor</b>      | The relative importance of the average execution time of an AppLogic in computing the AppLogic execution performance. This figure is specified as a percent. The sum of all agent load factors, ResultCached, AvgExecTime, LastExecTime, and ServerLoad must equal 100%. |
| <b>nasKesLastExecTimeFactor</b>     | The relative importance of the last execution time of an AppLogic in computing the AppLogic execution performance. This figure is specified as a percent. The sum of all agent load factors, ResultCached, AvgExecTime, LastExecTime, and ServerLoad must equal 100%.    |

| Measurement                            | Description  |
|--|--|
| <b>nasKesHitsFactor</b>                | The relative importance of the number of AppLogics running in computing the AppLogic execution performance. This figure is specified as a percent. The sum of all agent load factors, ResultCached, AvgExecTime, LastExecTime, and ServerLoad must equal 100%. |
| <b>nasKesServerLoadFactor</b>          | The relative importance of the server load (computed using the four server load factors) in computing AppLogic execution performance. The sum of all agent load factors, ResultCached, AvgExecTime, LastExecTime, and ServerLoad must equal 100%.              |
| <b>nasKesBroadcastInterval</b>         | The length of time in seconds, between each broadcast attempt from the load balancer daemon.   |
| <b>nasKesApplogicBroadcastInterval</b> | The length of time in seconds, between each broadcast of AppLogics load information across all the server in the cluster. This should be greater than nasKesBroadcastInterval.   |
| <b>nasKesServerBroadcastInterval</b>   | The length of time in seconds, between each broadcast of server load information across all the server in the cluster. This should be greater than nasKesBroadcastInterval.  |
| <b>nasKesServerLoadUpdateInterval</b>  | The length of time in seconds between each update of server load informations. A server load update applies the server load data that has been sampled up until the moment when the update occurs.   |
| <b>nasKesCpuLoadUpdateInterval</b>     | The length of time, in seconds, between each sampling of CPU usage.  |
| <b>nasKesDiskLoadUpdateInterval</b>    | The length of time, in seconds, between each sampling of disk usage.   |
| <b>nasKesMemLoadUpdateInterval</b>     | The length of time, in seconds, between each sampling of memory thrashes.  |
| <b>nasKesTotalReqsUpdateInterval</b>   | The length of time, in seconds, between each sampling of the number of requests.   |
| <b>nasKesMaxHops</b>                   | The maximum number of times a request can be load-balanced to another server.  |



| Measurement                         | Description   |
|-------------------------------------|---|
| <b>nasKesODBCReqMinThread</b>       | The minimum number of threads reserved to process asynchronous requests.  |
| <b>nasKesODBCReqMaxThread</b>       | The maximum number of threads reserved to process asynchronous requests.  |
| <b>nasKesODBCCacheMaxConns</b>      | The maximum number of connections opened between NAS and the database.  |
| <b>nasKesODBCCacheFreeSlots</b>     | The minimum number of cached connections established between NAS and the database.  |
| <b>nasKesODBCCacheTimeout</b>       | The time after which an idle connection is dropped.   |
| <b>nasKesODBCCacheInterval</b>      | The interval in seconds at which the cache cleaner will try to disconnect connections already idle for longer than the specified timeout. |
| <b>nasKesODBCConnGivetimeupTime</b> | Maximum time the driver will try to connect to the database.  |
| <b>nasKesODBCCacheDebug</b>         | Turns on the connection cache debug information.  |
| <b>nasKesODBCResultSetInitRows</b>  | The number of rows fetched at once from the database.   |
| <b>nasKesODBCResultSetMaxRows</b>   | The maximum number of rows the cached result set can contain.   |
| <b>nasKesODBCResultSetMaxSize</b>   | The maximum size of result set the driver will cache  |
| <b>nasKesODBCSqlDebug</b>           | Turns on SQL debug information.   |
| <b>nasKesODBCEnableParser</b>       | Turns on SQL parsing.   |
| <b>nasKesORCLReqMinThread</b>       | The minimum number of threads reserved to process asynchronous requests.  |
| <b>nasKesORCLReqMaxThread</b>       | The maximum number of threads reserved to process asynchronous requests.  |

| Measurement                        | Description   |
|------------------------------------|---|
| <b>nasKesORCLCacheMaxConns</b>     | The maximum number of connections opened between NAS and the database.  |
| <b>nasKesORCLCacheFreeSlots</b>    | The minimum number of cached connections established between NAS and the database.  |
| <b>nasKesORCLCacheTimeout</b>      | The time after which an idle connection is dropped.   |
| <b>nasKesORCLCacheInterval</b>     | The interval in seconds at which the cache cleaner will try to disconnect connections already idle for longer than the specified timeout. |
| <b>nasKesORCLConnGiveupTime</b>    | The maximum time the driver will spend trying to obtain a connection to Oracle.   |
| <b>nasKesORCLCacheDebug</b>        | Turns on the connection cache debug information.  |
| <b>nasKesORCLResultSetInitRows</b> | The number of rows fetched at once from the database.   |
| <b>nasKesORCLResultSetMaxRows</b>  | The maximum number of rows the cached result set can contain.   |
| <b>nasKesORCLResultSetMaxSize</b>  | The maximum size of result set the driver will cache.   |
| <b>nasKesORCLSqlDebug</b>          | Turns on SQL debug information.   |
| <b>nasKesSYBReqMinThread</b>       | The minimum number of threads reserved to process asynchronous requests.  |
| <b>nasKesSYBReqMaxThread</b>       | The maximum number of threads reserved to process asynchronous request.   |
| <b>nasKesSYBCacheMaxConns</b>      | The maximum number of connections opened between NAS and the database.  |
| <b>nasKesSYBCacheFreeSlots</b>     | The minimum number of cached connections established between NAS and the database.  |
| <b>nasKesSYBCacheTimeout</b>       | The time after which an idle connection is dropped.   |

| Measurement                       | Description  |
|-----------------------------------|--|
| <b>nasKesSYBCacheInterval</b>     | The interval time between cached connections.  |
| <b>nasKesSYBConnGiveupTime</b>    | The maximum time the driver will spend trying to obtain a connection to Sybase before giving up. |
| <b>nasKesSYBCacheDebug</b>        | Turns on the connection cache debug information.   |
| <b>nasKesSYBResultSetInitRows</b> | The number of rows fetched at once from the database.  |
| <b>nasKesSYBResultSetMaxRows</b>  | The maximum number of rows the cached result set can contain.                                    |
| <b>nasKesSYBResultSetMaxSize</b>  | The maximum size of result set the driver will cache.  |

### Engine Performance Counters

| Measurement             | Description  |
|-------------------------|--|
| <b>nasEngKesPort</b>    | The port of the KXS this engine serves. This is supplied as part of the object ID and cannot be modified after creation.                   |
| <b>nasEngPort</b>       | The TCP/IP port this engine is listening on. The port can only be specified at the creation of the engine. It is not allowed to modify it. |
| <b>nasEngType</b>       | Type of the engine: executive(0), Java(1000), C++(3000).   |
| <b>nasEngId</b>         | The ID is an incremental number starting at 0. The ID cannot be modified.  |
| <b>nasEngName</b>       | The name of this engine. This is an informational string that contains kcs, kxs or kjs.  |
| <b>nasEngNewConsole</b> | Starts each engine in a new console window.  |
| <b>nasEngStatus</b>     | The status column used to add, remove, enable or disable an engine. To create an engine, one needs to set. This follows rfc1443.           |

| Measurement                 | Description   |
|-----------------------------|---|
| <b>nasEngMinThread</b>      | The default minimum number of threads per engine.                         |
| <b>nasEngMaxThread</b>      | The default maximum number of threads per engine.                         |
| <b>nasEngReqRate</b>        | The rate at which requests arrive.  |
| <b>nasEngTotalReq</b>       | The total number of requests processed since engine startup.              |
| <b>nasEngReqNow</b>         | The number of requests being processed.                                   |
| <b>nasEngReqWait</b>        | The requests waiting to be serviced.                                      |
| <b>nasEngReqReady</b>       | The requests that are ready to be serviced.                               |
| <b>nasEngAvgReqTime</b>     | The average request processing time.                                      |
| <b>nasEngThreadNow</b>      | Number of threads in use by the request manager.                          |
| <b>nasEngThreadWait</b>     | The number of idle threads.   |
| <b>nasEngWebReqQueue</b>    | The number of web requests that are queued.                               |
| <b>nasEngFailedReq</b>      | The number of requests that failed.                                       |
| <b>nasEngTotalConn</b>      | The total number of connections opened.                                   |
| <b>nasEngTotalConnNow</b>   | The total number of connections in use.                                   |
| <b>nasEngTotalAccept</b>    | The total number of connections listening to incoming requests.           |
| <b>nasEngTotalAcceptNow</b> | The total number of connections listening to incoming connections in use. |
| <b>nasEngTotalSent</b>      | The total number of packets sent.   |
| <b>nasEngTotalSentBytes</b> | The total number of bytes sent.   |
| <b>nasEngTotalRecv</b>      | The total number of packets received.                                     |
| <b>nasEngTotalRecvBytes</b> | The total number of bytes received.                                       |
| <b>nasEngBindTotal</b>      | The number of AppLogic bound since startup.                               |

| Measurement                         | Description  |
|-------------------------------------|--|
| <b>nasEngBindTotalCached</b>        | The number of AppLogic cached since startup.                       |
| <b>nasEngTotalThreads</b>           | Total number of threads created in this process.                   |
| <b>nasEngCurrentThreads</b>         | Total number of threads in use in this process.                    |
| <b>nasEngSleepingThreads</b>        | Number of threads sleeping in this process.                        |
| <b>nasEngDAETotalQuery</b>          | Total number of queries executed since startup.                    |
| <b>nasEngDAEQueryNow</b>            | The number of queries being processed.                             |
| <b>nasEngDAETotalConn</b>           | The number of logical connections created since startup.           |
| <b>nasEngDAEConnNow</b>             | The number of logical connections in use.                          |
| <b>nasEngDAECacheCount</b>          | The number of caches.  |
| <b>nasEngODBCQueryTotal</b>         | Total number of queries executed since startup.                    |
| <b>nasEngODBCPreparedQueryTotal</b> | Total number of odbc prepared queries executed since startup.      |
| <b>nasEngODBCConnTotal</b>          | Total number of connections opened since startup.                  |
| <b>nasEngODBCConnNow</b>            | Number of connections currently opened.                            |
| <b>nasEngORCLQueryTotal</b>         | Total number of queries executed since startup.                    |
| <b>nasEngORCLPreparedQueryTotal</b> | Total number of prepared queries executed since startup.           |
| <b>nasEngORCLConnTotal</b>          | Total number of connections established with Oracle since startup. |

| Measurement                        | Description   |
|------------------------------------|---|
| <b>nasEngORCLConnNow</b>           | Number of connections opened with Oracle now.               |
| <b>nasEngSYBQueryTotal</b>         | Total number of queries the driver processed since startup. |
| <b>nasEngSYBPreparedQueryTotal</b> | Total number of prepared queries processed since startup.   |
| <b>nasEngSYBConnTotal</b>          | Total number of connections opened since startup.           |
| <b>nasEngSYBConnNow</b>            | Number of SYB connections opened now.                       |
| <b>nasStatusTrapEntry</b>          | The KES definition.   |
| <b>nasTrapKesIpAddress</b>         | The IP Address of KES host.                                 |
| <b>nasTrapKesPort</b>              | The port of the main engine of this NAS.                    |
| <b>nasTrapEngPort</b>              | The port of the engine generating this event.               |
| <b>nasTrapEngState</b>             | The port of the engine generating this event.               |

- 6 To measure an object, select it, and click **Add**. Add all the desired resources to the list, and click **Close**.

---

**Note:** The iPlanet (NAS) monitor can only monitor up to 25 measurements.


---

- 7 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

---

**Note:** You can improve the level of measurement information for the iPlanet (NAS) monitor by enabling measurements with string values to be listed (in addition to measurements with numeric values), and by enabling the name modifier (which displays the string value as an identifying part of the measurement name).

In the following example of a measurement using the name modifier, the string value of ProcessName (sched) is displayed in addition to its instance ID (0):



To enable this feature, add the following line to the *<ProTune root folder>\dat\monitors\snmp.cfg* file:  
SNMP\_show\_string\_nodes=1

**Usage Notes:** You can select more than one name modifier, but the first in the hierarchy will be used. Each time the iPlanet (NAS) Add Measurements dialog box opens, the information is reread from the *snmp.cfg* file. You cannot add the same measurement twice (once with a name modifier and once without it). If you do so, an error message is issued.

---

## Configuring the Microsoft Active Server Pages Monitor

You select measurements to monitor the Microsoft ASP application server using the MS Active Server Pages dialog box.

---

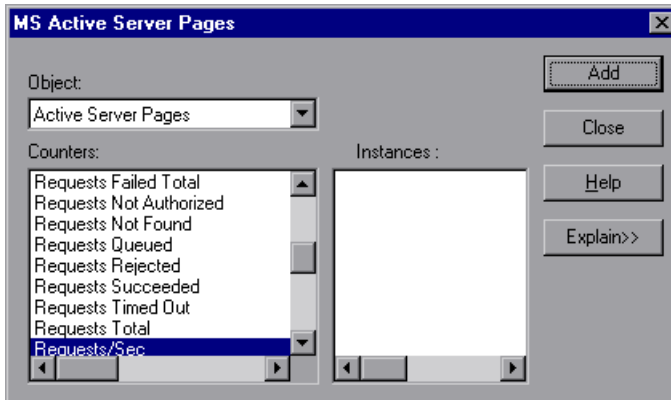
**Note:** To monitor an ASP server through a firewall, use TCP, port 139.

---

### To configure the ASP monitor:



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select MS Active Server Pages (in the Application Server category) and then click **Add**.
- 5 To select additional measurements, click **Add**. A dialog box displaying the Active Server Pages object, its counters, and instances opens.





The following table describes the default counters that can be monitored:

| Measurement                   | Description  |
|-------------------------------|--|
| Errors per Second             | The number of errors per second.   |
| Requests Wait Time            | The number of milliseconds the most recent request was waiting in the queue.                         |
| Requests Executing            | The number of requests currently executing.  |
| Requests Queued               | The number of requests waiting in the queue for service.   |
| Requests Rejected             | The total number of requests not executed because there were insufficient resources to process them. |
| Requests Not Found            | The number of requests for files that were not found.  |
| Requests/sec                  | The number of requests executed per second.  |
| Memory Allocated              | The total amount of memory, in bytes, currently allocated by Active Server Pages.                    |
| Errors During Script Run-Time | The number of failed requests due to run-time errors.  |
| Sessions Current              | The current number of sessions being serviced.   |
| Transactions/sec              | The number of transactions started per second.   |

---

**Note:** To change the default counters for the Microsoft ASP monitor, see “Changing a Monitor’s Default Counters” on page 661.

---

- 6** Select a counter and instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.
- 7** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.

- 8 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the Oracle9iAS HTTP Monitor

You select measurements to monitor the Oracle9iAS HTTP server using the Oracle HTTP Server Monitor Configuration dialog box. Note that you must start running the Oracle9iAS HTTP server before you begin selecting the measurements you want to monitor.

---

**Note:** The port you use to monitor an Oracle9iAS HTTP server through a firewall depends on the configuration of your server.

---

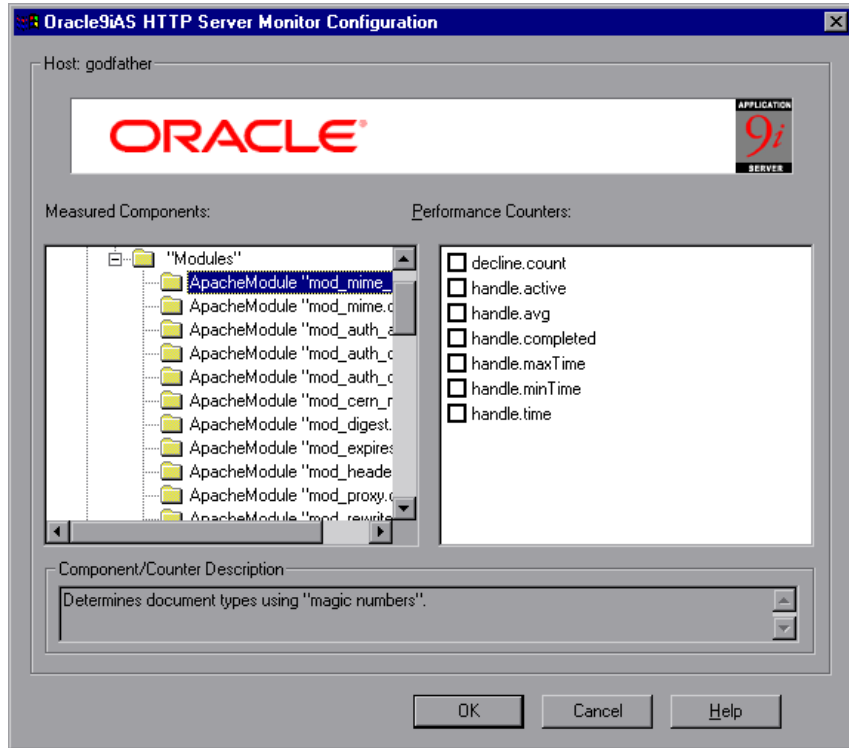
### To configure the Oracle9iAS HTTP monitor:



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select Oracle9iAS HTTP Server (in the Application Server category) and then click **Add**.

The Oracle9iAS HTTP Server Monitor Configuration dialog box opens, displaying the counters that can be monitored.

## 5 Browse the Measured Components tree.



## 6 Check the required machine processing counters, or application server performance counters and modules, in the Oracle HTTP Server Monitor Configuration window's right pane.

The following table describes some of the modules that can be monitored:

| Measurement                    | Description   |
|--------------------------------|---|
| <code>mod_mime.c</code>        | Determines document types using file extensions       |
| <code>mod_mime_magic.c</code>  | Determines document types using "magic numbers"       |
| <code>mod_auth_anon.c</code>   | Provides anonymous user access to authenticated areas |
| <code>mod_auth_dbm.c</code>    | Provides user authentication using DBM files          |
| <code>mod_auth_digest.c</code> | Provides MD5 authentication                           |

| Measurement                   | Description  |
|-------------------------------|--|
| <code>mod_cern_meta.c</code>  | Supports HTTP header metafiles   |
| <code>mod_digest.c</code>     | Provides MD5 authentication (deprecated by <code>mod_auth_digest</code> )      |
| <code>mod_expires.c</code>    | Applies Expires: headers to resources  |
| <code>mod_headers.c</code>    | Adds arbitrary HTTP headers to resources                                       |
| <code>mod_proxy.c</code>      | Provides caching proxy abilities   |
| <code>mod_rewrite.c</code>    | Provides powerful URI-to-filename mapping using regular expressions            |
| <code>mod_speling.c</code>    | Automatically corrects minor typos in URLs                                     |
| <code>mod_info.c</code>       | Provides server configuration information                                      |
| <code>mod_status.c</code>     | Displays server status   |
| <code>mod_usertrack.c</code>  | Provides user tracking using cookies   |
| <code>mod_dms.c</code>        | Provides access to DMS Apache statistics                                       |
| <code>mod_perl.c</code>       | Allows execution of perl scripts   |
| <code>mod_fastcgi.c</code>    | Supports CGI access to long-lived programs                                     |
| <code>mod_ssl.c</code>        | Provides SSL support   |
| <code>mod_plsql.c</code>      | Handles requests for Oracle stored procedures                                  |
| <code>mod_isapi.c</code>      | Provides Windows ISAPI extension support                                       |
| <code>mod_setenvif.c</code>   | Sets environment variables based on client information                         |
| <code>mod_actions.c</code>    | Executes CGI scripts based on media type or request method                     |
| <code>mod_imap.c</code>       | Handles imagemap files   |
| <code>mod_asis.c</code>       | Sends files that contain their own HTTP headers                                |
| <code>mod_log_config.c</code> | Provides user-configurable logging replacement for <code>mod_log_common</code> |
| <code>mod_env.c</code>        | Passes environments to CGI scripts   |

| Measurement              | Description  |
|--------------------------|--|
| <b>mod_alias.c</b>       | Maps different parts of the host file system in the document tree, and redirects URLs  |
| <b>mod_userdir.c</b>     | Handles user home directories  |
| <b>mod_cgi.c</b>         | Invokes CGI scripts  |
| <b>mod_dir.c</b>         | Handles the basic directory  |
| <b>mod_autoindex.c</b>   | Provides automatic directory listings  |
| <b>mod_include.c</b>     | Provides server-parsed documents   |
| <b>mod_negotiation.c</b> | Handles content negotiation  |
| <b>mod_auth.c</b>        | Provides user authentication using text files  |
| <b>mod_access.c</b>      | Provides access control based on the client hostname or IP address   |
| <b>mod_so.c</b>          | Supports loading modules (.so on UNIX, .dll on Win32) at run-time  |
| <b>mod_oprocmgr.c</b>    | Monitors JServ processes and restarts them if they fail  |
| <b>mod_jserv.c</b>       | Routes HTTP requests to JServ server processes. Balances load across multiple JServs by distributing new requests in round-robin order |
| <b>mod_ose.c</b>         | Routes requests to the JVM embedded in Oracle's database server  |
| <b>http_core.c</b>       | Handles requests for static Web pages  |

The following table describes the counters that can be monitored:

| Measurement           | Description  |
|-----------------------|--|
| <b>handle.minTime</b> | The minimum time spent in the module handler                   |
| <b>handle.avg</b>     | The average time spent in the module handler                   |
| <b>handle.active</b>  | The number of threads currently in the handle processing phase |

| Measurement                 | Description   |
|-----------------------------|---|
| <b>handle.time</b>          | The total amount of time spent in the module handler                                |
| <b>handle.completed</b>     | The number of times the handle processing phase was completed                       |
| <b>request.maxTime</b>      | The maximum amount of time required to service an HTTP request                      |
| <b>request.minTime</b>      | The minimum amount of time required to service an HTTP request                      |
| <b>request.avg</b>          | The average amount of time required to service an HTTP request                      |
| <b>request.active</b>       | The number of threads currently in the request processing phase                     |
| <b>request.time</b>         | The total amount of time required to service an HTTP request                        |
| <b>request.completed</b>    | The number of times the request processing phase was completed                      |
| <b>connection.maxTime</b>   | The maximum amount of time spent servicing any HTTP connection                      |
| <b>connection.minTime</b>   | The minimum amount of time spent servicing any HTTP connection                      |
| <b>connection.avg</b>       | The average amount of time spent servicing HTTP connections                         |
| <b>connection.active</b>    | The number of connections with currently open threads                               |
| <b>connection.time</b>      | The total amount of time spent servicing HTTP connections                           |
| <b>connection.completed</b> | The number of times the connection processing phase was completed                   |
| <b>numMods.value</b>        | The number of loaded modules  |
| <b>childFinish.count</b>    | The number of times the Apache parent server started a child server, for any reason |

| Measurement                   | Description   |
|-------------------------------|---|
| <b>childStart.count</b>       | The number of times “children” finished “gracefully.” There are some ungraceful error/crash cases that are not counted in childFinish.count |
| <b>Decline.count</b>          | The number of times each module declined HTTP requests  |
| <b>internalRedirect.count</b> | The number of times that any module passed control to another module using an “internal redirect”   |
| <b>cpuTime.value</b>          | The total CPU time utilized by all processes on the Apache server (measured in CPU milliseconds)  |
| <b>heapSize.value</b>         | The total heap memory utilized by all processes on the Apache server (measured in kilobytes)  |
| <b>pid.value</b>              | The process identifier of the parent Apache process   |
| <b>upTime.value</b>           | The amount of time the server been running (measured in milliseconds)   |

- 7 Click **OK** in the Oracle9iAS HTTP Server Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the Oracle9iAS HTTP monitor.

## Configuring the SilverStream Monitor

To monitor a SilverStream server you need to know the server statistics information URL. A simple way to verify the statistics URL is to access it from a browser.

The URL should be in the following format:

`http://<server_name/IP_address>:<port_number>/SilverStream/Statistics`

for example:

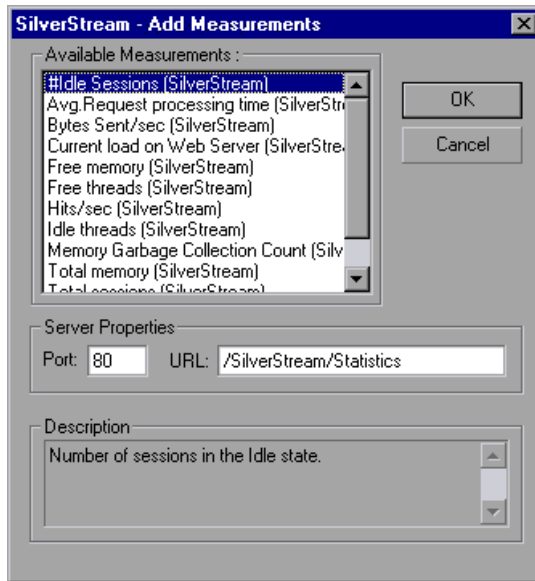
`http://199.203.78.57:80/SilverStream/Statistics`

**To configure the SilverStream monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select SilverStream (in the Application Server category) and then click **Add**.

A dialog box displaying the available measurements and server properties opens.



Select the required measurements. You can select multiple measurements using the **Ctrl** key.



The following table describes the measurements and server properties that can be monitored:

| Measurement                     | Description  |
|---------------------------------|--|
| #Idle Sessions                  | The number of sessions in the Idle state.  |
| Avg. Request processing time    | The average request processing time.   |
| Bytes Sent/sec                  | The rate at which data bytes are sent from the Web server.   |
| Current load on Web Server      | The percentage of load utilized by the SilverStream server, scaled at a factor of 25.                    |
| Hits/sec                        | The HTTP request rate.   |
| Total sessions                  | The total number of sessions.  |
| Free memory                     | The total amount of memory in the Java Virtual Machine currently available for future allocated objects. |
| Total memory                    | The total amount of memory in the Java Virtual Machine.  |
| Memory Garbage Collection Count | The total number of times the JAVA Garbage Collector has run since the server was started.               |
| Free threads                    | The current number of threads not associated with a client connection and available for immediate use.   |
| Idle threads                    | The number of threads associated with a client connection, but not currently handling a user request.    |
| Total threads                   | The total number of client threads allocated.  |

- 5** In the Server Properties section, enter the Port number and URL (without the server name), and click **OK**. The default URL is /SilverStream/Statistics.
- 6** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

**Note:** The default port number and URL can vary from one server to another. Please consult the Web server administrator.

---

**To change the default server properties:**

**1** Open the *SilverStream.cfg* file in the *<ProTune root folder>\dat\monitors\* directory.

**2** Edit the following parameters at the end of the file:

|                     |   |
|---------------------|---|
| <b>InfoURL</b>      | server statistics information URL   |
| <b>ServerPort</b>   | server port number  |
| <b>SamplingRate</b> | rate (milliseconds) at which the ProTune monitor will poll the server for the statistics information. If this value is greater than 1000, ProTune will use it as its sampling rate. Otherwise, it will use the sampling rate defined in the Monitors tab of the Options dialog box. |

---

**Note:** To monitor a SilverStream server through a firewall, use the Web server port (by default, port 80).

---

## Configuring the WebLogic (SNMP) Monitor

The WebLogic (SNMP) monitor uses SNMP to retrieve server statistics. To use this monitor, you must make sure that a version prior to WebLogic 6.0 is installed on your server, and that the SNMP agent is installed and activated on the server. For instructions on installing the SNMP agent, see <http://www.weblogic.com/docs51/admindocs/snmpagent.html>.

---

**Note:** To monitor a WebLogic (SNMP) server, use port 161 or 162, depending on the configuration of the agent.

---

### To configure the WebLogic (SNMP) monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select WebLogic (in the Application Server category) and then click **Add**.

The WebLogic SNMP Resources dialog box displays the available measurements.

---

**Note:** You need to define the port number if the WebLogic SNMP agent is running on a different port than the default SNMP port. You can define the default port for your WebLogic server in the configuration file, *snmp.cfg*, located in *<ProTune root folder>\dat\monitors*. For example, if the port used by the SNMP agent on your WebLogic server is 8888, you should edit the *snmp.cfg* file as follows:

```
; WebLogic
[cm_snmp_mon_isp]
port=8888
```

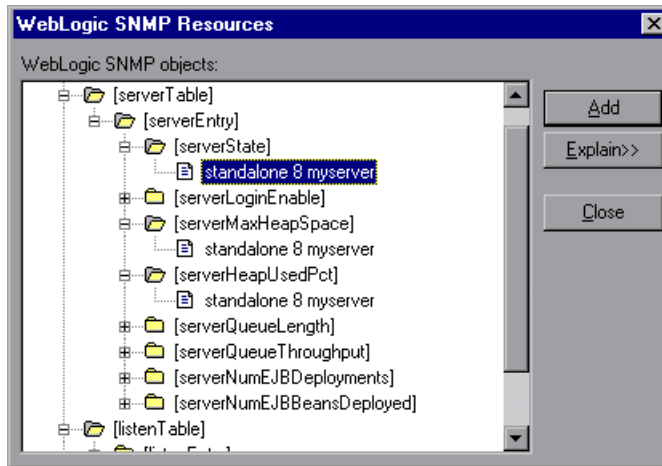
---

---

**Note:** The WebLogic (SNMP) monitor can only monitor up to 25 measurements.

---

**5** Browse the WebLogic SNMP Objects tree.



**6** To measure an object, select it, and click **Add**.

The following tables describe the measurements and server properties that can be monitored:

### Server Table

The Server Table lists all WebLogic (SNMP) servers that are being monitored by the agent. A server must be contacted or be reported as a member of a cluster at least once before it will appear in this table. Servers are only reported as a member of a cluster when they are actively participating in the cluster, or shortly thereafter.

| Measurement                      | Description   |
|----------------------------------|---|
| <b>ServerState</b>               | The state of the WebLogic server, as inferred by the SNMP agent. <i>Up</i> implies that the agent can contact the server. <i>Down</i> implies that the agent cannot contact the server. |
| <b>ServerLoginEnable</b>         | This value is true if client logins are enabled on the server.  |
| <b>ServerMaxHeapSpace</b>        | The maximum heap size for this server, in KB  |
| <b>ServerHeapUsedPct</b>         | The percentage of heap space currently in use on the server   |
| <b>ServerQueueLength</b>         | The current length of the server execute queue  |
| <b>ServerQueueThroughput</b>     | The current throughput of execute queue, expressed as the number of requests processed per second   |
| <b>ServerNumEJBDeployment</b>    | The total number of EJB deployment units known to the server  |
| <b>ServerNumEJBBeansDeployed</b> | The total number of EJB beans actively deployed on the server   |

### Listen Table

The Listen Table is the set of protocol, IP address, and port combinations on which servers are listening. There will be multiple entries for each server: one for each protocol, ipAddr, port combination. If clustering is used, the clustering-related MIB objects will assume a higher priority.

| Measurement          | Description   |
|----------------------|---|
| <b>ListenPort</b>    | Port number.  |
| <b>ListenAdminOK</b> | True if admin requests are allowed on this (protocol, ipAddr, port); otherwise false  |
| <b>ListenState</b>   | Listening if the (protocol, ipAddr, port) is enabled on the server; not Listening if it is not. The server may be listening but not accepting new clients if its server Login Enable state is false. In this case, existing clients will continue to function, but new ones will not. |

### ClassPath Table

The ClassPath Table is the table of classpath elements for Java, WebLogic (SNMP) server, and servlets. There are multiple entries in this table for each server. There may also be multiple entries for each path on a server. If clustering is used, the clustering-related MIB objects will assume a higher priority.

| Measurement    | Description  |
|----------------|--|
| <b>CPTYPE</b>  | The type of CP element: Java, WebLogic, servlet. A Java CPTYPE means the cpElement is one of the elements in the normal Java classpath. A WebLogic CPTYPE means the cpElement is one of the elements in weblogic.class.path. A servlet CPTYPE means the cpElement is one of the elements in the dynamic servlet classpath. |
| <b>CPIndex</b> | The position of an element within its path. The index starts at 1.   |

- 7 After selecting and adding the required objects, click **Close**.
- 8 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the WebLogic (JMX) Monitor

The BEA WebLogic (JMX) monitor uses the Java JMX interface to access runtime MBeans on the server. An MBean is a container that holds the performance data.

Before using the WebLogic (JMX) monitor, you must install Java 1.3 or later on the Console machine. If Java 1.3 or later is already installed, but is not the default Java version being used, specify the full path to the updated version. You specify the path in the *<ProTune root folder>\dat\monitors\WebLogicMon.ini* file. Edit the JVM entry in the [WebLogicMon] section. For example:

```
JVM="E:\Program Files\JavaSoft\JRE\1.3.1\bin\javaw.exe
```

---

**Note:** To use the WebLogic (JMX) monitor, you must make sure that WebLogic 6.0 or above is installed on your server.

---

### Setting Permissions for Monitoring

You must set certain permissions for a user to be able to monitor MBeans.

**To set permissions:**

- 1 Open the WebLogic console (<http://<host:port>/console>).
- 2 In the tree on the left, select **Security > ACLs**.  
If you are working with the WebLogic 6.1 console, click **Create a new ACL...** in the screen on the right.
- 3 In the New ACL Name box, type `weblogic.admin.mbean`, and click **Create**.

If you are working with the WebLogic 6.1 console, click **Add a new Permission...** in the screen on the right.

- 4** In the New Permission box (or Permission box, in the WebLogic 6.1 console), type `access`. In the WebLogic 6.0 console, click **Create**.
- 5** In the Users box and Groups box, enter the name of any user or group you want to use for monitoring.
- 6** Click **Grant Permission** in the WebLogic 6.0 console. In the WebLogic 6.1 console, click **Apply**.

### Loading Classes from the Server

The WebLogic (JMX) monitor utilizes a built-in server called the ClasspathServlet to load classes directly and automatically from the server. The advantages of this are easy installation and version independence. The disadvantages are a slight decrease in performance when loading classes for the first time (due to the size of the servlet), and the possibility of the servlet becoming disabled.

If the servlet is disabled, or if you do not want to use the servlet, you can load classes directly from the file system.

#### To load classes directly from the file system:

- 1** Copy the `weblogic.jar` file from the application server install folder (under the lib folder) to `<ProTune root folder>\classes`.
- 2** If the classes file is not located in the default `<ProTune root folder>` folder, you need to specify the full path to it in the `<ProTune root folder>\dat\monitors\WebLogicMon.ini` file. In this file, change the line `Weblogic=weblogic.jar` to `Weblogic=<full path to weblogic.jar>`.



## Configuring the WebLogic (JMX) Monitor

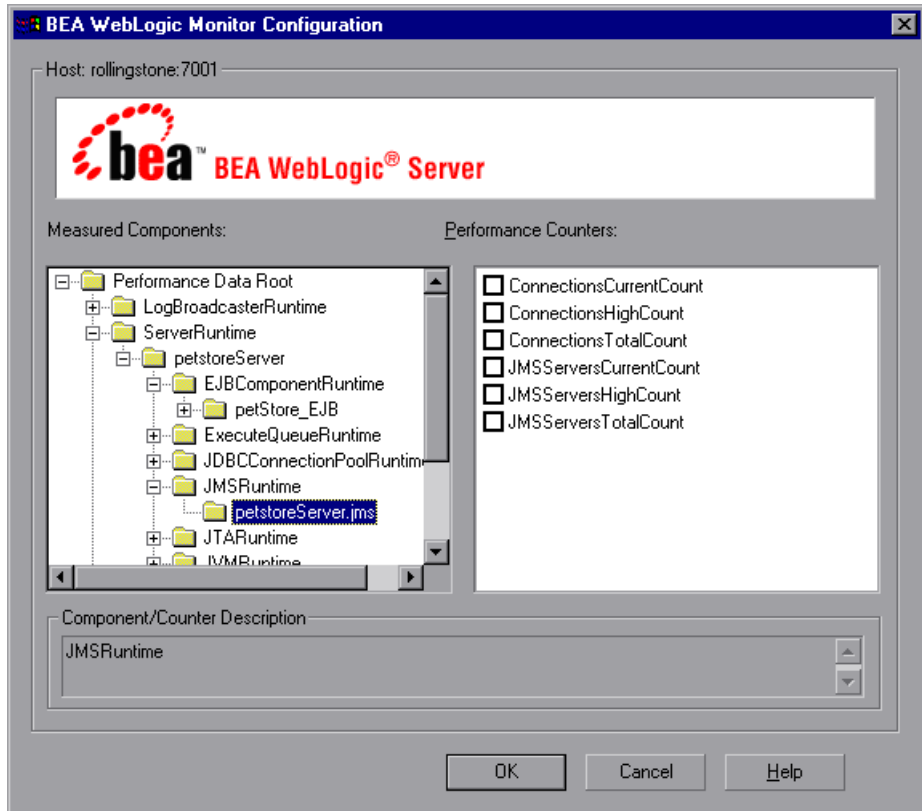
You select measurements to monitor the WebLogic (JMX) application server using the BEA WebLogic Monitor Configuration dialog box.

### To configure the WebLogic (JMX) Monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select BEA WebLogic (JMX) (in the Application Server category) and then click **Add**.
- 5** In the Enter Login Information dialog box, enter the username and password of a user with administrative privileges to the WebLogic server. The BEA WebLogic Monitor Configuration dialog box opens. For details on creating user permissions, see “Setting Permissions for Monitoring” on page 357.

6 Browse the Measured Components tree.



7 Check the required performance counters in the BEA WebLogic Monitor Configuration window's right pane.

8 Click **OK** in the BEA WebLogic Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the WebLogic (JMX) monitor.

The following measurements are available for the WebLogic (JMX) server:

### LogBroadcasterRuntime

| Measurement     | Description   |
|-----------------|---|
| MessagesLogged  | The number of total log messages generated by this instance of the WebLogic server. |
| Registered      | Returns “false” if the MBean represented by this object has been unregistered.      |
| CachingDisabled | Private property that disables caching in proxies.                                  |

### ServerRuntime

For more information on the measurements contained in each of the following measurement categories, see Mercury Interactive’s Load Testing Monitors Web site ([http://www-heva.mercuryinteractive.com/products/loadrunner/load\\_testing\\_monitors/bealogic.html](http://www-heva.mercuryinteractive.com/products/loadrunner/load_testing_monitors/bealogic.html)).

- ServletRuntime
- WebAppComponentRuntime
- EJBStatefulHomeRuntime
- JTARuntime
- JVMRuntime
- EJBEntityHomeRuntime.
- DomainRuntime
- EJBComponentRuntime
- DomainLogHandlerRuntime
- JDBCConnectionPoolRuntime
- ExecuteQueueRuntime
- ClusterRuntime
- JMSRuntime

- TimeServiceRuntime
- EJBStatelessHomeRuntime
- WLECConnectionServiceRuntime

### ServerSecurityRuntime

| Measurement                               | Description   |
|---|---|
| <b>UnlockedUsersTotalCount</b>            | Returns the number of times a user has been unlocked on the server                                |
| <b>InvalidLoginUsersHighCount</b>         | Returns the high-water number of users with outstanding invalid login attempts for the server     |
| <b>LoginAttemptsWhileLockedTotalCount</b> | Returns the cumulative number of invalid logins attempted on the server while the user was locked |
| <b>Registered</b>                         | Returns "false" if the MBean represented by this object has been unregistered.                    |
| <b>LockedUsersCurrentCount</b>            | Returns the number of currently locked users on the server  |
| <b>CachingDisabled</b>                    | Private property that disables caching in proxies.  |
| <b>InvalidLoginAttemptsTotalCount</b>     | Returns the cumulative number of invalid logins attempted on the server                           |
| <b>UserLockoutTotalCount</b>              | Returns the cumulative number of user lockouts done on the server                                 |

## Configuring the WebSphere Monitor

The WebSphere 3.x, 4.x, and 5.x application servers have different monitor installation requirements.

To monitor versions 3.02, 3.5, and 3.5.x of the IBM WebSphere application server, you must first install the appropriate IBM WebSphere servlet patch on the WebSphere machine.

WebSphere versions 4.x and 5.x contain the performance servlet within the default installation. To monitor WebSphere version 5.x, you need to deploy the performance servlet on the application server using the IBM WebSphere "Installing a New Application" wizard.

**To install the IBM WebSphere servlet patch for the WebSphere 3.x server:**

- 1 Open Mercury Interactive's Customer Support site, and select **Downloads > Patches** from the tree on the left.

Please make sure to install the appropriate patch according to your WebSphere version:

**WebSphere version 3.02**      IBM\_WebSphere3.02\_Servlet.zip

**WebSphere version 3.5**      IBM\_WebSphere3.5\_Servlet.zip

**WebSphere version 3.5.x**      IBM\_WebSphere3.5.x\_Servlet.zip

- 2 Unzip the *IBM\_WebSphere<version#>\_Servlet.zip* file in the ProTune Performance Monitors section.
- 3 Copy *xml4j.jar*, *performance.dtd* and *perf.jar* (version 3.02), or *perf35.jar* (version 3.5) or *perf35x.jar* (version 3.5.2 and 3.5.3) into the *default\_host\default\_app servlets* directory on the monitored machine.

To find the *default\_app servlets* directory of the Web application, check the Web application's classpath. From the admin console, select the Web application in the tree and click the Advanced tab. You should see the classpath for the Web application.

For example:

- ▶ **Microsoft Windows Platforms:** if the IBM WebSphere directory is installed under drive E, then the files should be copied to:  
*E:\WebSphere\AppServer\hosts\default\_host\default\_app\servlets*
- ▶ **IBM iSeries Platforms:** files should be copied to:  
*/QIBM/UserData/WebASAdv/<instance>/hosts/default\_host/default\_app/servlets*
- ▶ **UNIX/Linux Platforms:** files should be copied to:  
*/opt/IBMWebAS/hosts/default\_host/default\_app/servlets*

---

**Note:** If you want to monitor additional Web applications that are not on the same machine, copy the above files to the Servlets folder of the application you want to monitor.

Add the `com.ibm.ivb.epm.servlet.PerformanceServlet` to the classpath configuration in the WebSphere console for each Web application.

---

- 4** After copying the files, verify that the servlet is running properly and that the performance data is being generated. A simple way to verify that the performance data is accessible is to display it in a Web browser. The URL must be in the following format:

`http://<server name:port number>/<servlet_folder>/com.ibm.ivb.epm.servlet.PerformanceServlet`

For example: `http://websphere.mercury.co.il:81/servlet/com.ibm.ivb.epm.servlet.PerformanceServlet`

---

**Note:** Only browsers that are XML-compatible will allow you to view the performance XML file.

---

- 5** Open the `<ProTune root folder>\dat\monitors\xmlmonitorshared.ini` file and add the following line to the [WebSphereMonitor] section:

`QueryLoginInfo=1`

**To set up the environment to monitor the WebSphere 4.0 server:**

- 1** Open the `<ProTune root folder>\dat\monitors\xmlmonitorshared.ini` file.
- 2** Add the following lines to the [WebSphereMonitor] section:
 

```
ServletName=com.ibm.ws.pmi.perfServlet.PerformanceServlet
ServletAlias=wasPerfTool/servlet QueryLoginInfo=1
```

---

**Note:** After making this change to the `xmlmonitorshared.ini` file, the Console will only be able to monitor WebSphere version 4.0. Previous versions of WebSphere will not be supported.

---

**To deploy the performance servlet on the application server for WebSphere 5.x:**

- 1** From the administrative console, click **Applications > Install New Application** in the console navigation tree.
- 2** For Path, specify the full path name of the source application file ("PerfServletApp.ear") on the server machine and click **Next**.
- 3** Select the **Generate Default Bindings** check box and click **Next**.
- 4** On the Install New Application page, click **Summary**, and select the **Cell/Node/Server** option. Click **Click here**.
- 5** On the **Map modules to application servers** panel, select the server onto which the application files will install from the **Clusters and Servers** list, and select **Module** to select all of the application modules.
- 6** Click **Next**, and in the Summary panel click **Finish**.
- 7** Verify that the servlet is running properly and that the performance data is being generated. A simple way to verify that the performance data is accessible is to display it in a Web browser. The URL must be in the following format:

```
http://<server name:port number>/<servlet_folder>/com.ibm.ivb.epm.servlet.
PerformanceServlet
```

For example: `http://websphere.mercury.co.il:81/servlet/com.ibm.ivb.epm.servlet.PerformanceServlet`

**Note:** Only browsers that are XML-compatible will allow you to view the performance XML file.

---

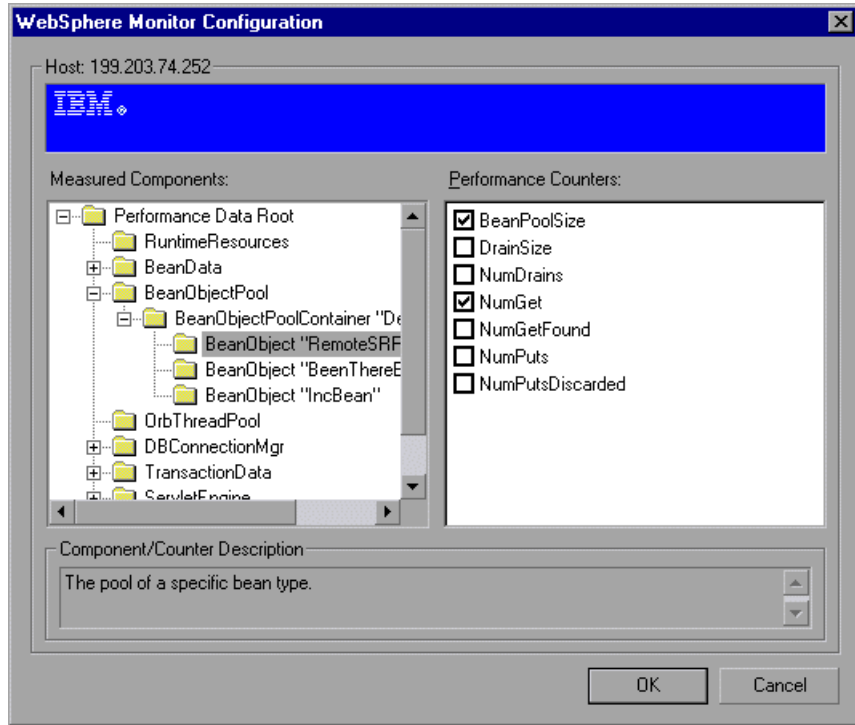
**To configure the WebSphere 3.x, and WebSphere 4.x - 5.x monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select WebSphere or WebSphere 4.x - 5.x (in the Application Server category) and then click **Add**. The WebSphere Monitor Configuration dialog box displays the available measurements.



## 5 Browse the Measured Components tree.



- 6 Check the required performance counters in the WebSphere Monitor Configuration window's right pane. For a list of the available performance counters, see page 370.
- 7 Click **OK** in the WebSphere Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the WebSphere monitor.

---

**Note:** The port you use to monitor a WebSphere server through a firewall depends on the configuration of your server.

---

**To specify another Web alias for the servlet directory:**

By default, ProTune uses the alias `servlet` as the servlet directory Web alias. For example, if the WebSphere Server machine is named `mercury` and the path for the servlets directory is: `E:\AppServer\hosts\default_host\default_app\servlets`, ProTune will request the XML file in the following URL:  
`http://mercury/servlet/com.ibm.ivb.epm.servlet.PerformanceServlet`, where `servlet` is the default web alias for the servlet directory.

If the Web alias for the servlet directory is not `servlet`, you must specify the servlet directory Web alias in the Add Machine dialog box according to the following format:

`http://<server name:port number>/<servlet_dir_alias>`

For example: `http://mercury/servlet2`

Using this method, you can monitor as many application servers as you want—whether they are installed on the same machine, or on different machines.

**To monitor other applications, in addition to the default application:**

You can monitor as many applications as you want, regardless of whether they are installed on the same machine or different machines.

- 1** Copy the same files that you copied to the Servlets directory for the Default application to the Servlets directories for any other Web applications that you want to monitor.
- 2** Add the `com.ibm.ivb.epm.servlet.PerformanceServlet` to the configuration in the WebSphere Console for each Web application.
- 3** Add the Web application to be monitored to the WebSphere Performance Monitor using the following format:

`http://<server:port_number>/<servlet_dir_alias>/servlet`

For example: `http://mercury/servlet3/servlet`

**To work with WebSphere version 3.5.x**

- 1** The EPM counters in 3.5.x are by default set to "none". To enable the counters, choose the application server you are monitoring in the WebSphere Administrator's Console browser.
- 2** Right-click the application server and select **Performance**. Select Performance Modules from the pop-up window.
- 3** Right-click Performance Modules to choose a performance level. Selecting various levels of counters enables the application server to manage varying levels of performance data.
- 4** Click the **Set** button.
- 5** In versions 3.5.2 and 3.5.3 the Servlet counters have been disabled. To enable the Servlet counters, you need to modify the contents of the `com/ibm/servlet/appserver.properties` file located in "`<WAS_HOME>\lib\ibmwebas.jar`".

Extract the *jar* file and modify the `appserver.properties` as follows:

```
#listeners.application=com.ibm.servlet.engine.EPMApplicationListener
com.ibm.servlet.debug.OLTServletManager
listeners.application=
```

Should be:

```
listeners.application=com.ibm.servlet.engine.EPMApplicationListener
com.ibm.servlet.debug.OLTServletManager
#listeners.application=
```

- 6** Repackage the *jar* file.

## WebSphere Counters

The following tables describe the counters that can be monitored:

### Run-Time Resources

Contains resources related to the Java Virtual Machine run-time, as well as the ORB.

| Measurement | Description   |
|-------------|---|
| MemoryFree  | The amount of free memory remaining in the Java Virtual Machine |
| MemoryTotal | The total memory allocated for the Java Virtual Machine         |
| MemoryUse   | The total memory in use within the Java Virtual Machine         |

### BeanData

Every home on the server provides performance data, depending upon the type of bean deployed in the home. The top level bean data holds an aggregate of all the containers.

| Measurement         | Description  |
|---------------------|--|
| BeanCreates         | The number of beans created. Applies to an individual bean that is either 'stateful' or 'entity'   |
| EntityBeanCreates   | The number of entity beans created   |
| BeanRemoves         | The number of entity beans pertaining to a specific bean that have been removed. Applies to an individual bean that is either 'stateful' or 'entity' |
| EntityBeanRemoves   | The number of entity beans removed   |
| StatefulBeanCreates | The number of stateful beans created   |
| StatefulBeanRemoves | The number of stateful bean removed  |

| Measurement                      | Description  |
|----------------------------------|--|
| <b>BeanPassivates</b>            | The number of bean passivates pertaining to a specific bean. Applies to an individual bean that is either 'stateful' or 'entity' |
| <b>EntityBeanPassivates</b>      | The number of entity bean passivates   |
| <b>StatefulBeanPassivates</b>    | The number of stateful bean passivates   |
| <b>BeanActivates</b>             | The number of bean activates pertaining to a specific bean. Applies to an individual bean that is either 'stateful' or 'entity'  |
| <b>EntityBeanActivates</b>       | The number of entity bean activates  |
| <b>StatefulBeanActivates</b>     | The number of stateful bean activates  |
| <b>BeanLoads</b>                 | The number of times the bean data was loaded. Applies to entity  |
| <b>BeanStores</b>                | The number of times the bean data was stored in the database. Applies to entity  |
| <b>BeanInstantiates</b>          | The number of times a bean object was created. This applies to an individual bean, regardless of its type.                       |
| <b>StatelessBeanInstantiates</b> | The number of times a stateless session bean object was created  |
| <b>StatefulBeanInstantiates</b>  | The number of times a stateful session bean object was created   |
| <b>EntityBeanInstantiates</b>    | The number of times an entity bean object was created  |
| <b>BeanDestroys</b>              | The number of times an individual bean object was destroyed. This applies to any bean, regardless of its type                    |
| <b>StatelessBeanDestroys</b>     | The number of times a stateless session bean object was destroyed  |
| <b>StatefulBeanDestroys</b>      | The number of times a stateful session bean object was destroyed   |

| Measurement                | Description   |
|----------------------------|---|
| <b>EntityBeanDestroys</b>  | The number of times an entity bean object was destroyed   |
| <b>BeansActive</b>         | The average number of instances of active beans pertaining to a specific bean. Applies to an individual bean that is either 'stateful' or 'entity'                |
| <b>EntityBeansActive</b>   | The average number of active entity beans   |
| <b>StatefulBeansActive</b> | The average number of active session beans  |
| <b>BeansLive</b>           | The average number of bean objects of this specific type that are instantiated but not yet destroyed. This applies to an individual bean, regardless of its type. |
| <b>StatelessBeansLive</b>  | The average number of stateless session bean objects that are instantiated but not yet destroyed  |
| <b>StatefulBeansLive</b>   | The average number of stateful session bean objects that are instantiated but not yet destroyed   |
| <b>EntityBeansLive</b>     | The average number of entity bean objects that are instantiated but not yet destroyed   |
| <b>BeanMethodRT</b>        | The average method response time for all methods defined in the remote interface to this bean. Applies to all beans   |
| <b>BeanMethodActive</b>    | The average number of methods being processed concurrently. Applies to all beans  |
| <b>BeanMethodCalls</b>     | The total number of method calls against this bean  |

## BeanObjectPool

The server holds a cache of bean objects. Each home has a cache and there is therefore one BeanObjectPoolContainer per container. The top level BeanObjectPool holds an aggregate of all the containers data.

| Measurement                    | Description   |
|--------------------------------|---|
| <b>BeanObjectPoolContainer</b> | The pool of a specific bean type  |
| <b>BeanObject</b>              | The pool specific to a home   |
| <b>NumGet</b>                  | The number of calls retrieving an object from the pool  |
| <b>NumGetFound</b>             | The number of calls to the pool that resulted in finding an available bean                                      |
| <b>NumPuts</b>                 | The number of beans that were released to the pool  |
| <b>NumPutsDiscarded</b>        | The number of times releasing a bean to the pool resulted in the bean being discarded because the pool was full |
| <b>NumDrains</b>               | The number of times the daemon found the pool was idle and attempted to clean it                                |
| <b>DrainSize</b>               | The average number of beans discarded by the daemon during a clean  |
| <b>BeanPoolSize</b>            | The average number of beans in the pool   |

## OrbThreadPool

These are resources related to the ORB thread pool that is on the server.

| Measurement             | Description   |
|-------------------------|---|
| <b>ActiveThreads</b>    | The average number of active threads in the pool  |
| <b>TotalThreads</b>     | The average number of threads in the pool   |
| <b>PercentTimeMaxed</b> | The average percent of the time that the number of threads in the pool reached or exceeded the desired maximum number |

| Measurement       | Description                                     |
|-------------------|---|
| ThreadCreates     | The number of threads created                   |
| ThreadDestroys    | The number of threads destroyed                 |
| ConfiguredMaxSize | The configured maximum number of pooled threads |

### DBConnectionMgr

These are resources related to the database connection manager. The manager consists of a series of data sources, as well as a top-level aggregate of each of the performance metrics.

| Measurement            | Description   |
|------------------------|---|
| DataSource             | Resources related to a specific data source specified by the "name" attribute |
| ConnectionCreates      | The number of connections created   |
| ConnectionDestroys     | The number of connections released  |
| ConnectionPoolSize     | The average size of the pool, i.e., number of connections                     |
| ConnectionAllocates    | The number of times a connection was allocated                                |
| ConnectionWaiters      | The average number of threads waiting for a connection                        |
| ConnectionWaitTime     | The average time, in seconds, of a connection grant                           |
| ConnectionTime         | The average time, in seconds, that a connection is in use                     |
| ConnectionPercentUsed  | The average percentage of the pool that is in use                             |
| ConnectionPercentMaxed | The percentage of the time that all connections are in use                    |



## TransactionData

These are resources that pertain to transactions.

| Measurement                   | Description  |
|-------------------------------|--|
| <b>NumTransactions</b>        | The number of transactions processed                                 |
| <b>ActiveTransactions</b>     | The average number of active transactions                            |
| <b>TransactionRT</b>          | The average duration of each transaction                             |
| <b>BeanObjectCount</b>        | The average number of bean object pools involved in a transaction    |
| <b>RolledBack</b>             | The number of transactions rolled back                               |
| <b>Committed</b>              | The number of transactions committed                                 |
| <b>LocalTransactions</b>      | The number of transactions that were local                           |
| <b>TransactionMethodCount</b> | The average number of methods invoked as part of each transaction    |
| <b>Timeouts</b>               | The number of transactions that timed out due to inactivity timeouts |
| <b>TransactionSuspended</b>   | The average number of times that a transaction was suspended         |

## ServletEngine

These are resources that are related to servlets and JSPs.

| Measurement            | Description   |
|------------------------|---|
| <b>ServletsLoaded</b>  | The number of servlets currently loaded                     |
| <b>ServletRequests</b> | The number of requests serviced                             |
| <b>CurrentRequests</b> | The number of requests currently being serviced             |
| <b>ServletRT</b>       | The average response time for each request                  |
| <b>ServletsActive</b>  | The average number of servlets actively processing requests |

| Measurement               | Description  |
|---------------------------|--|
| <b>ServletIdle</b>        | The amount of time that the server has been idle (i.e., time since last request) |
| <b>ServletErrors</b>      | The number of requests that resulted in an error or an exception                 |
| <b>ServletBeanCalls</b>   | The number of bean method invocations that were made by the servlet              |
| <b>ServletBeanCreates</b> | The number of bean references that were made by the servlet                      |
| <b>ServletDBCalls</b>     | The number of database calls made by the servlet                                 |
| <b>ServletDBConAlloc</b>  | The number of database connections allocated by the servlet                      |
| <b>SessionLoads</b>       | The number of times the servlet session data was read from the database          |
| <b>SessionStores</b>      | The number of times the servlet session data was stored in the database          |
| <b>SessionSize</b>        | The average size, in bytes, of a session data                                    |
| <b>LoadedSince</b>        | The time that has passed since the server was loaded (UNC time)                  |

### Sessions

These are general metrics regarding the HTTP session pool.

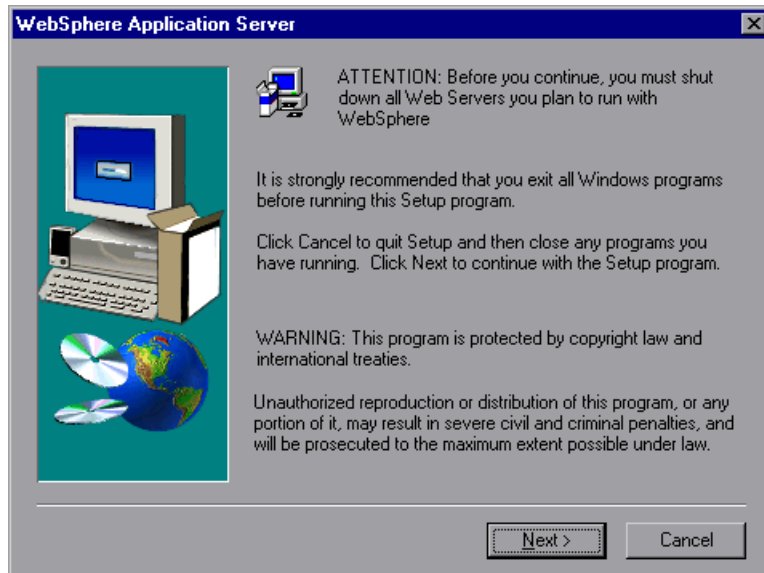
| Measurement                | Description  |
|----------------------------|--|
| <b>SessionsCreated</b>     | The number of sessions created on the server   |
| <b>SessionsActive</b>      | The number of currently active sessions  |
| <b>SessionsInvalidated</b> | The number of invalidated sessions. May not be valid when using sessions in the database mode                    |
| <b>SessionLifetime</b>     | Contains statistical data of sessions that have been invalidated. Does not include sessions that are still alive |

## Configuring the WebSphere (EPM) Monitor

To monitor the IBM WebSphere application server (3.5.x), you must first install the IBM WebSphere Administrator's Console on the Console machine. You may also need to copy the security keyring.

### To install the IBM WebSphere Administrator's Console:

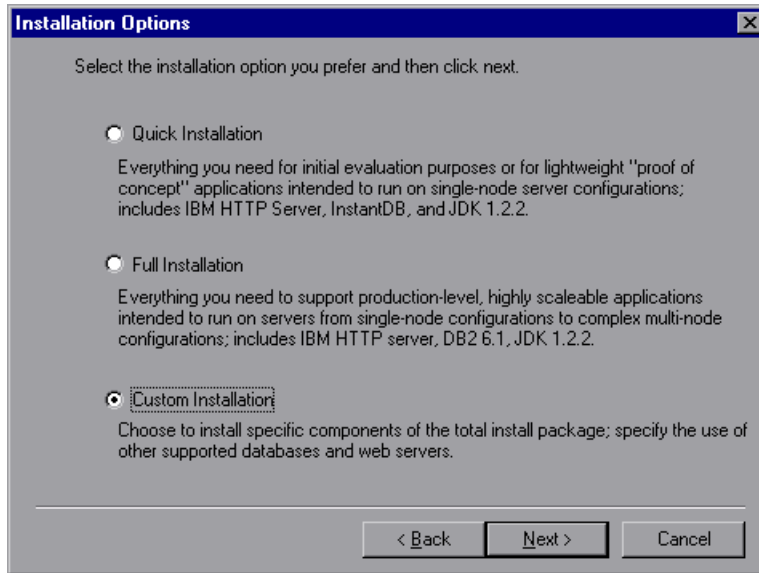
- 1 Start the WebSphere installation program from the WebSphere 3.5 Windows NT distribution CD-ROM. The WebSphere Application Server dialog box opens.



- 2 Disregard the instruction to shut down all Web servers that you plan to run with WebSphere. This is not relevant to the Administrator's Console installation.

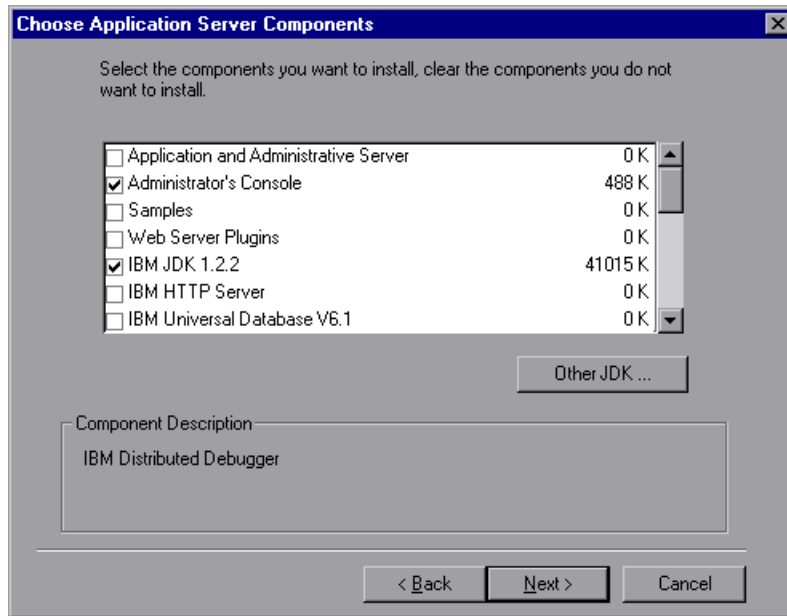
Click **Next**.

**3** The Installation Options dialog box opens. Select **Custom Installation**.



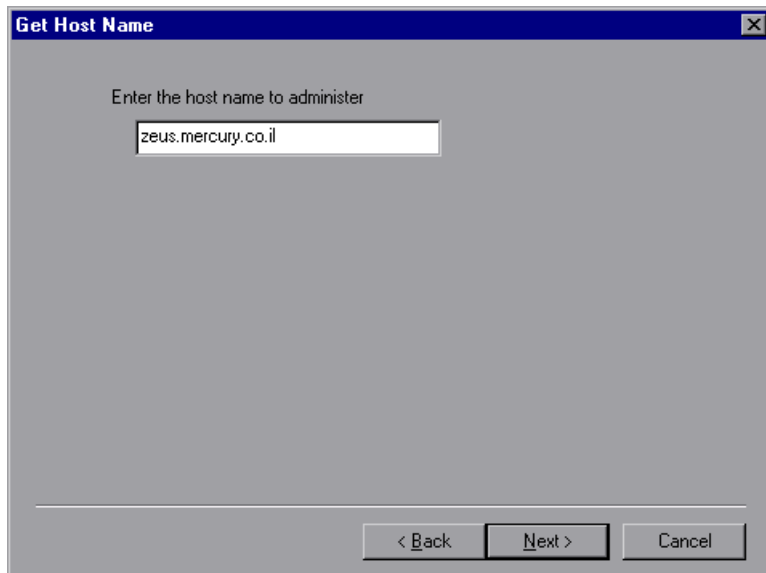
Click **Next**.

- 4 The Choose Application Server Components dialog box opens. Select **Administrator's Console** and **IBM JDK 1.2.2**. Clear all the other options.



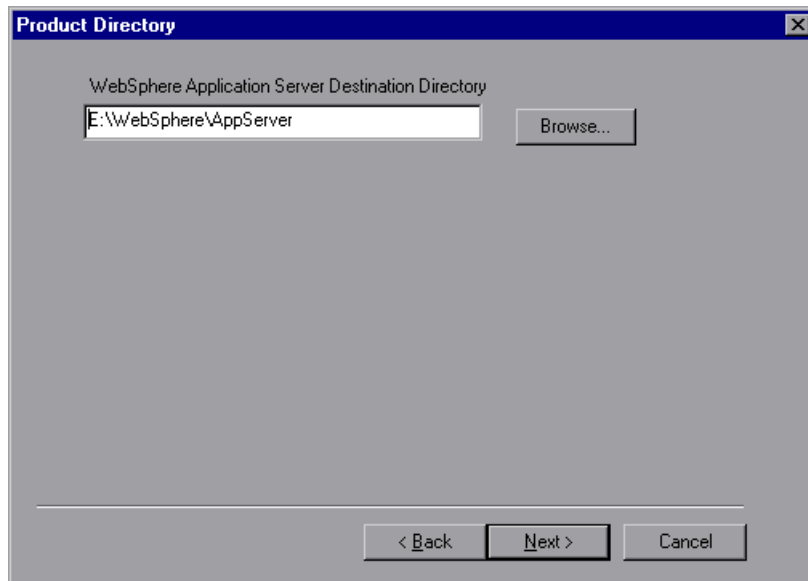
Click **Next**.

- 5 The Get Host Name dialog box opens. Type the name of the machine that you want to monitor.




Click **Next**.

- 6 The Product Directory dialog box opens. Specify the folder in which to install the Administrator's Console. To select a different location, click **Browse**, choose a folder other than the default folder, and click **OK**.




Click **Next**.

- 7 The Select Program Folder dialog box opens. Specify a program folder, or accept the default folder, IBM WebSphere\Application Server V3.5.



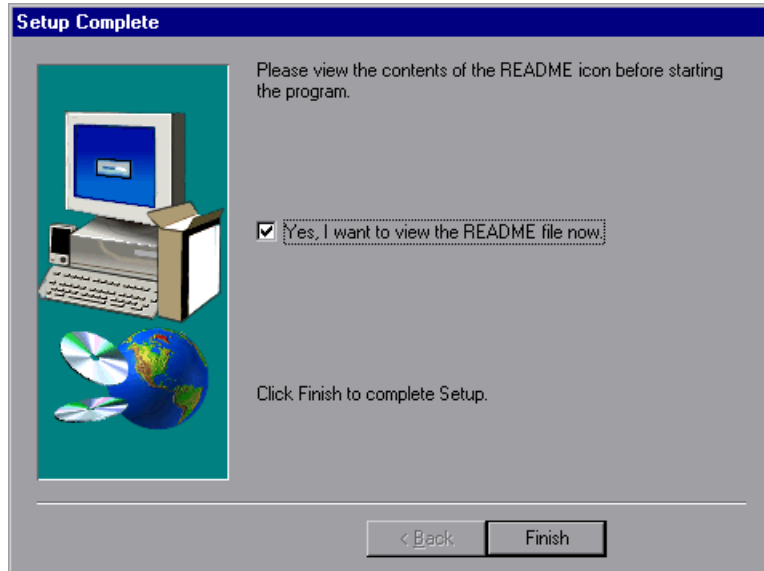
Click **Next**.

The installation process begins. To pause or quit the installation, click **Cancel**.






- 8 When the installation is complete, the Setup Complete dialog box opens. Select the check box to view the readme file before starting the program. You can view the readme file at any time by selecting **Start > Programs > Application Server V3.5 > IBM WebSphere > README**.



- 9 Click **Finish** to complete the installation program. The Restarting Windows dialog box opens.



- 10 Select either to restart your computer and complete the installation now (recommended) or to wait and complete the installation later.
- 11 Click **OK** to complete the installation of the Administrator's Console.

### **Copying the Security Keyring**

If you enabled security on the WebSphere server, you must copy the security keyring from the server to the admin client. (One way to tell whether security is enabled is to see whether the Administrator's Console can connect to the admin server.) A keyring is a certification used by the server to identify the client.

You need to copy the *jar* file containing the keyring from the server lib folder to the client lib folder. You also need to add the *jar* file containing the keyring to the monitoring client command line.

---

**Note:** The keyring used in this file (*353Keyring.jar*) is the IBM dummy keyring that must be installed on servers using versions 3.52 and below. If your server is using the IBM dummy keyring and is version 3.52 or below, you do not need to change the line. If you are using the dummy keyring and are running version 3.53 or later, you do not need to do anything.

---

#### **To copy the keyring:**

- 1 Copy the keyring *jar* file from the server to the admin client lib folder (by default, C:\Websphere\Appserver\lib):

The *jar* file containing the keyring, *xxxKeyring.jar*, is located by default in the following location:

|             |                             |
|-------------|-----------------------------|
| NT Server   | C:\Websphere\Appserver\lib  |
| UNIX Server | OPT/websphere/Appserver/lib |

- 2 Open the <ProTune root folder>\dat\monitors\WebSphere35Mon.ini file in a text editor.
- 3 Locate the following line:  
JVM\_CLASSES4=C:\WebSphere\AppServer\lib\353Keyring.jar

---

**Note:** If you did not use the default location for the WebSphere installation, the line will be different.

---

- 4 Change *353Keyring.jar* to the keyring you are using.

### Enabling EPM Counters on the WebSphere 3.5.x Server

To enable the EPM counters, which are by default set to "none," right-click the application you are monitoring in the WebSphere Administrator's Console browser, and select **Performance**. Expand the Performance Modules tree in the dialog box that opens. In order to manage different levels of performance data, right-click the performance modules and choose a performance level. Click the **Set** button.

Alternatively, ensure that the application server is started, select the **Advanced** tab in the WebSphere Administrator's Console browser, and in the EPM Specification box, type:  
epm=high:epm.beanMethodData=none

### Activating the WebSphere (EPM) Monitor

Once you have installed the WebSphere Administrator's Console and enabled the EPM counters, you can activate the WebSphere (EPM) monitor.

**To activate the WebSphere EPM monitor:**

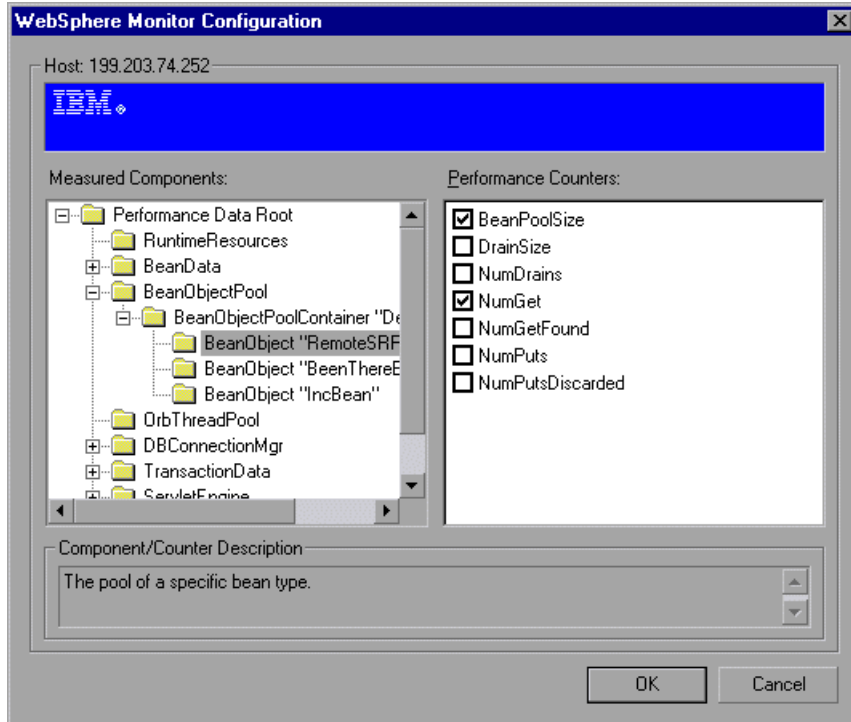


- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.

- 4 In the left section of the dialog box, select WebSphere (EPM) (in the Application Server category) and then click **Add**.

The WebSphere Monitor Configuration dialog box displays the available measurements.

- 5 Browse the Measured Components tree.



- 6 Check the required performance counters in the WebSphere Monitor Configuration window's right pane. For a list of the available performance counters, see page 370.
- 7 Click **OK** in the WebSphere Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the WebSphere (EPM) monitor.

# 23

---

## Database Resource Monitoring

You can monitor DB2, Oracle, SQL Server, or Sybase database resource usage during a session step run using ProTune's Database Server Resource monitors.

This chapter describes:

- Configuring the DB2 Monitor
- Configuring the Oracle Monitor
- Configuring the SQL Server Monitor
- Configuring the Sybase Monitor

### About Database Resource Monitoring

The DB2, Oracle, SQL Server, or Sybase database server resource monitors measure statistics for DB2, Oracle, SQL Server, or Sybase database servers. During a session step run, you use these monitors to isolate database server performance bottlenecks.

For each database server, you configure the measurements you want to monitor before running your session step. Note that in order to run the DB2, Oracle, and Sybase monitors, you must also install the client libraries on the database server you want to monitor.

## Configuring the DB2 Monitor

The DB2 database server monitor measures the resource usage on a DB2 database during a session step run.

---

**Note:** If there is no application working with a database, you can only monitor the database manager instance.

---

Before you can monitor a DB2 database server, you must set up the DB2 monitor environment.

### To set up the DB2 monitor environment:

- 1 Install all the client files and libraries on the Console machine.
- 2 Select **Start > Programs > DB2 for Windows NT > Control Center**. Enter your DB2 server username and password (with administrative privileges).
- 3 In the console that opens, right-click **Systems**, and select **Add**.
- 4 Enter the following settings in the dialog box:
  - System Name:** *<server name>*
  - Remote Instance:** DB2
  - Host Name:** *<server name>*
  - Service Name:** the DB2 server port. The default value is 50000.
- 5 Click **Retrieve**, and then **OK**.

---

**Note:** If you receive an error message after clicking **Retrieve**, repeat steps 3 and 4, and click **OK**.

---

- 6 Expand the *<server name>* node in the console tree.
- 7 Right-click **Instance**, and select **Add**.

**8** Enter the following settings in the dialog box:

**Remote Instance:** DB2

**Instance Name:** the database instance to be called from the Console

**Host Name:** <server name>

**Service Name:** the DB2 server port. The default value is 50000.

**9** Click **OK** and close the Control Center.

---

**Note:** You can only work with a single Database Manager instance during each monitoring session.

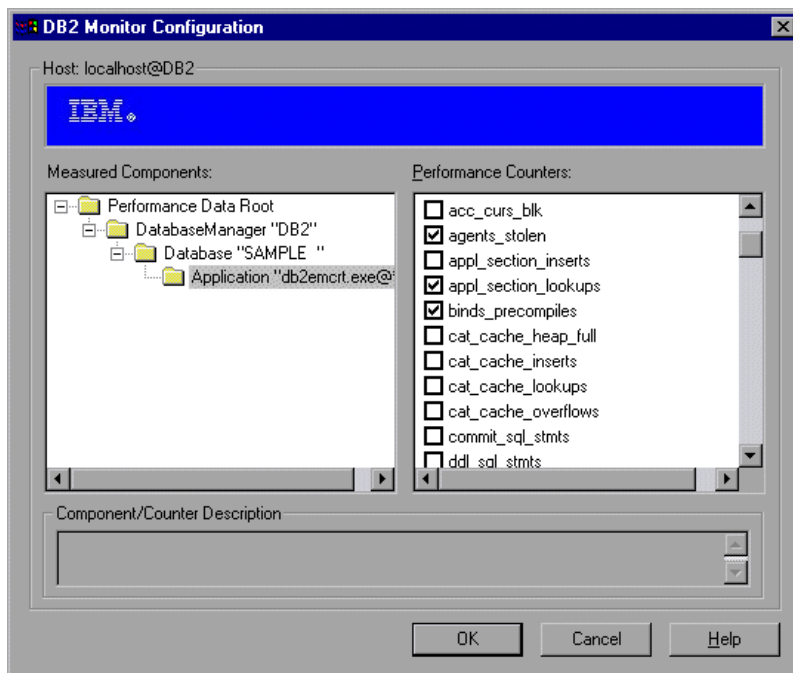
---

#### To configure the DB2 monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select DB2 (in the Database Server Resource category) and then click **Add**.
- 5** In the dialog box that opens, enter your DB2 server username and password, and click **OK**. The DB2 Monitor Configuration dialog box opens.

- Expand the Measured Components tree and select the methods and counters you want to monitor.





The following tables describe the default counters that can be monitored.

### DatabaseManager

| Measurement                      | Description  |
|----------------------------------|--|
| <b>rem_cons_in</b>               | The current number of connections initiated from remote clients to the instance of the database manager that is being monitored.   |
| <b>rem_cons_in_exec</b>          | The number of remote applications that are currently connected to a database and are currently processing a unit of work within the database manager instance being monitored. |
| <b>local_cons</b>                | The number of local applications that are currently connected to a database within the database manager instance being monitored.  |
| <b>local_cons_in_exec</b>        | The number of local applications that are currently connected to a database within the database manager instance being monitored and are currently processing a unit of work.  |
| <b>con_local_databases</b>       | The number of local databases that have applications connected.  |
| <b>agents_registered</b>         | The number of agents registered in the database manager instance that is being monitored (coordinator agents and subagents).   |
| <b>agents_waiting_on_token</b>   | The number of agents waiting for a token so they can execute a transaction in the database manager.  |
| <b>idle_agents</b>               | The number of agents in the agent pool that are currently unassigned to an application and are therefore "idle".   |
| <b>agents_from_pool</b>          | The number of agents assigned from the agent pool  |
| <b>agents_created_empty_pool</b> | The number of agents created because the agent pool was empty.   |

| Measurement                  | Description   |
|------------------------------|---|
| <b>agents_stolen</b>         | The number of times that agents are stolen from an application. Agents are stolen when an idle agent associated with an application is reassigned to work on a different application. |
| <b>comm_private_mem</b>      | The amount of private memory that the instance of the database manager has currently committed at the time of the snapshot.   |
| <b>inactive_gw_agents</b>    | The number of DRDA agents in the DRDA connections pool that are primed with a connection to a DRDA database, but are inactive.  |
| <b>num_gw_conn_switches</b>  | The number of times that an agent from the agents pool was primed with a connection and was stolen for use with a different DRDA database.  |
| <b>sort_heap_allocated</b>   | The total number of allocated pages of sort heap space for all sorts at the level chosen and at the time the snapshot was taken.  |
| <b>post_threshold_sorts</b>  | The number of sorts that have requested heaps after the sort heap threshold has been reached.   |
| <b>piped_sorts_requested</b> | The number of piped sorts that have been requested.   |
| <b>piped_sorts_accepted</b>  | The number of piped sorts that have been accepted.  |

### Database

| Measurement           | Description  |
|-----------------------|--|
| <b>appls_cur_cons</b> | Indicates the number of applications that are currently connected to the database.   |
| <b>appls_in_db2</b>   | Indicates the number of applications that are currently connected to the database, and for which the database manager is currently processing a request. |

| Measurement                      | Description  |
|----------------------------------|--|
| <b>total_sec_cons</b>            | The number of connections made by a sub-agent to the database at the node.   |
| <b>num_assoc_agents</b>          | At the application level, this is the number of sub-agents associated with an application. At the database level, it is the number of sub-agents for all applications. |
| <b>sort_heap_allocated</b>       | The total number of allocated pages of sort heap space for all sorts at the level chosen and at the time the snapshot was taken.                                       |
| <b>total_sorts</b>               | The total number of sorts that have been executed.   |
| <b>total_sort_time</b>           | The total elapsed time (in milliseconds) for all sorts that have been executed.  |
| <b>sort_overflows</b>            | The total number of sorts that ran out of sort heap and may have required disk space for temporary storage.  |
| <b>active_sorts</b>              | The number of sorts in the database that currently have a sort heap allocated.   |
| <b>total_hash_joins</b>          | The total number of hash joins executed.   |
| <b>total_hash_loops</b>          | The total number of times that a single partition of a hash join was larger than the available sort heap space.  |
| <b>hash_join_overflows</b>       | The number of times that hash join data exceeded the available sort heap space   |
| <b>hash_join_small_overflows</b> | The number of times that hash join data exceeded the available sort heap space by less than 10%.   |
| <b>pool_data_l_reads</b>         | Indicates the number of logical read requests for data pages that have gone through the buffer pool.   |
| <b>pool_data_p_reads</b>         | The number of read requests that required I/O to get data pages into the buffer pool.  |

| Measurement                    | Description   |
|--------------------------------|---|
| <b>pool_data_writes</b>        | Indicates the number of times a buffer pool data page was physically written to disk.   |
| <b>pool_index_l_reads</b>      | Indicates the number of logical read requests for index pages that have gone through the buffer pool.   |
| <b>pool_index_p_reads</b>      | Indicates the number of physical read requests to get index pages into the buffer pool.   |
| <b>pool_index_writes</b>       | Indicates the number of times a buffer pool index page was physically written to disk.  |
| <b>pool_read_time</b>          | Provides the total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from disk to buffer pool.  |
| <b>pool_write_time</b>         | Provides the total amount of time spent physically writing data or index pages from the buffer pool to disk.  |
| <b>files_closed</b>            | The total number of database files closed.  |
| <b>pool_async_data_reads</b>   | The number of pages read asynchronously into the buffer pool.   |
| <b>pool_async_data_writes</b>  | The number of times a buffer pool data page was physically written to disk by either an asynchronous page cleaner, or a pre-fetcher. A pre-fetcher may have written dirty pages to disk to make space for the pages being pre-fetched.  |
| <b>pool_async_index_writes</b> | The number of times a buffer pool index page was physically written to disk by either an asynchronous page cleaner, or a pre-fetcher. A pre-fetcher may have written dirty pages to disk to make space for the pages being pre-fetched. |
| <b>pool_async_index_reads</b>  | The number of index pages read asynchronously into the buffer pool by a pre-fetcher.  |
| <b>pool_async_read_time</b>    | The total elapsed time spent reading by database manager pre-fetchers.  |

| Measurement                      | Description  |
|----------------------------------|--|
| <b>pool_async_write_time</b>     | The total elapsed time spent writing data or index pages from the buffer pool to disk by database manager page cleaners.                     |
| <b>pool_async_data_read_reqs</b> | The number of asynchronous read requests.  |
| <b>pool_lsn_gap_clns</b>         | The number of times a page cleaner was invoked because the logging space used had reached a pre-defined criterion for the database.          |
| <b>pool_drty_pg_steal_clns</b>   | The number of times a page cleaner was invoked because a synchronous write was needed during the victim buffer replacement for the database. |
| <b>pool_drty_pg_thrsh_clns</b>   | The number of times a page cleaner was invoked because a buffer pool had reached the dirty page threshold criterion for the database.        |
| <b>prefetch_wait_time</b>        | The time an application spent waiting for an I/O server (pre-fetcher) to finish loading pages into the buffer pool.                          |
| <b>pool_data_to_estore</b>       | The number of buffer pool data pages copied to extended storage.   |
| <b>pool_index_to_estore</b>      | The number of buffer pool index pages copied to extended storage.  |
| <b>pool_data_from_estore</b>     | The number of buffer pool data pages copied from extended storage.   |
| <b>pool_index_from_estore</b>    | The number of buffer pool index pages copied from extended storage.  |
| <b>direct_reads</b>              | The number of read operations that do not use the buffer pool.   |
| <b>direct_writes</b>             | The number of write operations that do not use the buffer pool.  |
| <b>direct_read_reqs</b>          | The number of requests to perform a direct read of one or more sectors of data.  |
| <b>direct_write_reqs</b>         | The number of requests to perform a direct write of one or more sectors of data.   |

| Measurement                    | Description   |
|--------------------------------|---|
| <b>direct_read_time</b>        | The elapsed time (in milliseconds) required to perform the direct reads.  |
| <b>direct_write_time</b>       | The elapsed time (in milliseconds) required to perform the direct writes.   |
| <b>cat_cache_lookups</b>       | The number of times that the catalog cache was referenced to obtain table descriptor information.   |
| <b>cat_cache_inserts</b>       | The number of times that the system tried to insert table descriptor information into the catalog cache.  |
| <b>cat_cache_overflows</b>     | The number of times that an insert into the catalog cache failed due the catalog cache being full.  |
| <b>cat_cache_heap_full</b>     | The number of times that an insert into the catalog cache failed due to a heap-full condition in the database heap.   |
| <b>pkg_cache_lookups</b>       | The number of times that an application looked for a section or package in the package cache. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset. |
| <b>pkg_cache_inserts</b>       | The total number of times that a requested section was not available for use and had to be loaded into the package cache. This count includes any implicit prepares performed by the system.                                |
| <b>pkg_cache_num_overflows</b> | The number of times that the package cache overflowed the bounds of its allocated memory.   |
| <b>appl_section_lookups</b>    | Lookups of SQL sections by an application from its SQL work area.   |
| <b>appl_section_inserts</b>    | Inserts of SQL sections by an application from its SQL work area.   |
| <b>sec_logs_allocated</b>      | The total number of secondary log files that are currently being used for the database.   |

| Measurement             | Description   |
|-------------------------|---|
| <b>log_reads</b>        | The number of log pages read from disk by the logger.   |
| <b>log_writes</b>       | The number of log pages written to disk by the logger.  |
| <b>total_log_used</b>   | The total amount of active log space currently used (in bytes) in the database.   |
| <b>locks_held</b>       | The number of locks currently held.   |
| <b>lock_list_in_use</b> | The total amount of lock list memory (in bytes) that is in use.   |
| <b>deadlocks</b>        | The total number of deadlocks that have occurred.   |
| <b>lock_escals</b>      | The number of times that locks have been escalated from several row locks to a table lock.  |
| <b>x_lock_escals</b>    | The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock. |
| <b>lock_timeouts</b>    | The number of times that a request to lock an object timed-out instead of being granted.  |
| <b>lock_waits</b>       | The total number of times that applications or connections waited for locks.  |
| <b>lock_wait_time</b>   | The total elapsed time waited for a lock.   |
| <b>locks_waiting</b>    | Indicates the number of agents waiting on a lock.   |
| <b>rows_deleted</b>     | The number of row deletions attempted.  |
| <b>rows_inserted</b>    | The number of row insertions attempted.   |
| <b>rows_updated</b>     | The number of row updates attempted.  |
| <b>rows_selected</b>    | The number of rows that have been selected and returned to the application.   |

| Measurement               | Description  |
|---------------------------|--|
| <b>int_rows_deleted</b>   | The number of rows deleted from the database as a result of internal activity.                         |
| <b>int_rows_updated</b>   | The number of rows updated from the database as a result of internal activity.                         |
| <b>int_rows_inserted</b>  | The number of rows inserted into the database as a result of internal activity caused by triggers.     |
| <b>static_sql_stmts</b>   | The number of static SQL statements that were attempted.   |
| <b>dynamic_sql_stmts</b>  | The number of dynamic SQL statements that were attempted.  |
| <b>failed_sql_stmts</b>   | The number of SQL statements that were attempted, but failed.  |
| <b>commit_sql_stmts</b>   | The total number of SQL COMMIT statements that have been attempted.                                    |
| <b>rollback_sql_stmts</b> | The total number of SQL ROLLBACK statements that have been attempted.                                  |
| <b>select_sql_stmts</b>   | The number of SQL SELECT statements that were executed.  |
| <b>uid_sql_stmts</b>      | The number of SQL UPDATE, INSERT, and DELETE statements that were executed.                            |
| <b>ddl_sql_stmts</b>      | This element indicates the number of SQL Data Definition Language (DDL) statements that were executed. |
| <b>int_auto_rebinds</b>   | The number of automatic rebinds (or recompiles) that have been attempted.                              |
| <b>int_commits</b>        | The total number of commits initiated internally by the database manager.                              |
| <b>int_rollbacks</b>      | The total number of rollbacks initiated internally by the database manager.                            |



| Measurement                   | Description   |
|-------------------------------|---|
| <b>int_deadlock_rollbacks</b> | The total number of forced rollbacks initiated by the database manager due to a deadlock. A rollback is performed on the current unit of work in an application selected by the database manager to resolve the deadlock. |
| <b>binds_precompiles</b>      | The number of binds and pre-compiles attempted.   |

### Application

| Measurement                | Description   |
|----------------------------|---|
| <b>agents_stolen</b>       | The number of times that agents are stolen from an application. Agents are stolen when an idle agent associated with an application is reassigned to work on a different application. |
| <b>num_assoc_agents</b>    | At the application level, this is the number of sub-agents associated with an application. At the database level, it is the number of sub-agents for all applications.                |
| <b>total_sorts</b>         | The total number of sorts that have been executed.  |
| <b>total_sort_time</b>     | The total elapsed time (in milliseconds) for all sorts that have been executed.   |
| <b>sort_overflows</b>      | The total number of sorts that ran out of sort heap and may have required disk space for temporary storage.   |
| <b>total_hash_joins</b>    | The total number of hash joins executed.  |
| <b>total_hash_loops</b>    | The total number of times that a single partition of a hash join was larger than the available sort heap space.   |
| <b>hash_join_overflows</b> | The number of times that hash join data exceeded the available sort heap space  |

| Measurement                      | Description  |
|----------------------------------|--|
| <b>hash_join_small_overflows</b> | The number of times that hash join data exceeded the available sort heap space by less than 10%.   |
| <b>pool_data_l_reads</b>         | Indicates the number of logical read requests for data pages that have gone through the buffer pool.   |
| <b>pool_data_p_reads</b>         | The number of read requests that required I/O to get data pages into the buffer pool.  |
| <b>pool_data_writes</b>          | Indicates the number of times a buffer pool data page was physically written to disk.  |
| <b>pool_index_l_reads</b>        | Indicates the number of logical read requests for index pages that have gone through the buffer pool.  |
| <b>pool_index_p_reads</b>        | Indicates the number of physical read requests to get index pages into the buffer pool.  |
| <b>pool_index_writes</b>         | Indicates the number of times a buffer pool index page was physically written to disk.   |
| <b>pool_read_time</b>            | Provides the total amount of elapsed time spent processing read requests that caused data or index pages to be physically read from disk to buffer pool. |
| <b>prefetch_wait_time</b>        | The time an application spent waiting for an I/O server (pre-fetcher) to finish loading pages into the buffer pool.                                      |
| <b>pool_data_to_estore</b>       | The number of buffer pool data pages copied to extended storage.   |
| <b>pool_index_to_estore</b>      | The number of buffer pool index pages copied to extended storage.  |
| <b>pool_data_from_estore</b>     | The number of buffer pool data pages copied from extended storage.   |
| <b>pool_index_from_estore</b>    | The number of buffer pool index pages copied from extended storage.  |

| Measurement                | Description   |
|----------------------------|---|
| <b>direct_reads</b>        | The number of read operations that do not use the buffer pool.  |
| <b>direct_writes</b>       | The number of write operations that do not use the buffer pool.   |
| <b>direct_read_reqs</b>    | The number of requests to perform a direct read of one or more sectors of data.   |
| <b>direct_write_reqs</b>   | The number of requests to perform a direct write of one or more sectors of data.  |
| <b>direct_read_time</b>    | The elapsed time (in milliseconds) required to perform the direct reads.  |
| <b>direct_write_time</b>   | The elapsed time (in milliseconds) required to perform the direct writes.   |
| <b>cat_cache_lookups</b>   | The number of times that the catalog cache was referenced to obtain table descriptor information.   |
| <b>cat_cache_inserts</b>   | The number of times that the system tried to insert table descriptor information into the catalog cache.  |
| <b>cat_cache_overflows</b> | The number of times that an insert into the catalog cache failed due the catalog cache being full.  |
| <b>cat_cache_heap_full</b> | The number of times that an insert into the catalog cache failed due to a heap-full condition in the database heap.   |
| <b>pkg_cache_lookups</b>   | The number of times that an application looked for a section or package in the package cache. At a database level, it indicates the overall number of references since the database was started, or monitor data was reset. |
| <b>pkg_cache_inserts</b>   | The total number of times that a requested section was not available for use and had to be loaded into the package cache. This count includes any implicit prepares performed by the system.                                |

| <b>Measurement</b>          | <b>Description</b>  |
|-----------------------------|---|
| <b>appl_section_lookups</b> | Lookups of SQL sections by an application from its SQL work area.   |
| <b>appl_section_inserts</b> | Inserts of SQL sections by an application from its SQL work area.   |
| <b>uow_log_space_used</b>   | The amount of log space (in bytes) used in the current unit of work of the monitored application.   |
| <b>locks_held</b>           | The number of locks currently held.   |
| <b>deadlocks</b>            | The total number of deadlocks that have occurred.   |
| <b>lock_escals</b>          | The number of times that locks have been escalated from several row locks to a table lock.  |
| <b>x_lock_escals</b>        | The number of times that locks have been escalated from several row locks to one exclusive table lock, or the number of times an exclusive lock on a row caused the table lock to become an exclusive lock. |
| <b>lock_timeouts</b>        | The number of times that a request to lock an object timed-out instead of being granted.  |
| <b>lock_waits</b>           | The total number of times that applications or connections waited for locks.  |
| <b>lock_wait_time</b>       | The total elapsed time waited for a lock.   |
| <b>locks_waiting</b>        | Indicates the number of agents waiting on a lock.   |
| <b>uow_lock_wait_time</b>   | The total amount of elapsed time this unit of work has spent waiting for locks.   |
| <b>rows_deleted</b>         | The number of row deletions attempted.  |
| <b>rows_inserted</b>        | The number of row insertions attempted.   |
| <b>rows_updated</b>         | The number of row updates attempted.  |
| <b>rows_selected</b>        | The number of rows that have been selected and returned to the application.   |

| Measurement              | Description  |
|--------------------------|--|
| <b>rows_written</b>      | The number of rows changed (inserted, deleted or updated) in the table.  |
| <b>rows_read</b>         | The number of rows read from the table.  |
| <b>int_rows_deleted</b>  | The number of rows deleted from the database as a result of internal activity.   |
| <b>int_rows_updated</b>  | The number of rows updated from the database as a result of internal activity.   |
| <b>int_rows_inserted</b> | The number of rows inserted into the database as a result of internal activity caused by triggers.                           |
| <b>open_rem_curs</b>     | The number of remote cursors currently open for this application, including those cursors counted by 'open_rem_curs_blk'.    |
| <b>open_rem_curs_blk</b> | The number of remote blocking cursors currently open for this application.   |
| <b>rej_curs_blk</b>      | The number of times that a request for an I/O block at server was rejected and the request was converted to non-blocked I/O. |
| <b>acc_curs_blk</b>      | The number of times that a request for an I/O block was accepted.  |
| <b>open_loc_curs</b>     | The number of local cursors currently open for this application, including those cursors counted by 'open_loc_curs_blk'.     |
| <b>open_loc_curs_blk</b> | The number of local blocking cursors currently open for this application.  |
| <b>static_sql_stmts</b>  | The number of static SQL statements that were attempted.   |
| <b>dynamic_sql_stmts</b> | The number of dynamic SQL statements that were attempted.  |
| <b>failed_sql_stmts</b>  | The number of SQL statements that were attempted, but failed.  |

| Measurement                   | Description   |
|-------------------------------|---|
| <b>commit_sql_stmts</b>       | The total number of SQL COMMIT statements that have been attempted.   |
| <b>rollback_sql_stmts</b>     | The total number of SQL ROLLBACK statements that have been attempted.   |
| <b>select_sql_stmts</b>       | The number of SQL SELECT statements that were executed.   |
| <b>uid_sql_stmts</b>          | The number of SQL UPDATE, INSERT, and DELETE statements that were executed.   |
| <b>ddl_sql_stmts</b>          | This element indicates the number of SQL Data Definition Language (DDL) statements that were executed.  |
| <b>int_auto_rebinds</b>       | The number of automatic rebinds (or recompiles) that have been attempted.   |
| <b>int_commits</b>            | The total number of commits initiated internally by the database manager.   |
| <b>int_rollbacks</b>          | The total number of rollbacks initiated internally by the database manager.   |
| <b>int_deadlock_rollbacks</b> | The total number of forced rollbacks initiated by the database manager due to a deadlock. A rollback is performed on the current unit of work in an application selected by the database manager to resolve the deadlock. |
| <b>binds_precompiles</b>      | The number of binds and pre-compiles attempted.   |

- 7** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.
- 8** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the Oracle Monitor

The Oracle server measures information from the V\$SESSTAT and V\$SYSSTAT Oracle V\$ tables, and other table counters defined by the user in the custom query. In order to monitor the Oracle server, you must set up the monitoring environment as described below.

---

**Note:** The port you use to monitor an Oracle server through a firewall depends on the configuration of the Oracle server. Configuration information for the connection between the client and server is located in the Oracle client *tnsnames.ora* file.

---

### To set up the Oracle monitor environment:

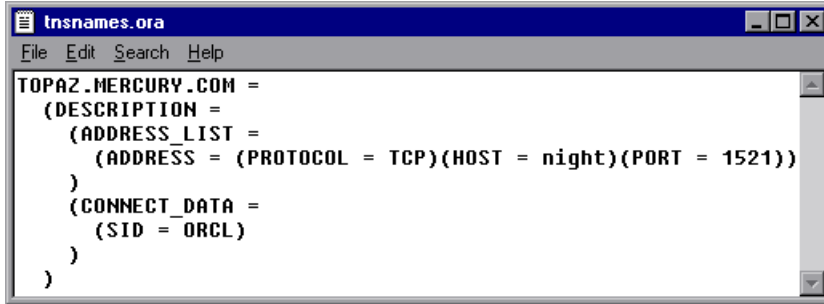
- 1** Ensure that the Oracle client libraries are installed on the Console machine.
- 2** Verify that `%OracleHome%\bin` is included in the path environment variable. If it is not, add it.
- 3** Configure the *tnsnames.ora* file on the Console machine so that the Oracle client can communicate with the Oracle server(s) you plan to monitor.

You can configure connection parameters either manually, by editing the *tnsnames.ora* file in a text editor, or using the Oracle service configuration tool (for example, select **Start > Programs > Oracle for Windows NT > Oracle Net8 Easy Config**).

You specify:

- a new service name (TNS name) for the Oracle instance
- TCP protocol
- the host name (name of monitored server machine)
- the port number (usually 1521)
- the database SID (the default SID is ORCL)

For example:



```

tnsnames.ora
File Edit Search Help
TOPAZ.MERCURY.COM =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = night)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SID = ORCL)
    )
  )

```

---

**Note:** Only the 32-bit Oracle client should be installed on the Console machine running the Oracle monitor. If you have a 16-bit and a 32-bit Oracle client installation on the Console machine, the 16-bit installation should be uninstalled.

---

- 4** Obtain a username and password for the service from your database administrator, and ensure that the Console has database administrator privileges for the Oracle V\$ tables (V\$SESSTAT, V\$SYSSTAT, V\$STATNAME, V\$INSTANCE, V\$SESSION).
  - 5** Verify connection with the Oracle server by performing *tns ping* from the Console machine. Note that there may be a problem connecting if the Oracle server is behind a DMZ/firewall that limits its communication to application servers accessing it.
  - 6** Ensure that the registries are updated for the version of Oracle that you are using and that they have the following key:  
HKEY\_LOCAL\_MACHINE\SOFTWARE\ORACLE
  - 7** Verify that the Oracle server you want to monitor is up and running.
- 

**Note:** It is possible to monitor several Oracle database servers concurrently.

---



- 8 Run SQL\*Plus from the Console and attempt to log in to the Oracle server(s) with the desired username/password/server combination.
- 9 Type `SELECT * FROM V$SYSSTAT` to verify that you can view the V\$SYSSTAT table on the Oracle server. Use similar queries to verify that you can view the V\$SESSTAT, V\$SESSION, V\$INSTANCE, V\$STATNAME, and V\$PROCESS tables on the server. Make sure that the Oracle bin directory is in the search path.
- 10 To change the length of each monitoring sample (in seconds), you need to edit the `dat\monitors\vmmon.cfg` file in the ProTune root folder. The default rate is 10 seconds.

---

**Note:** The minimum sampling rate for the Oracle Monitor is 10 seconds. If you set the sampling rate at less than 10 seconds, the Oracle Monitor will continue to monitor at 10 second intervals.

---

---

**Note:** If a problem occurs in setting up the Oracle environment, view the error message issued by the Oracle server.

---


#### To configure the Oracle monitor:




- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.

- 4 In the left section of the dialog box, select Oracle (in the Database Server Resource category) and then click **Add**.

The Oracle Logon dialog box opens.



- 5 Enter your Login Name, Password, and Server Name, and click **OK**. The Add Oracle Measurements dialog box opens.



- 6 Select an object, a measurement, and an instance. You can select multiple measurements using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted measurement are running. For a description of each measurement, click **Explain>>** to expand the dialog box. For instructions on creating custom queries, see “Custom Queries,” on page 410.

The following measurements are most commonly used when monitoring the Oracle server (from the V\$SYSSTAT table):

| Measurement                                   | Description   |
|---|---|
| <b>CPU used by this session</b>               | This is the amount of CPU time (in 10s of milliseconds) used by a session between the time a user call started and ended. Some user calls can be completed within 10 milliseconds and, as a result, the start and end user-call time can be the same. In this case, 0 milliseconds are added to the statistic. A similar problem can exist in the operating system reporting, especially on systems that suffer from many context switches. |
| <b>Bytes received via SQL*Net from client</b> | The total number of bytes received from the client over Net8  |
| <b>Logons current</b>                         | The total number of current logons  |
| <b>Opens of replaced files</b>                | The total number of files that needed to be reopened because they were no longer in the process file cache  |
| <b>User calls</b>                             | Oracle allocates resources (Call State Objects) to keep track of relevant user call data structures every time you log in, parse, or execute. When determining activity, the ratio of user calls to RPI calls gives you an indication of how much internal work gets generated as a result of the type of requests the user is sending to Oracle.   |
| <b>SQL*Net roundtrips to/from client</b>      | The total number of Net8 messages sent to, and received from, the client  |
| <b>Bytes sent via SQL*Net to client</b>       | The total number of bytes sent to the client from the foreground process(es)  |
| <b>Opened cursors current</b>                 | The total number of current open cursors  |

| Measurement             | Description  |
|-------------------------|--|
| <b>DB block changes</b> | Closely related to consistent changes, this statistic counts the total number of changes that were made to all blocks in the SGA that were part of an update or delete operation. These are changes that are generating redo log entries and hence will be permanent changes to the database if the transaction is committed. This statistic is a rough indication of total database work and indicates (possibly on a per-transaction level) the rate at which buffers are being dirtied. |
| <b>Total file opens</b> | The total number of file opens being performed by the instance. Each process needs a number of files (control file, log file, database file) in order to work against the database.  |

- 7** Click **Add** to place the selected measurement on the resource list. Add all the desired resources to the list, and click **Close**.
- 8** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

---

**Note:** By default, the database returns the absolute value of a counter. However, by changing the `IsRate` setting in the `dat\monitors\vmmon.cfg` file to 1, you can instruct the database to report a counter's rate value—the change in the counter per unit time.

---

### Custom Queries

Using the custom query feature, you can define your own query to the Oracle database and view the result of this query—a single numerical value—in the Oracle online monitor graph. By defining your own query, you can monitor not only the V\$SYSSTAT and V\$SESSTAT table counters that are currently provided by the Oracle monitor, but other tables that contain useful performance information as well.

**To create a custom query:**

- 1** In the third line of the *vmon.cfg* file, `CustomCounters=`, indicate the number of custom counters you want to create.
- 2** Create a new section in the *vmon.cfg* file for the new counter. Each section has the following format:

```
[Custom2]
```

```
Name=Number of sessions
```

```
Description=This counter returns the number of sessions active.
```

```
Query=SELECT COUNT(*) FROM V$SESSION
```

```
IsRate=1
```

- 3** In the `[Custom#]` line, assign the next number in the sequence of counters to the new custom counter. Note that the custom counters must be in consecutive order, beginning with the number 0.
- 4** In the `Name` line, enter the name of the new counter.
- 5** In the `Description` line, enter the description of the counter that you want the help message to contain.
- 6** In the `Query` line, enter the text of the SQL query (on one line of the *vmon.cfg* file) that returns exactly one row from the database. This row must contain one column, a numerical value.

---

**Note:** Custom queries should not exceed 512 characters.

---

- 7** In the `IsRate` line, enter 0 if you want the database to report the counter as an absolute number. If you want the database to report the change in the counter per unit time, enter 1.

---

**Note:** Custom queries cannot return negative values.

---

## Configuring the SQL Server Monitor

The SQL Server monitor measures the standard Windows resources on the SQL server machine.

---

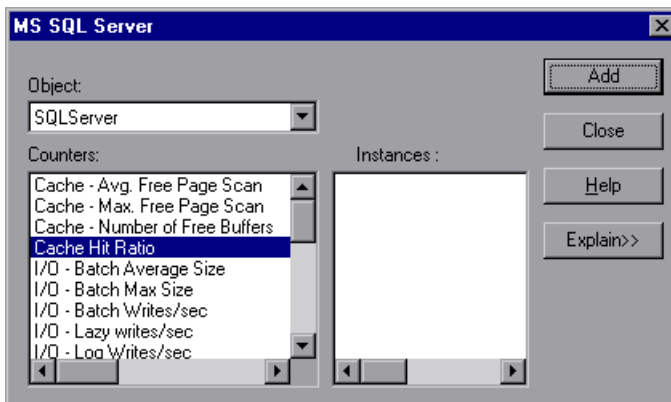
**Note:** To monitor an SQL server through a firewall, use TCP, port 139.

---

### To configure the SQL server monitor:



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select MS SQL Server (in the Database Server Resource category) and then click **Add**. A dialog box displaying the SQL Server object, its counters, and instances opens.



The following table describes the default counters that can be monitored on version 6.5 of the SQL Server:

| Measurement                        | Description   |
|------------------------------------|---|
| <b>% Total Processor Time (NT)</b> | The average percentage of time that all the processors on the system are busy executing non-idle threads. On a multi-processor system, if all processors are always busy, this is 100%, if all processors are 50% busy this is 50% and if 1/4th of the processors are 100% busy this is 25%. It can be viewed as the fraction of the time spent doing useful work. Each processor is assigned an Idle thread in the Idle process which consumes those unproductive processor cycles not used by any other threads.  |
| <b>% Processor Time (Win 2000)</b> | The percentage of time that the processor is executing a non-idle thread. This counter was designed as a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the idle process in each sample interval, and subtracting that value from 100%. (Each processor has an idle thread which consumes cycles when no other threads are ready to run). It can be viewed as the percentage of the sample interval spent doing useful work. This counter displays the average percentage of busy time observed during the sample interval. It is calculated by monitoring the time the service was inactive, and then subtracting that value from 100%. |
| <b>Cache Hit Ratio</b>             | The percentage of time that a requested data page was found in the data cache (instead of being read from disk)   |
| <b>I/O - Batch Writes/sec</b>      | The number of 2K pages written to disk per second, using Batch I/O. The checkpoint thread is the primary user of Batch I/O.   |
| <b>I/O - Lazy Writes/sec</b>       | The number of 2K pages flushed to disk per second by the Lazy Writer  |
| <b>I/O - Outstanding Reads</b>     | The number of physical reads pending  |
| <b>I/O - Outstanding Writes</b>    | The number of physical writes pending   |

| Measurement            | Description  |
|------------------------|--|
| I/O - Page Reads/sec   | The number of physical page reads per second                   |
| I/O - Transactions/sec | The number of Transact-SQL command batches executed per second |
| User Connections       | The number of open user connections                            |

---

**Note:** To change the default counters for the SQL Server monitor, see “Changing a Monitor’s Default Counters,” on page 661.

---

- 5** Select a counter and an instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.
- 6** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.
- 7** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

---

**Note:** Certain measurements or counters are especially useful for determining server performance and isolating the cause of a bottleneck during an initial stress test on the SQL Server. For more information about these counters, see “Useful Counters for Stress Testing,” on page 662.

---



## Configuring the Sybase Monitor

The Sybase monitor enables monitoring of Sybase Adaptive Server Enterprise (Sybase ASE) servers (version 11 or later) on Windows and UNIX. The monitor connects to the Sybase ASE server (via the Adaptive Server Enterprise Monitor Server) and retrieves metrics from the server using standard, Sybase-provided libraries.

---

**Note:** When connecting to the monitored server, you connect to the Adaptive Server Enterprise Monitor Server, not the Sybase ASE server. The Adaptive Server Enterprise Monitor Server is an application that runs on the same machine as Sybase ASE server and retrieves performance information from it. The Adaptive Server Enterprise Monitor Server usually has the same name as the Sybase server, but with the suffix *\_ms*.

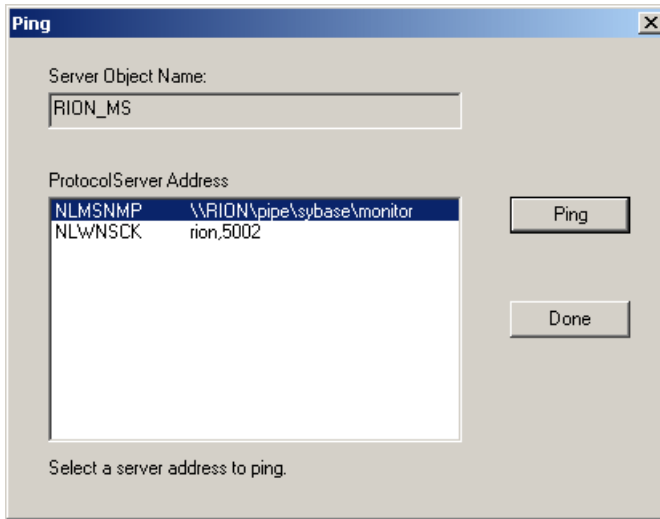
---

In order to monitor the Sybase ASE server, you must first set up the Sybase monitor environment.

**To set up the Sybase monitor environment:**

- 1 Install the Sybase client files and libraries on the Console machine.

- 2 Verify a connection between the client and server on the Console machine. To do so, use the Sybase client's *dsedit* tool to ping the Adaptive Server Enterprise Monitor Server.



---

**Note:** The port you use to monitor a Sybase server through a firewall depends on the configuration of the Sybase server. Configuration information for the connection between the client and server is located in the Sybase client *sql.ini* file.

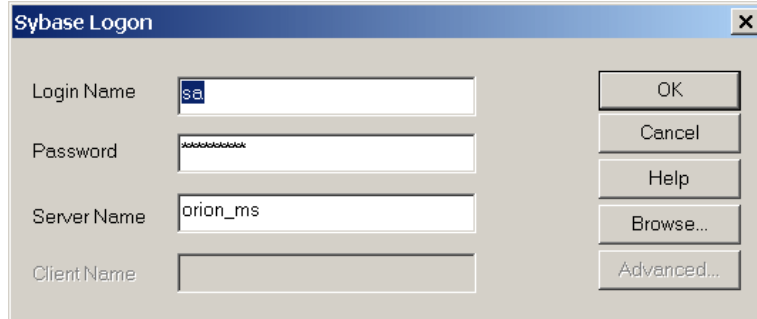
---

**To configure the Sybase ASE monitor:**

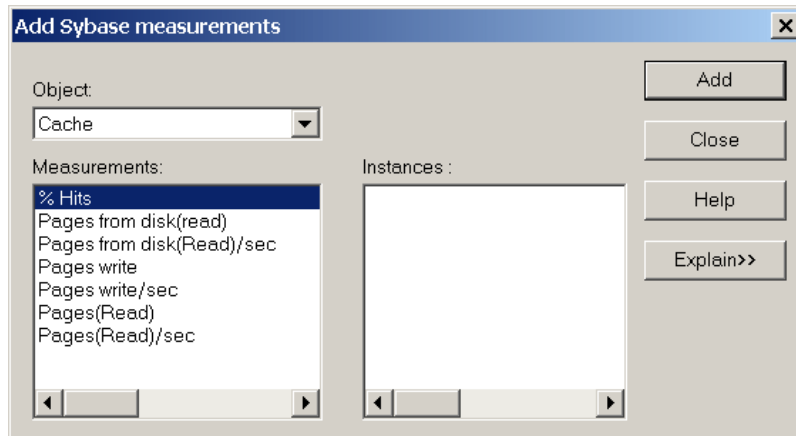


- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select Sybase (in the Database Server Resource category) and then click **Add**.

The Sybase Logon dialog box opens.



- 5 Enter the login name and password of a user that has administrative privileges on the Sybase ASE server, as well as the Adaptive Server Enterprise Monitor Server name (usually the same name as the Sybase server but with the suffix *\_ms*).
- 6 Click **OK**. The Add Sybase Measurements dialog box opens.



- 7 Select an object, measurement, and instance. You can select multiple measurements using the **CTRL** key. The instance is relevant only if multiple instances of the highlighted measurement are running. For a description of the measurements, click **Explain>>** to expand the dialog box.

The following measurements are available when monitoring a Sybase server:

| Object  | Measurement                | Description  |
|---------|----------------------------|--|
| Network | Average packet size (Read) | Reports the number of network packets received                             |
|         | Average packet size (Send) | Reports the number of network packets sent                                 |
|         | Network bytes (Read)       | Reports the number of bytes received, over the sampling interval           |
|         | Network bytes (Read)/sec   | Reports the number of bytes received, per second                           |
|         | Network bytes (Send)       | Reports the number of bytes sent, over the sampling interval               |
|         | Network bytes (Send)/sec   | Reports the number of bytes sent, per second                               |
|         | Network packets (Read)     | Reports the number of network packets received, over the sampling interval |
|         | Network packets (Read)/sec | Reports the number of network packets received, per second                 |
|         | Network packets (Send)     | Reports the number of network packets sent, over the sampling interval     |
|         | Network packets (Send)/sec | Reports the number of network packets sent, per second                     |
| Memory  | Memory                     | Reports the amount of memory, in bytes, allocated for the page cache       |
| Disk    | Reads                      | Reports the number of reads made from a database device                    |
|         | Writes                     | Reports the number of writes made to a database device                     |
|         | Waits                      | Reports the number of times that access to a device had to wait            |

| Object            | Measurement                        | Description   |
|-------------------|------------------------------------|---|
| Disk              | Grants                             | Reports the number of times access to a device was granted                                    |
| Engine            | Server is busy (%)                 | Reports the percentage of time during which the Adaptive Server is in a "busy" state          |
|                   | CPU time                           | Reports how much "busy" time was used by the engine   |
|                   | Logical pages (Read)               | Reports the number of data page reads, whether satisfied from cache or from a database device |
|                   | Pages from disk (Read)             | Reports the number of data page reads that could not be satisfied from the data cache         |
|                   | Pages stored                       | Reports the number of data pages written to a database device                                 |
| Stored Procedures | Executed (sampling period)         | Reports the number of times a stored procedure was executed, over the sampling interval       |
|                   | Executed (session)                 | Reports the number of times a stored procedure was executed, during the session               |
|                   | Average duration (sampling period) | Reports the time, in seconds, spent executing a stored procedure, over the sampling interval  |
|                   | Average duration (session)         | Reports the time, in seconds, spent executing a stored procedure, during the session          |
| Locks             | % Requests                         | Reports the percentage of successful requests for locks                                       |
|                   | Locks count                        | Reports the number of locks. This is an accumulated value.                                    |

| Object  | Measurement               | Description  |
|---------|---------------------------|--|
| Locks   | Granted immediately       | Reports the number of locks that were granted immediately, without having to wait for another lock to be released          |
|         | Granted after wait        | Reports the number of locks that were granted after waiting for another lock to be released                                |
|         | Not granted               | Reports the number of locks that were requested but not granted  |
|         | Wait time (avg.)          | Reports the average wait time for a lock   |
| SqlSrvr | Locks/sec                 | Reports the number of locks. This is an accumulated value.   |
|         | % Processor time (server) | Reports the percentage of time that the Adaptive Server is in a "busy" state   |
|         | Transactions              | Reports the number of committed Transact-SQL statement blocks (transactions)   |
|         | Deadlocks                 | Reports the number of deadlocks  |
| Cache   | % Hits                    | Reports the percentage of times that a data page read could be satisfied from cache without requiring a physical page read |
|         | Pages (Read)              | Reports the number of data page reads, whether satisfied from cache or from a database device                              |
|         | Pages (Read)/sec          | Reports the number of data page reads, whether satisfied from cache or from a database device, per second                  |

| Object      | Measurement                | Description   |
|-------------|----------------------------|---|
| Cache       | Pages from disk (Read)     | Reports the number of data page reads that could not be satisfied from the data cache   |
|             | Pages from disk (Read)/sec | Reports the number of data page reads, per second, that could not be satisfied from the data cache  |
|             | Pages (Write)              | Reports the number of data pages written to a database device   |
|             | Pages (Write)/sec          | Reports the number of data pages written to a database device, per second   |
| Process     | % Processor time (process) | Reports the percentage of time that a process running a given application was in the "Running" state (out of the time that all processes were in the "Running" state) |
|             | Locks/sec                  | Reports the number of locks, by process. This is an accumulated value.  |
|             | % Cache hit                | Reports the percentage of times that a data page read could be satisfied from cache without requiring a physical page read, by process                                |
|             | Pages (Write)              | Reports the number of data pages written to a database device, by process   |
| Transaction | Transactions               | Reports the number of committed Transact-SQL statement blocks (transactions), during the session  |
|             | Rows (Deleted)             | Reports the number of rows deleted from database tables during the session  |

| Object      | Measurement          | Description   |
|-------------|----------------------|---|
| Transaction | Inserts              | Reports the number of insertions into a database table during the session   |
|             | Updates              | Reports the updates to database tables during the session   |
|             | Updates in place     | Reports the sum of expensive, in-place and not-in-place updates (everything except updates deferred) during the session |
|             | Transactions/sec     | Reports the number of committed Transact-SQL statement blocks (transactions) per second                                 |
|             | Rows (Deleted)/sec   | Reports the number of rows deleted from database tables, per second   |
|             | Inserts/sec          | Reports the number of insertions into a database table, per second  |
|             | Updates/sec          | Reports the updates to database tables, per second  |
|             | Updates in place/sec | Reports the sum of expensive, in-place and not-in-place updates (everything except updates deferred), per second        |

- 8 Click **Add** to place the selected measurement on the resource list. Add all the desired resources to the list, and click **Close**.
- 9 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.



# 24

---

## Streaming Media Monitoring

During a session run, you can monitor the Windows Media Server and RealPlayer audio/video servers, as well as the RealPlayer and Media Player clients, in order to isolate server and client performance bottlenecks.

This chapter describes:

- ▶ Configuring the Windows Media Server Monitor
- ▶ Configuring the RealPlayer Server Monitor
- ▶ Viewing the RealPlayer Client Online Graph
- ▶ Viewing the Media Player Client Online Graph

---

**Note:** For instructions on recording a script containing streaming media functions, refer to the *ProTune Virtual User Generator User's Guide*.

---

### About Streaming Media Monitoring

The streaming media monitors provide you with performance information for the Windows Media Server and RealPlayer audio/video servers, as well as the RealPlayer and Media Player clients. In order to obtain data for the Windows Media Server and RealPlayer Server, you need to activate the streaming media monitor before executing the session, and indicate which statistics and measurements you want to monitor. The RealPlayer Client and Media Player Client do not require pre-session activation or configuration.

## Configuring the Windows Media Server Monitor

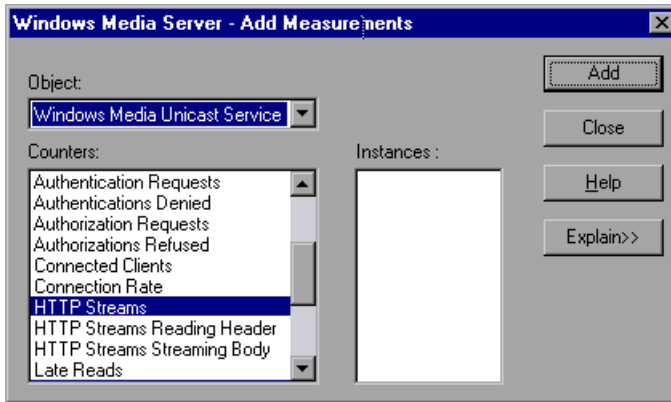
To monitor the Windows Media Server, you must first select the counters you want the Windows Media Server monitor to measure. You select these counters using the Windows Media Server dialog box.

**To configure the Windows Media Server monitor:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Windows Media Server (in the Streaming Media Server category) and then click **Add**.

The Windows Media Server dialog box opens, displaying the Windows Media Unicast Service object, its counters, and instances.



The following table describes the default counters that can be monitored:

| Measurement                                  | Description  |
|--|--|
| <b>Active Live Unicast Streams (Windows)</b> | The number of live unicast streams that are being streamed |
| <b>Active Streams</b>                        | The number of streams that are being streamed              |

| Measurement         | Description  |
|---------------------|--|
| Active TCP Streams  | The number of TCP streams that are being streamed  |
| Active UDP Streams  | The number of UDP streams that are being streamed  |
| Aggregate Read Rate | The total, aggregate rate (bytes/sec) of file reads  |
| Aggregate Send Rate | The total, aggregate rate (bytes/sec) of stream transmission   |
| Connected Clients   | The number of clients connected to the server  |
| Connection Rate     | The rate at which clients are connecting to the server   |
| Consoles            | The number of Consoles currently connected to the server   |
| HTTP Streams        | The number of HTTP streams being streamed  |
| Late Reads          | The number of late read completions per second   |
| Pending Connections | The number of clients that are attempting to connect to the server, but are not yet connected. This number may be high if the server is running near maximum capacity and cannot process a large number of connection requests in a timely manner. |
| Stations            | The number of station objects that currently exist on the server   |
| Streams             | The number of stream objects that currently exist on the server  |
| Stream Errors       | The cumulative number of errors occurring per second   |

- 5 Select a counter and an instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.

- 6 Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.
- 7 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the RealPlayer Server Monitor

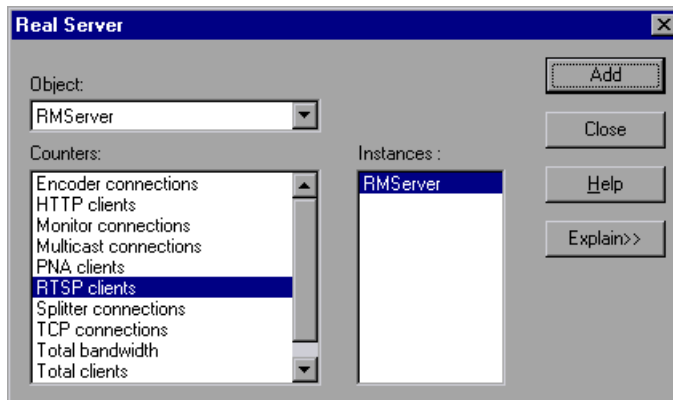
To monitor the RealPlayer Server, you must first select the counters you want the RealPlayer Server monitor to measure. You select these counters using the Real Server dialog box.

**To configure the RealPlayer Server monitor:**



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select RealPlayer Server (in the Streaming Media Server category) and then click **Add**.

The Real Server dialog box opens, displaying the counters that can be monitored.



- 5** Select a counter and an instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.

The following table describes the default counters that can be monitored:

| Measurement           | Description                                     |
|-----------------------|---|
| Encoder Connections   | The number of active encoder connections        |
| HTTP Clients          | The number of active clients using HTTP         |
| Monitor Connections   | The number of active server monitor connections |
| Multicast Connections | The number of active multicast connections      |
| PNA Clients           | The number of active clients using PNA          |
| RTSP Clients          | The number of active clients using RTSP         |
| Splitter Connections  | The number of active splitter connections       |
| TCP Connections       | The number of active TCP connections            |
| Total Bandwidth       | The number of bits per second being consumed    |
| Total Clients         | The total number of active clients              |
| UDP Clients           | The number of active UDP connections            |

- 6** Click **Add** to place the selected counter on the resource list. Add all the desired resources to the list, and click **Close**.
- 7** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Viewing the RealPlayer Client Online Graph

You can view the RealPlayer Client online monitor graph by dragging it from the Available Graphs tree into the right pane of the Execute tab.

The following table describes the RealPlayer Client measurements that are monitored:

| Measurement                                  | Description   |
|--|---|
| <b>Current Bandwidth (Kbits/sec)</b>         | The number of kilobytes in the last second  |
| <b>Buffering Event Time (sec)</b>            | The average time spent on buffering   |
| <b>Network Performance</b>                   | The ratio (percentage) between the current bandwidth and the actual bandwidth of the clip |
| <b>Percentage of Recovered Packets</b>       | The percentage of error packets that were recovered                                       |
| <b>Percentage of Lost Packets</b>            | The percentage of packets that were lost  |
| <b>Percentage of Late Packets</b>            | The percentage of late packets  |
| <b>Time to First Frame Appearance (sec)</b>  | The time for first frame appearance (measured from the start of the replay)               |
| <b>Number of Buffering Events</b>            | The average number of all buffering events  |
| <b>Number of Buffering Seek Events</b>       | The average number of buffering events resulting from a seek operation                    |
| <b>Buffering Seek Time</b>                   | The average time spent on buffering events resulting from a seek operation                |
| <b>Number of Buffering Congestion Events</b> | The average number of buffering events resulting from network congestion                  |
| <b>Buffering Congestion Time</b>             | The average time spent on buffering events resulting from network congestion              |

| Measurement                                  | Description  |
|--|--|
| <b>Number of Buffering Live Pause Events</b> | The average number of buffering events resulting from live pause     |
| <b>Buffering Live Pause Time</b>             | The average time spent on buffering events resulting from live pause |

## Viewing the Media Player Client Online Graph

You can view the Windows Media Player Client online monitor graph by dragging it from the Available Graphs tree into the right pane of the Execute tab.

The following table describes the Media Player Client measurements that are monitored:

| Measurement                              | Description  |
|--|--|
| <b>Stream Quality (Packet-level)</b>     | The percentage ratio of packets received to total packets  |
| <b>Current bandwidth (Kbits/sec)</b>     | The number of kbits per second received  |
| <b>Stream Packet Rate</b>                | The number of packets received   |
| <b>Total number of recovered packets</b> | The number of lost packets that were recovered. This value is only relevant during network playback.     |
| <b>Total number of lost packets</b>      | The number of lost packets that were not recovered. This value is only relevant during network playback. |
| <b>Stream Quality (Sampling-level)</b>   | The percentage of stream samples received on time (no delays in reception)                               |





# 25

---

## ERP/CRM Server Resource Monitoring

During a session step run, you can monitor ERP server resources in order to isolate server performance bottlenecks.

This chapter explains how to configure the SAP, SAP Portal, Siebel, and Siebel Server Manager ERP/CRM Server Resource Monitors.

It describes:

- Setting up the Monitoring Environment
- Configuring the SAP Monitor
- Configuring the SAP Portal Monitor
- Configuring the Siebel Monitor
- Configuring the Siebel Server Manager Monitor

### About ERP/CRM Server Resource Monitoring

Siebel, Siebel Server Manager, and SAP Portal monitors are ERP/CRM Server Resource Monitors, that provide performance information such as the number of open sessions, active transactions, and database connections for the Siebel and SAP Portal ERP/CRM application servers. You can monitor the server resources on a machine during a session step run, and determine why a bottleneck occurred on a particular machine.

## Setting up the Monitoring Environment

To monitor the Siebel, Siebel Server Manager, SAP, and SAP Portal server performance, you must first install and configure SiteScope. SiteScope is the application that is used to monitor these servers.

### Before setting up the ERP/CRM Server Resource monitors:

- 1 Make sure that SiteScope has been installed on a dedicated machine.

---

**Note:** It is recommended that you install SiteScope and the ProTune Console on different machines.

---

- 2 On the machine where SiteScope is installed, configure SiteScope to monitor the required ERP/CRM server machines. For more information on configuring the SiteScope server, refer to the SiteScope User Guide (<http://www.freshwater.com/SiteScope/UGtoc.htm>).

---

**Note:** SiteScope's default sampling rate is 10 minutes, and its minimum rate 15 seconds.

---

- 3 Verify that SiteScope is collecting the required data from the servers it is monitoring.
- 4 In the System Topology window, add an element representing the server running Siebel. Note that this element must not be a Load Generator.

## Setting Up the SAP Monitor

Before monitoring a SAP R/3 system server, you must set up the server monitor environment.

**To set up the SAP monitor environment:**

- 1** Install the SAP GUI client on the Console machine.
- 2** Click **F6** to check whether you can access the st03 transaction and query for *last minute load* information. If this functionality is not already enabled, enable it from the SAP R/3 client on the Console machine, using the username and password defined in the Console.

## Configuring the SAP Monitor

To monitor a SAP R/3 system server, you must select the counters you want the SAP monitor to measure. You select these counters using the Add SAP Monitor Measurements dialog box.

---

**Note:** The SAP R/3 performance monitor supports SAP server versions 3.1 to 4.6, regardless of the SAP R/3 server's operating system and the platform on which it is installed.

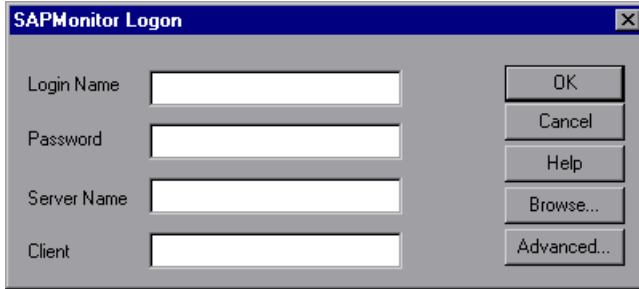
---

**To configure the SAP monitor on the Console machine:**



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** From the Server list, choose the server running Siebel.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select SAP (in the ERP/CRM Server Resource Graphs category) and then click **Add**.

The SAP Monitor Logon dialog box opens.



- 5 Enter your Login Name, Password, Server Name, and Client.

---

**Note:** If you want to connect to the SAP monitor through a router, you need to enter the router string into the Server Name field. A router string has the format:

*<RouterString/ServerIP/S/sapdpxx>*

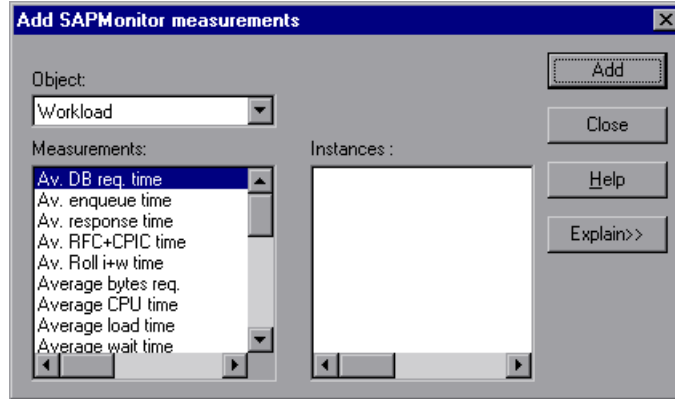
where RouterString is /H/<IP\_ADDRESS>/H/<IP\_ADDRESS>/H/  
ServerIP is the application server IP address  
and xx is the system number.

For example, if the router string = /H/199.35.107.9/H/204.79.199.244/H/,  
application server IP address = 172.20.11.6, and the system number = 00, you  
should enter the following string into the Server Name field:

*/H/199.35.107.9/H/204.79.199.244/H/172.20.11.6/S/sapdp00*

---

- 6 Click **OK**. The Add SAP Monitor Measurements dialog box opens.



- 7 Select an object, a measurement, and an instance. You can select multiple measurements using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted measurement are running. For a description of each measurement, click **Explain>>** to expand the dialog box.

The following are the most commonly monitored counters:

| Measurement           | Description  |
|-----------------------|--|
| Average CPU time      | The average CPU time used in the work process.   |
| Average response time | The average response time, measured from the time a dialog sends a request to the dispatcher work process, through the processing of the dialog, until the dialog is completed and the data is passed to the presentation layer. The response time between the SAP GUI and the dispatcher is not included in this value.   |
| Average wait time     | The average amount of time that an unprocessed dialog step waits in the dispatcher queue for a free work process. Under normal conditions, the dispatcher work process should pass a dialog step to the application process immediately after receiving the request from the dialog step. Under these conditions, the average wait time would be a few milliseconds. A heavy load on the application server or on the entire system causes queues at the dispatcher queue.                                 |
| Average load time     | The time needed to load and generate objects, such as ABAP source code and screen information, from the database.  |
| Database calls        | The number of parsed requests sent to the database.  |
| Database requests     | The number of logical ABAP requests for data in the database. These requests are passed through the R/3 database interface and parsed into individual database calls. The proportion of database calls to database requests is important. If access to information in a table is buffered in the SAP buffers, database calls to the database server are not required. Therefore, the ratio of calls/requests gives an overall indication of the efficiency of table buffering. A good ratio would be 1:10. |
| Roll ins              | The number of rolled-in user contexts.   |
| Roll outs             | The number of rolled-out user contexts.  |

| Measurement                      | Description  |
|----------------------------------|--|
| Roll in time                     | The processing time for roll ins.  |
| Roll out time                    | The processing time for roll outs.   |
| Roll wait time                   | The queue time in the roll area. When synchronous RFCs are called, the work process executes a roll out and may have to wait for the end of the RFC in the roll area, even if the dialog step is not yet completed. In the roll area, RFC server programs can also wait for other RFCs sent to them.   |
| Average time per logical DB call | The average response time for all commands sent to the database system (in milliseconds). The time depends on the CPU capacity of the database server, the network, the buffering, and on the input/output capabilities of the database server. Access times for buffered tables are many magnitudes faster and are not considered in the measurement. |

- 8 Click **Add** to place the selected measurement on the resource list. Add all the desired resources to the list, and click **Close**.
- 9 Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the SAP Portal Monitor

To use the SAP Portal monitor, you must first configure the SAP Portal monitor on the SiteScope machine, and then select the counters you want the SAP Portal monitor to measure. You select these counters using the Console's SAP Portal dialog box.

### Configuring the SAP Portal monitor on the SiteScope machine:

- 1 Restart SiteScope after installing the monitor add-in, and verify connectivity to the SAP Portal SWSE page by opening the following url from the machine where sitescope is installed:

`http://<your_SAP_Portal_server>/sapportal`

- 2 In the SiteScope Panel, click **Create Group**, enter a name for the group, and click **Add**.
- 3 In the **Add to Group** section, click **Monitor** and select **SAP Portal Monitor** from the list of monitors.
- 4 In the server field, enter the url to the SWE stats page you want to monitor. For example:

`http://<your_SAP_Portal_server>/sapportal`

- 5 Enter the name of the application you want to monitor in the Application field.
- 6 Enter the user name and password (if applicable).
- 7 Click **Choose Counters** and select your desired counters.
- 8 Click **Choose Counters** again, and then click **Add Monitor**.

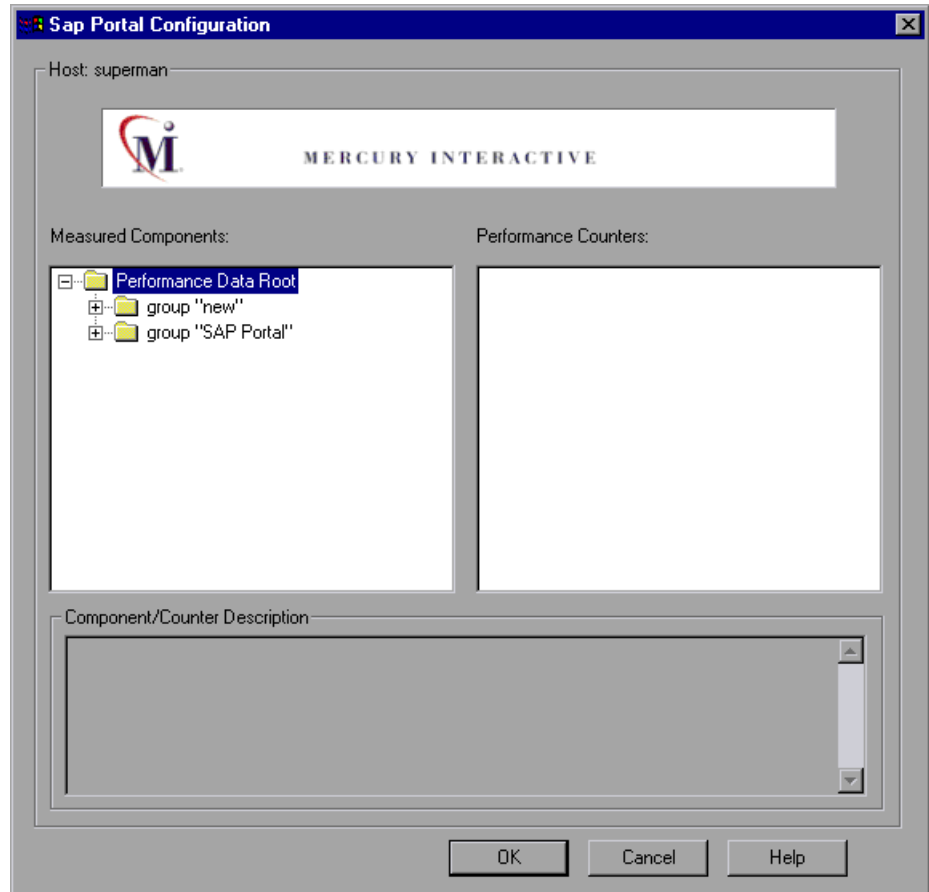
### To configure the SAP Portal monitor on the Console machine:



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 From the Server list, choose the server running Siebel.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select SAP Portal (in the ERP/CRM Server Resource Graphs category) and then click **Add**.



The SAP Portal Monitor Configuration dialog box is displayed.



- 5** In the Measured Components pane, locate the SAP Portal measurement that you are monitoring and click it. The performance counters that SAP Portal is monitoring on the selected component are displayed in the Performance Counters pane.

The following table shows the default counters that can be measured:

| Measurement   | Description  |
|---|--|
| Accumulated Amount of Outbound Data (bytes)           | The accumulated amount of outbound data, measured in bytes.              |
| Time for all Requests (ms)                            | The total time, in milliseconds, taken for processing all requests.      |
| Average Amount of Outbound Data per Request (bytes)   | The average amount of outbound data per request, measured in bytes.      |
| Average Number of Component Calls per Request (bytes) | The average number of component calls per request, measured in bytes.    |
| Average Time of a Request (ms)                        | The average amount of time, in milliseconds, taken to process a request. |
| Number of Calls with Outbound Data                    | The total number of calls with outbound data.                            |
| Number of Component Calls for all Requests            | The total number of component calls for all requests.                    |
| Number of Requests since First Request                | The total number of requests since the first request was made.           |
| Requests per Second                                   | The number of requests made per second.                                  |
| Time Stamp of First Request                           | The time stamp of the first request.                                     |

- 6 Check the required performance counters in the Performance Counters pane. When you have selected the performance counters for the SAP Portal measurements you are monitoring, click **OK** to close the SAP Portal Configuration dialog box. The **Select Measurements to Monitor** dialog box appears with the selected SAP Portal measurements in the **Selected Measurements** pane.
- 7 Click **OK** in the **Select Measurements to Monitor** dialog box, and click **OK** in the **Monitors Configuration** dialog box, to activate the SAP Portal monitor.

---

**Note:** For troubleshooting tips and limitations, see “Troubleshooting Server Resource Monitors,” on page 541.

---

## Configuring the Siebel Monitor

To use the Siebel monitor, you must first configure the Siebel monitor on the SiteScope machine, and then select the counters you want the Siebel monitor to measure. You select these counters using the Console's Siebel dialog box.

### Configuring the Siebel monitor on the SiteScope machine:

- 1** Restart SiteScope after installing the monitor add-in, and verify connectivity to the Siebel SWSE page by opening the following url from the machine where sitescope is installed:

`http://<your_siebel_server>/callcenter/_stats.swe`

- 2** In the SiteScope Panel, click **Create Group**, enter a name for the group, and click **Add**.
- 3** In the **Add to Group** section, click **Monitor** and select **Siebel Web Server** from the list of monitors.
- 4** In the server field, enter the url to the SWE stats page you want to monitor. For example:

`http://<your_siebel_server>/callcenter/_stats.swe`

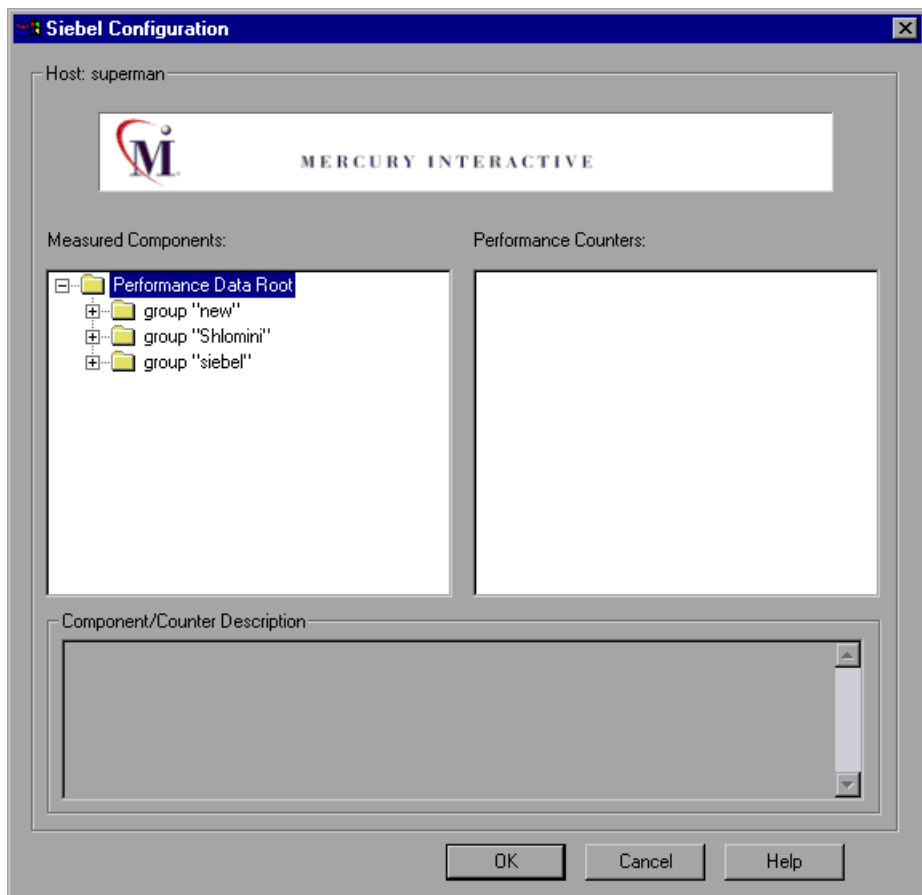
- 5** Enter the name of the application you want to monitor in the Application field. For example: Callcenter.
- 6** Enter the user name and password (if applicable).
- 7** Click **Choose Counters** and select your desired counters.
- 8** Click **Choose Counters** again, and then click **Add Monitor**.

**To configure the Siebel monitor:**

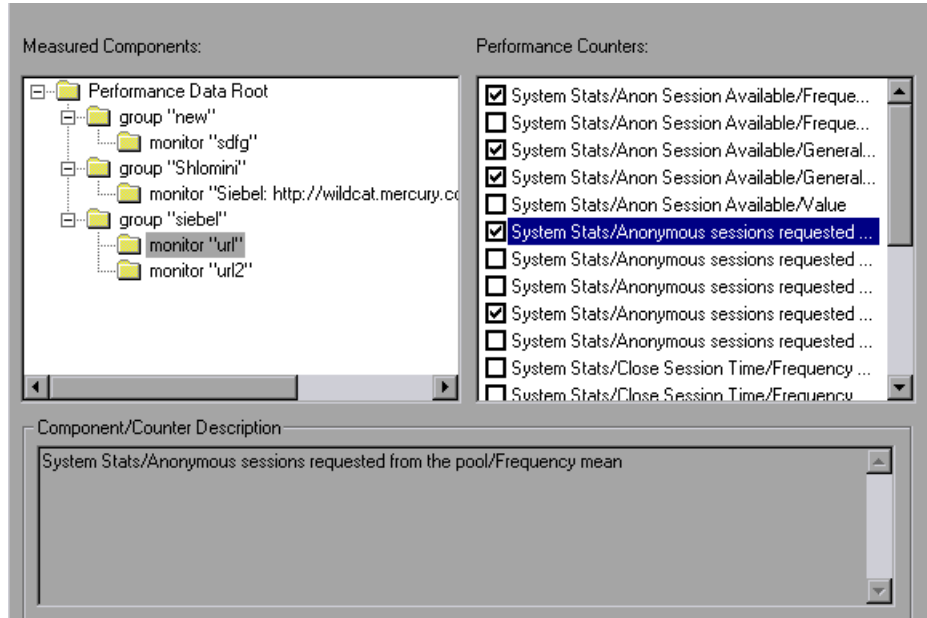


- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 From the Server list, choose the server running Siebel.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select Siebel (in the ERP/CRM Server Resource Graphs category) and then click **Add**.

The Siebel Monitor Configuration dialog box is displayed.



- In the Measured Components pane, locate the Siebel measurement that you are monitoring and click it. The performance counters that Siebel is monitoring on the selected component are displayed in the Performance Counters pane.



- Check the required performance counters in the Siebel Monitor Configuration window's right pane.

The following table shows the default counters that can be measured:

| Measurement                                       | Description   |
|---|---|
| <b>Anonymous sessions requested from the pool</b> | The number of anonymous sessions requested from the pool. |
| <b>Open Session Time</b>                          | The time users experience logging on to the system.       |
| <b>Anon Session Removed</b>                       | The number of anonymous sessions removed from the pool.   |

| Measurement                            | Description   |
|--|---|
| Anon Session Available                 | The number of anonymous sessions available in the pool. |
| Anonymous sessions returns to the pool | The number of anonymous sessions returned to the pool.  |
| Response Time                          | The time taken to respond to a user request.            |
| Close Session Time                     | The time users experience logging off the system.       |
| Request Time                           | The time taken to process the user request.             |

- 7 When you have selected the performance counters for the Siebel measurements you are monitoring, click **OK** to close the Siebel Monitor Configuration dialog box. The **Select Measurements to Monitor** dialog box appears with the selected Siebel measurements in the **Selected Measurements** pane.
- 8 Click **OK** in the **Select Measurements to Monitor** dialog box, and click **OK** in the **Monitors Configuration** dialog box, to activate the Siebel monitor.

---

**Note:** For troubleshooting tips and limitations, see “Troubleshooting Server Resource Monitors,” on page 541.

---

## Configuring the Siebel Server Manager Monitor

To monitor the Siebel Server Manager performance, you must first install the Siebel Server Manager client on the SiteScope machine. Then select the counters you want the Siebel Server Manager monitor to measure. You select these counters using the Console's Siebel Server Manager dialog box.

### Configuring the Siebel Server Manager client:

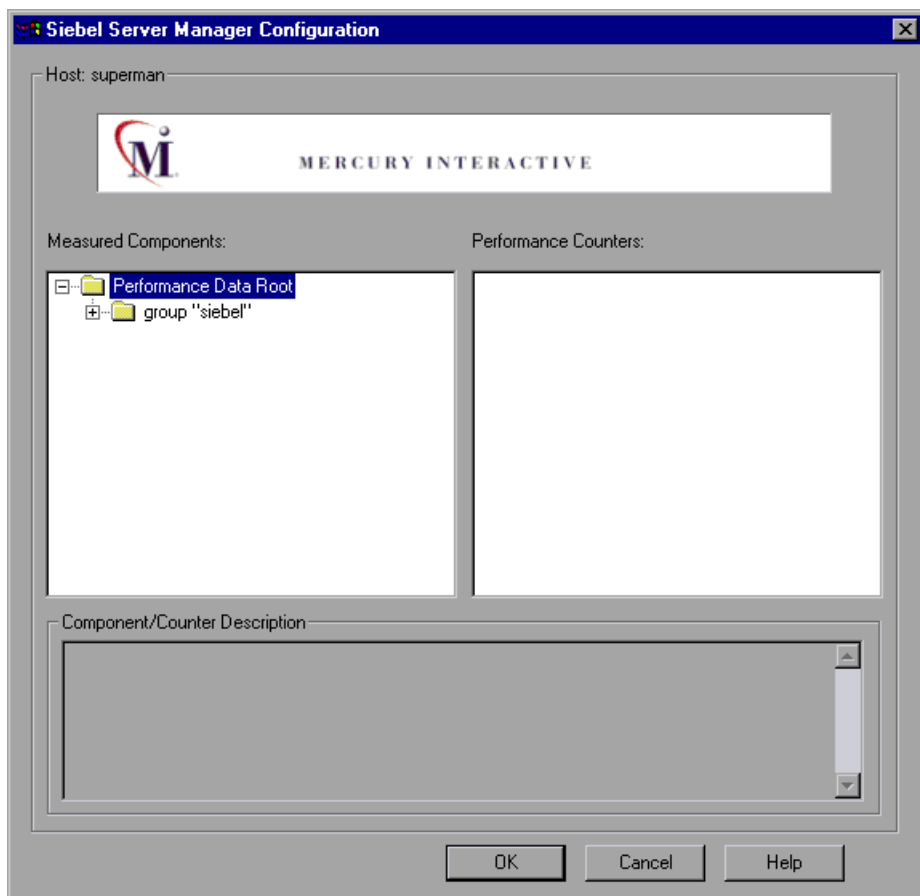
- 1** Restart SiteScope and verify connectivity to the Siebel SWE page by opening the following url from the machine where sitescope is installed:  
`http://<your_siebel_server>/<application_name>/_stats.swe`
- 2** In the SiteScope Panel, click **Create Group**, enter a name for the group, and click **Add**.
- 3** In the **Add to Group** section, click **Monitor** and select **Siebel Server Manager** from the list of monitors.
- 4** Enter the name of the Siebel Server in the Application server field, the Enterprise Server in the Enterprise server field, and the Gateway Server in the Gateway server field.
- 5** Enter the path to Siebel Server Manager in the Path to Script field.
- 6** Enter the user name and password (if applicable).
- 7** Click **Choose Counters** and select your desired counters.
- 8** Click **Choose Counters** again, and then click **Add Monitor**.

### To configure the Siebel Server Manager monitor on the Console machine:



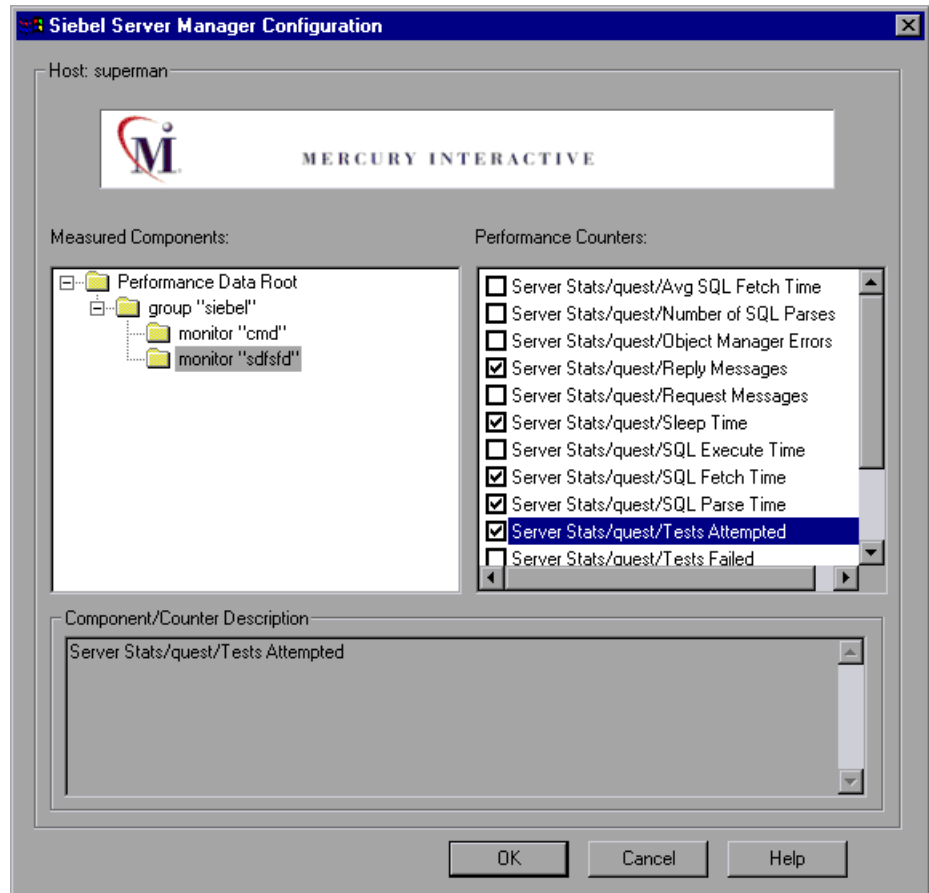
- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** From the Server list, choose the server running Siebel.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Siebel Server Manager (in the ERP/CRM Server Resource Graphs category) and then click **Add**.

The Siebel Server Manager Monitor Configuration dialog box is displayed.





- In the Measured Components pane, locate the Siebel Server Manager measurement that you are monitoring and click it. The performance counters that Siebel Server Manager is monitoring on the selected component are displayed in the Performance Counters pane.



- Check the required performance counters in the Siebel Server Manager Monitor Configuration window's right pane.

The following performance counters are available for the Siebel Server Manager monitor:

| <b>Measurement</b>                  | <b>Description</b>   |
|-------------------------------------|--|
| <b>Average Connect Time</b>         | The average connection time.   |
| <b>Average Reply Size</b>           | The average size of a user reply.  |
| <b>Average Request Size</b>         | The average size of a user request.  |
| <b>Average Requests Per Session</b> | The average number of user requests per session.                             |
| <b>Average Response Time</b>        | The average amount of time that it takes the server to respond to a request. |
| <b>Average Think Time</b>           | The average amount of think time taken to respond to a request.              |
| <b>Avg SQL Execute Time</b>         | The average SQL execute time.  |
| <b>Avg SQL Fetch Time</b>           | The average SQL fetch time.  |
| <b>Avg SQL Parse Time</b>           | The average SQL parse time.  |
| <b>CPU Time</b>                     | The CPU time used in the work process.                                       |
| <b>Elapsed Time</b>                 | The total amount of elapsed time.  |
| <b>Num of DBConn Retries</b>        | The number of database connection retries.                                   |
| <b>Num of DLRbk Retries</b>         | The number of DLRbk retries.   |
| <b>Num of Exhausted Retries</b>     | The total number of retries that expired.                                    |
| <b>Number of SQL Executes</b>       | The total number of SQL executes.  |
| <b>Number of SQL Fetches</b>        | The total number of SQL fetches.   |
| <b>Number of SQL Parses</b>         | The total number of SQL parses.  |
| <b>Number of Sleeps</b>             | The number of sleeps.  |

| Measurement           | Description                                |
|-----------------------|--|
| Object Manager Errors | The total number of object manager errors. |
| Reply Messages        | The total number of reply messages.        |
| Request Messages      | The total number of request messages.      |
| SQL Execute Time      | The total SQL execute time.                |
| SQL Fetch Time        | The total SQL fetch time.                  |
| SQL Parse Time        | The total SQL parse time.                  |
| Sleep Time            | The total sleep time.                      |
| Tests Attempted       | The number of tests attempted.             |
| Tests Failed          | The number of tests that failed.           |
| Tests Successful      | The number of tests that were successful.  |
| Total Reply Size      | The total reply size, measured in bytes.   |
| Total Request Size    | The total request size, measured in bytes. |
| Total Response Time   | The total response time.                   |
| Total Tasks           | The total number of tasks.                 |
| Total Think Time      | The total think time.                      |

- 7 When you have selected the performance counters for the Siebel measurements you are monitoring, click **OK** to close the Siebel Server Monitor Configuration dialog box. The **Select Measurements to Monitor** dialog box appears with the selected Siebel measurements in the **Selected Measurements** pane.
- 8 Click **OK** in the **Select Measurements to Monitor** dialog box, and click **OK** in the **Monitors Configuration** dialog box, to activate the Siebel Server Manager monitor.

---

**Note:** For troubleshooting tips and limitations, see “Troubleshooting Server Resource Monitors,” on page 541.

---

# 26

---

## Java Performance Monitoring

During a session step run, you can monitor the resource usage of Enterprise Java Bean (EJB) objects and Java-based applications, using the Java performance monitors:

This chapter describes:

- ▶ EJB Performance Monitoring
- ▶ JProbe Performance Monitoring
- ▶ Sitraka JMonitor Performance Monitoring

### About Java Performance Monitoring

The Java performance monitors provide you with performance information for Enterprise Java Bean (EJB) objects and Java-based applications, using the EJB, JProbe, and Sitraka JMonitor during session step execution. In order to obtain this data, you need to activate the Java performance monitors before executing the session step, and indicate which statistics and measurements you want to monitor.

## EJB Performance Monitoring

### Support Matrix:

| Application Server | Version                 | Platform              |
|--------------------|-------------------------|-----------------------|
| WebLogic           | 4.x; 5.1; 6.0; 6.1; 7.0 | Windows; Solaris; AIX |
| WebSphere          | 3.x; 4.x                | Windows; Solaris; AIX |
| Oracle 9i          | 1.0.2.2                 | Windows; Solaris; AIX |

You can monitor Enterprise Java Bean (EJB) objects on a WebLogic, WebSphere, or Oracle 9iAS application server during a session step run using the EJB performance monitor. In order to monitor EJB objects, you must first install the EJB monitor, run the monitor detector, and activate the EJB monitor on the application server machine. You then configure the EJB monitor on the client machine by selecting the counters you want the monitor to measure.

---

**Note:** The server side installation contains new EJBDetector support files for generating EJB Vuser scripts. For more information on the EJBDetector, refer to the *ProTune Virtual User Generator User's Guide*.

---

### Installing the EJB Monitor and Running the Monitor Detector

Before EJB objects can be monitored, you must install the EJB monitor support files, and verify that you have a valid JDK environment on the application server machine. You then prepare the EJB monitor for monitoring by running the monitor detector from the batch file, or from the command line.

**To install the EJB monitor support files:**

Create a home directory for the Mercury Interactive EJB support files—for example, `MERC_MONITOR_HOME`—and unzip the `<ProTune CD>add-ins\Monitors\J2EE\Windows\jmonitor_<platform>.jar` file into that directory.

On UNIX platforms, use the `jar` utility to extract the installation jar:

Change to the `MERC_MONITOR_HOME` directory and type the following command:

```
jar -xvf <path to your jmonitor_<platform>.jar>
```

**To run the monitor detector from the batch file:**

- 1 Open the `env.cmd` (NT) or `env.sh` (UNIX) file and set the following variables:
  - JAVA\_HOME** Specify the root directory of the JDK installation.
  - APP\_SERVER\_DRIVE** Specify the drive on which the application server is installed (for NT only).
  - DETECTOR\_INS\_DIR** Specify the root directory of the Detector installation.
  - APP\_SERVER\_ROOT** Follow these guidelines:
    - BEA WebLogic Servers 4.x and 5.x:** Specify the application server root directory.
    - BEA WebLogic Servers 6.x and 7.x:** Specify the full path of the domain folder.
    - WebSphere Servers 3.x and 4.0:** Specify the application server root directory.
    - Oracle OC4J:** Specify the application server root directory.
    - Sun J2EE Server:** Specify the full path to the deployable `.ear` file or directory containing a number of `.ear` files.
  - EJB\_DIR\_LIST (optional)** Specify a list of directories/files, separated by ‘;’ and containing deployable `.ear/.jar` files, and any additional classes directory or `.jar` files or used by your EJBs under test.

- 2 Run the *Mon\_Detector.cmd* (NT) or *Mon\_Detector.sh* (UNIX) batch file to collect information about the EJBs deployed. Running the monitor detector generates the following three files in the `<MERC_MONITOR_HOME>\dat` directory: *ejb\_monitor.hooks*; *cjhook.ini*; and *regmon.properties*. These files contain information about the EJBs detected on the application server.

---

**Note:** You must run the monitor detector each time you add, change, or delete EJBs on the application server.

---

**To run the monitor detector from a command line:**

- 1 Add `<MERC_MONITOR_HOME>\classes`, `<MERC_MONITOR_HOME>\dat`, and the `<MERC_MONITOR_HOME>\classes\xerces.jar` file to the CLASSPATH environment variable.
- 2 Use the `java MonDetect <search root dir>` command line to collect information about the EJBs deployed.

*<search root dir>* Specify one or more directories or files in which to search for EJBs (separated by semicolons). Follow these guidelines:

- BEA WebLogic Servers 4.x and 5.x:** Specify the application server root directory.
- BEA WebLogic Servers 6.x and 7.x:** Specify the full path of the domain folder followed by the root directory.
- WebSphere Servers 3.x and 4.0:** Specify the application server root directory.
- Oracle OC4J:** Specify the application server root directory.
- Sun J2EE Server:** Specify the full path to the deployable *.ear* file or directory containing a number of *.ear* files.

Note that you can also specify a search list of directories and/or files to search. If unspecified, the CLASSPATH will be searched.



Running the monitor detector generates the following three files in the `<MERC_MONITOR_HOME>\dat` directory: `ejb_monitor.hooks`; `cjhook.ini`; and `regmon.properties`. These files contain information about the EJBs detected on the application server.

---

**Note:** You must run the monitor detector each time you add, change, or delete EJBs on the application server.

---

### Configuring the EJB Monitor on the Application Server

After you have installed Mercury Interactive's EJB monitor support files on your WebLogic, WebSphere, or Oracle 9iAS machine, you must configure the application server to run with EJB monitor support.

---

**Note:** It is important to set the environment variables in the order in which they appear below.

---

#### WebLogic Server

The WebLogic 4.x-5.x server, WebLogic 6.x server, and the WebLogic 7.x server must be configured differently.

##### To configure the WebLogic 4.x-5.x server:

- 1 Copy the `<WebLogic Home>\startWeblogic.cmd` file into `<WebLogic Home>\startWeblogicMercury.cmd` so that the file is backed up.
- 2 Open the `<WebLogic Home>\startWeblogicMercury.cmd` file.
- 3 In the 'runWebLogicJava' section of the file, after the `WEBLOGIC_CLASSPATH` environment settings, set the following environment variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<EJB Monitor Home Directory>
set CLASSPATH=%MERC_MONITOR_HOME%\dat
```

```
set JAVA_CLASSPATH=%MERC_MONITOR_HOME%\dat;%MERC_MONITOR_HOME%\classes;%MERC_MONITOR_HOME%\classes\xerces.jar;%JAVA_CLASSPATH%
```

```
set PATH=%PATH%;%MERC_MONITOR_HOME%\bin
```

For UNIX platforms:

```
MERC_MONITOR_HOME <EJB Monitor Home Directory>
```

```
CLASSPATH ${MERC_MONITOR_HOME}/dat
```

```
JAVA_CLASSPATH ${MERC_MONITOR_HOME}/dat:${MERC_MONITOR_HOME}/classes:${MERC_MONITOR_HOME}/classes/xerces.jar:${JAVA_CLASSPATH}
```

```
LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:${MERC_MONITOR_HOME}/bin
```

```
export CLASSPATH
```

```
export LD_LIBRARY_PATH
```

```
export JAVA_CLASSPATH
```

---

**Note:** For IBM AIX platform replace LD\_LIBRARY\_PATH with LIBPATH. Replace <EJB Monitor Home Directory> with the EJB monitor installation root directory. Note that on UNIX platforms you may have to export the library path variables.

---

- 4** In the same section of the file, add a parameter to the command line:

```
-Xrunjdkhook.
```

For example on Windows platforms:

```
%JAVA_HOME%\bin\java -ms64m -mx64m -Xrunjdkhook -classpath  
%JAVA_CLASSPATH% -Dweblogic.class.path=%WEBLOGIC_CLASSPATH%  
-Dweblogic.home=. -Djava.security.manager  
-Djava.security.policy==.\weblogic.policy weblogic.Server
```

---

**Note:** For Solaris installation only.

If you are using JDK 1.2.x add a parameter to the command line:

```
-Dweblogic.classloader.preprocessor=com.mercuryinteractive.aim.  
MercuryWL5Preprocessor
```

for example, on Windows platforms:

```
%JAVA_HOME%\bin\java -ms64m -mx64m -classpath %JAVA_CLASSPATH%  
-Dweblogic.classloader.preprocessor=com.mercuryinteractive.aim.  
MercuryWL5Preprocessor  
-Dweblogic.class.path=%WEBLOGIC_CLASSPATH%  
-Dweblogic.home=. -Djava.security.manager  
-Djava.security.policy==.\weblogic.policy weblogic.Server
```

---

**5** Run the `<WebLogic Home>\startWeblogicMercury.cmd` file.

**To configure the WebLogic 6.x server:**

- 1** Copy the `<WebLogic Home>\config\<domain name>\startWeblogic.cmd` file into `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` so that the file is backed up.
- 2** Open the `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` file.
- 3** In the 'runWebLogic' section of the file, set the following environment variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<your MERC_MONITOR_HOME directory>  
set CLASSPATH=%CLASSPATH%;%MERC_MONITOR_HOME%\dat;%  
MERC_MONITOR_HOME%\classes;%MERC_MONITOR_HOME%\classes\  
xerces.jar  
set PATH=%PATH%;%MERC_MONITOR_HOME%\bin
```

For UNIX platforms:

```
MERC_MONITOR_HOME <EJB Monitor Home Directory>
```

```
CLASSPATH ${JAVA_CLASSPATH}:${MERC_MONITOR_HOME}/dat:${MERC_MONITOR_HOME}/classes:${MERC_MONITOR_HOME}/classes/xerces.jar
```

```
LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:${MERC_MONITOR_HOME}/bin
```

```
export CLASSPATH
```

```
export LD_LIBRARY_PATH
```

---

**Note:** For IBM AIX platform replace LD\_LIBRARY\_PATH with LIBPATH. Replace <EJB Monitor Home Directory> with the EJB monitor installation root directory. Note that on UNIX platforms you may have to export the library path variables.

---

- 4 In the same section of the file add a parameter to the command line:

```
-Xrunjdkhook.
```

for example, on Windows platforms:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -Xrunjdkhook -classpath  
%CLASSPATH% -Dweblogic.Domain=mydomain  
-Dweblogic.Name=myserver "-Dbea.home=f:\bea"  
"-Djava.security.policy==f:\bea\wlserver6.0\lib\weblogic.policy"  
-Dweblogic.management.password=%WLS_PW% weblogic.Server
```

- 5 Run the <WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd file.

**To configure the WebLogic 7.x server:**

- 1** Copy the `<WebLogic Home>\server\bin\startwls.cmd` file into `<WebLogic Home>\server\bin\startwlsMercury.cmd` so that the file is backed up.
- 2** Open the `<WebLogic Home>\server\bin\startwlsMercury.cmd` file.
- 3** In the 'runWebLogic' section of the file, set the following environment variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<your MERC_MONITOR_HOME directory>
set
CLASSPATH=%CLASSPATH%;%MERC_MONITOR_HOME%\dat;%MERC_M
ONITOR_HOME%\classes;%MERC_MONITOR_HOME%\classes\xerces.jar
set PATH=%PATH%;%MERC_MONITOR_HOME%\bin
```

For UNIX platforms:

```
MERC_MONITOR_HOME <EJB Monitor Home Directory>
CLASSPATH=$CLASSPATH:$MERC_MONITOR_HOME/dat:$MERC_
MONITOR_HOME/classes:$MERC_MONITOR_HOME/classes/xerces.jar
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$MERC_MONITOR_HOME/bin
export CLASSPATH
export LD_LIBRARY_PATH
```

---

**Note:** For IBM AIX platform replace LD\_LIBRARY\_PATH with LIBPATH. Replace `<EJB Monitor Home Directory>` with the EJB monitor installation root directory. Note that on UNIX platforms you may have to export the library path variables.

---

- 4 In the same section of the file add a parameter to the command line:

-Xrunjdkhook.

for example, on Windows platforms:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -Xrunjdkhook -classpath  
%CLASSPATH% -Dweblogic.Domain=mydomain  
-Dweblogic.Name=myserver "-Dbea.home=f:\bea"  
"-Djava.security.policy==f:\bea\wls\server6.0\lib\weblogic.policy"  
-Dweblogic.management.password=%WLS_PW% weblogic.Server
```

- 5 Copy the `<domain name>\startWeblogic.cmd` file into `<domain name>\startWeblogicMercury.cmd` so that the file is backed up.
- 6 Open the `<domain name>\startWeblogicMercury.cmd` file.
- 7 Find the call to the weblogic server. For example, call:  
D:\bea\weblogic700\server\bin\startWLS.cmd
- 8 Change the call from `startWLS.cmd` to `startWLSMercury.cmd`, and save the file.
- 9 Run the `<domain name>\startWeblogicMercury.cmd` file.

### WebSphere Server - Versions 3.0 and 3.5

By default, the WebSphere 3.x application server runs as an automatic service, upon machine startup. Since Mercury Interactive does not currently support ProTune EJB monitoring on a WebSphere server run as an automatic service, you must change the default WebSphere server startup to *manual*.

**To change the default WebSphere 3.x server startup:**

- 1 Select **Start > Settings > Control Panel**.
- 2 Double-click **Services**.
- 3 Select **IBM WS AdminServer**, and click the **Stop** button.
- 4 Double click **IBM WS AdminServer**, and select the **Manual** Startup Type.
- 5 Click **OK** to save your settings and close the dialog box.

You can now start the WebSphere Server from `<WebSphere Home>\AppServer\bin\debug\adminserver.bat`, instead of using the automatic service.

**To add ProTune EJB monitor support to the WebSphere 3.x server:**

- 1** Make a backup copy of the  
<WebSphere Home>\AppServer\bin\debug\adminserver.bat file.
- 2** Open the <WebSphere Home>\AppServer\bin\debug\adminserver.bat file.
- 3** Add the following environment variables at the end of the 'SET\_CP' section:

For Windows platforms:

```
set CLASSPATH=<MERC_MONITOR_HOME>\dat;<MERC_MONITOR_HOME>\classes;<MERC_MONITOR_HOME>\classes\xerces.jar;
%CLASSPATH%

set PATH=%PATH%;<MERC_MONITOR_HOME>\bin
```

For UNIX platforms:

```
CLASSPATH ${MERC_MONITOR_HOME}/dat:${MERC_MONITOR_HOME}/
classes:${MERC_MONITOR_HOME}/classes/xerces.jar:${CLASSPATH}

LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:${MERC_MONITOR_HOME}/bin

export CLASSPATH

export LD_LIBRARY_PATH
```

---

**Note:** For IBM AIX platform replace LD\_LIBRARY\_PATH with LIBPATH. Replace <EJB Monitor Home Directory> with the EJB monitor installation root directory. Note that on UNIX platforms you may have to export the library path variables.

---

**Note:** For Solaris installation only.

If you are working with JRE1.2.x, you must download the patch file, PQ46831.jar, from IBM's Web site or FTP site:

<http://www-3.ibm.com/software/webservers/appserv/efix-archive.html>

<ftp://ftp.software.ibm.com/software/websphere/appserv/support/fixes/pq46831/>

Make sure to download the version that corresponds to your server version.

Add the patch file to the classpath:

```
setenv CLASSPATH PQ46831.jar:${CLASSPATH}
```

---

- 4** Run the *adminserver.bat* file.
- 5** Open the WebSphere Advanced Administrative Console, and select **View > Topology**.
- 6** Expand the WebSphere Administrative Domain tree by selecting **<server machine name> > Default Server**.
- 7** Select the **General** tab in the Application Server:Default Server window.
- 8** Type `-Xrunjdkhook` in the command line Arguments box, and click **Apply**.

If you are working with a WebSphere 3.0 Server with JDK1.1.7 IBM, double-click on **Environment**. Type `_CLASSLOAD_HOOK` in the Variable Name box, and `jdkhook` in the Value box. Click the **Add**, **OK**, and **Apply** buttons.

---

**Note:** For Solaris installation only.

If you are working with a WebSphere 3.5 Server with J2RE1.2.x, in the Command Line Arguments box, type the following and click **Apply**:

```
-Dcom.ibm.ejs.sm.server.ServiceInitializer=com.ibm.ejs.sm.server.WilyInitializer
```

```
-Dcom.ibm.websphere.introscope.implClass=com.mercuryinteractive.aim.
```

```
MercuryWASPreprocessor
```

---

- 9** Close the WebSphere Advanced Administrative Console.
- 10** Close and restart the *adminserver.bat* file.



### WebSphere Server - Version 4.0

You can start the WebSphere 4.0 server using the startServerBasic.bat file or the startServer.bat file.

#### To configure the WebSphere 4.0 server:

- 1** Ensure that the WebSphere Administrative Server is running, and start the Administrator Console.
- 2** In the WebSphere Administrative Domain tree, expand the Nodes, hostname, and Application Servers subtrees, and select the Default Server (or the Application Server you wish to use with JMonitor).
- 3** For Windows 2000/NT or Solaris, click the **General** tab, and add the following variables to the **Environment** box:

---

**Note:** Replace *<EJB Monitor Home Directory>* with the EJB monitor installation root directory.

---

For Windows 2000/NT:

name=PATH

value=*<EJB Monitor Home Directory>*\bin

For Solaris:

name=LD\_LIBRARY\_PATH

value=*<EJB Monitor Home Directory>*/bin

Click **OK** to close the **Environment Editor** dialog box.

For AIX:

If the LIBPATH environment variable has been changed, you need to link the EJB monitor libraries to the /usr/lib directory.

Add the following command:

```
#ln -s <EJB Monitor Home Directory>/bin/libcjhooke_mon.so  
/usr/lib/libcjhooke_mon.so  
  
#ln -s <EJB Monitor Home Directory>/bin/libconfig.so /usr/lib/libconfig.so  
  
#ln -s <EJB Monitor Home Directory>/bin/libjdkhook.so /usr/lib/libjdkhook.so  
  
#ln -s <EJB Monitor Home Directory>/bin/libmllib_ds.so /usr/lib/libcjhooke_mon.so  
  
#ln -s <EJB Monitor Home Directory>/bin/libmosifs.so /usr/lib/libmosifs.so  
  
#ln -s <EJB Monitor Home Directory>/bin/libthrdutil.so /usr/lib/libthrdutil.so
```

---

**Note:** You will likely require root permissions in order to create the link. Alternatively, you can place the link in WebSphere's /bin directory (usually /usr/WebSphere/AppServer/bin).

---

- 4 Click the **JVM Settings** tab in the WebSphere Administrative Console, and add the following values to the classpath:

---

**Note:** Replace <EJB Monitor Home Directory> with the EJB monitor installation root directory.

---

For Windows 2000/NT:

```
<EJB Monitor Home Directory>\dat  
<EJB Monitor Home Directory>\classes  
<EJB Monitor Home Directory>\classes\xerces.jar
```

For Solaris or AIX:

```
<EJB Monitor Home Directory>/dat  
<EJB Monitor Home Directory>/classes  
<EJB Monitor Home Directory>/classes/xerces.jar
```

---

**Note:** For Solaris installation only.

If you are working with JRE1.2.x, you must download the patch file, PQ46831.jar, from IBM's Web site or FTP site:

<http://www-3.ibm.com/software/webservers/appserv/efix-archive.html>

<ftp://ftp.software.ibm.com/software/websphere/appserv/support/fixes/pq46831/>

Make sure to download the version that corresponds to your server version. Add the following value to the classpath:

`<EJB Monitor Home Directory>/classes/PQ46831.jar`

---

- 5 Click the **Advanced JVM Settings** button. In the Command line arguments field, add the following value for Windows 2000/NT, Solaris, and AIX:

`-Xrunjdkhook`

---

**Note:** For Solaris installation only.

If you are working with JRE1.2.x, instead of `-Xrunjdkhook` add the following value:

`-Dcom.ibm.ejs.sm.server.ServiceInitializer=com.ibm.ejs.sm.server.WilyInitializer`

`-Dcom.ibm.websphere.introscope.implClass=com.mercuryinteractive.aim.MercuryWASPreprocessor`

---

- 6 Click the **OK** and **Apply** buttons to save the changes for the Application server. You can now start and stop your WebSphere server using the ProTune EJB Monitor.

### Oracle 9iAS Server

Once you have configured the support files and set up the JDK environment on the Oracle 9iAS application server, run the *oc4jMonitor.cmd* file on an NT machine, or the *oc4jMonitor.sh* file on a UNIX machine. The application server starts running with EJB monitor support.

### Configuring the EJB Monitor on the Client Machine

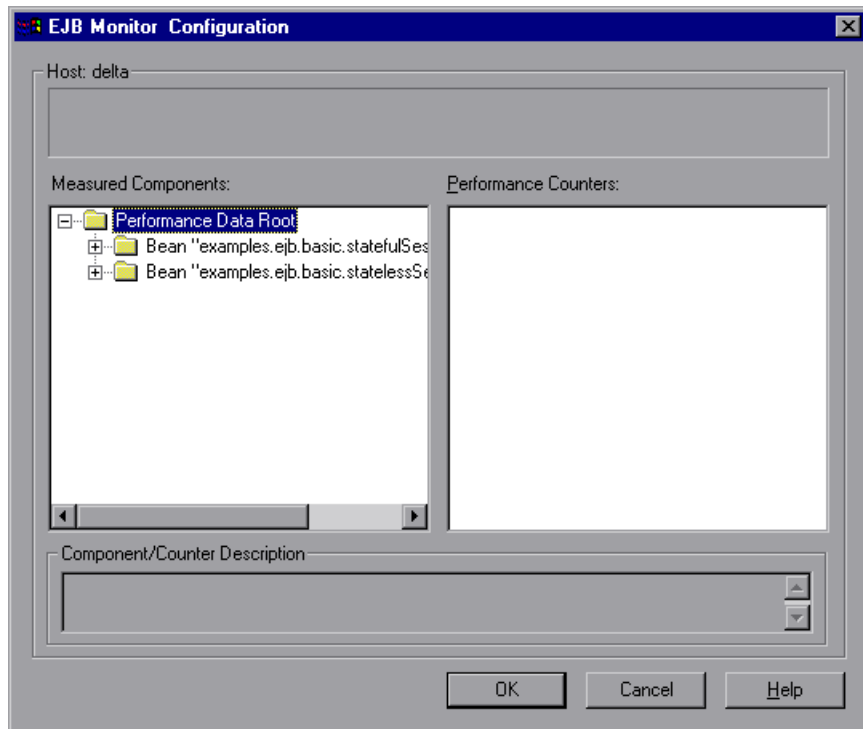
To monitor EJB performance, you must select the counters you want the EJB monitor to measure. You select these counters using the Console's EJB Monitor Configuration dialog box.

#### To configure the EJB monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select EJB (in the Java Technology category) and then click **Add**.

The EJB Monitor Configuration dialog box opens, displaying the available EJBs.



- Expand the Measured Components tree and select the methods and counters you want to monitor. The following counters can be monitored for each method:

| Measurement             | Description  |
|-------------------------|--|
| Average Response Time   | The average response time, in milliseconds, of the EJB object being monitored. |
| Method Calls per Second | The number of EJB object method calls per second.                              |

- Click **OK** in the EJB Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the EJB monitor.

## JProbe Performance Monitoring

You can monitor Java-based applications during a session step run using the JProbe Java profiling tool. In order to monitor Java-based applications, you must first configure the JProbe tool to monitor the specified application. You then use the Console to select the counters you want the JProbe Java profiling tool to measure.

---

**Note:** You must run the JPlauncher before opening the Console.

---

### Configuring the JProbe Tool to Monitor an Application

In order to monitor Java-based applications—for example, a WebLogic server—you must first configure the JProbe tool to monitor the specified application.

**To configure the JProbe tool:**

- 1 Launch JProbe Profiler 3.0.**
- 2 Select Program > Open JProbe Launch Pad.** Select the **Program** tab, and enter the following settings:

**Target Server:** <server name>—for example, BEA WebLogic 5.1.0

**Server Home Directory:** for example, G:\Weblogic

**Working Directory:** for example, G:\Weblogic

**Classpath:** for example, %CLASSPATH%

- 3 Select the VM (Virtual Machine) tab,** and enter the following settings:

**Virtual Machine Type:** for example, Java 2

**VM Path:** for example, g:\JDK1.2\bin\java.exe

**VM Arguments:** for example, -ms64m -mx64m -Dweblogic.home= -Djava.security.manager -Djava.security.policy=. \Weblogic.policy

**Snapshot Directory:** for example, G:\TEMP

- 4 Select the **Attach** tab. In the JProbe Console section, select the option that corresponds to your setup (for example, choose **Local** if the JProbe tool is on the same machine as the Console). In the section relating to the port to be used, select **Use Default Port**.

---

**Note:** The Console usually uses port 4444 as the default port for the JProbe tool. To change the default port, edit the *<installation root>\dat\monitors\jprobe.cfg* file, or specify the server machine name in the Console's Add Machine dialog box as host:port.

---

- 5 Click the **Save As** button, and save the file in JProbe Launch Pad (JPL) format.
- 6 Close the JProbe Profiler.
- 7 In DOS, enter `cd <JProbe Profiler Dir>/jplauncher`. At the prompt, type the following:

```
jplauncher -jp_input=< .JPL file> -jp_socket="<Console machine>:<Port>"
```

The JPlauncher establishes communication with the Console machine. The Console receives data from the JProbe tool through the above default port.

### Configuring the JProbe Tool in the Console

To monitor a Java-based application, you must select the counters you want the JProbe tool to measure. You select these counters using the Console's JProbe dialog box.

#### To configure the JProbe tool:



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.
- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select JProbe (in the Java Technology category) and then click **Add**.

The following two measurements appear.

| Measurement             | Description   |
|-------------------------|---|
| Allocated Memory (heap) | The amount of allocated memory in the heap (in bytes) |
| Available Memory (heap) | The amount of available memory in the heap (in bytes) |

- 5 Click **OK** in the JProbe dialog box to activate the monitor.

## Sitraka JMonitor Performance Monitoring

### Support Matrix:

| Application Server | Version      | Platform              |
|--------------------|--------------|-----------------------|
| WebLogic           | 6.0; 6.1     | Windows; Solaris; AIX |
| WebSphere          | 4.0          | Windows; Solaris; AIX |
| Tomcat             | 3.2.3, 4.0.3 | Windows; Solaris; AIX |

### Installation Options

The Sitraka JMonitor can be configured and run together with the EJB Monitor, or as a standalone monitor.

If you want to install the Sitraka JMonitor with the EJB Monitor, see “Configuring the EJB Monitor and the Sitraka JMonitor on the Application Server,” on page 479.



## Installing the Sitraka JMonitor on the Application Server

To install the Sitraka JMonitor, create a home directory for the Sitraka JMonitor support files—for example, `MERC_MONITOR_HOME`—and unzip the installation file `<ProTune_Installation>\Add-ins\J2EE\Sitraka\jmonitor_<platform>.jar` file into that directory.

On UNIX platforms, use the `jar` utility to extract the installation jar.

Change to the Sitraka JMonitor installation directory and type the following command:

```
jar -xvf <path to your Sitraka JMonitor installation jar>
```

## Configuring the Sitraka JMonitor on the Application Server

After you have installed the Sitraka JMonitor support files on your WebLogic, WebSphere, or Tomcat server, you must configure the application server to run with Sitraka JMonitor support.

---

**Note:** It is important to set the environment variables in the order in which they appear below.

---

### To configure the WebLogic 6.0/6.1 server:

- 1 Make a backup copy of the WebLogic startup script.

For Windows 2000/NT, this will be:

```
<WebLogic Home>\config\<domain name>\startWebLogic.cmd
```

Name the new file `startWebLogicJMonitor.cmd`

For Solaris or AIX, this will be in:

```
$BEA_HOME/wlserver6.0/config/<domain name>/startWebLogic.sh
```

```
or $BEA_HOME/wlserver6.1/config/<domain name>/startWebLogic.sh
```

Name the new file `startWebLogicJMonitor.sh`

- 2 Open the *startWebLogicJMonitor.cmd* or *startWebLogicJMonitor.sh* file.
- 3 In the 'runWebLogic' section of the file (or just prior to the call to invoke the JVM), set the following environment variables:

---

**Note:** Replace any instances of *<install directory>* with the JMonitor installation directory. Note that on Unix platforms you may have to export the library path variables.

---

For Windows 2000/NT:

```
set JMONITOR_HOME=<install directory>

set CLASSPATH=%JMONITOR_HOME%\lib;%JMONITOR_HOME%\lib\
miniwebserver.jar;%JMONITOR_HOME%\lib\jlrutils.jar;%JMONITOR_HOME%\
lib\jmonitor.jar;%CLASSPATH%

set PATH=%PATH%;%JMONITOR_HOME%\bin\win32_ia32
```

For Solaris:

```
JMONITOR_HOME=<install directory>

CLASSPATH=$JMONITOR_HOME/lib:$JMONITOR_HOME/lib/miniwebserver.j
ar:$JMONITOR_HOME/lib/jlrutils.jar:$JMONITOR_HOME/lib/jmonitor.jar:$CLA
SSPATH

LD_LIBRARY_PATH=$JMONITOR_HOME/bin/solaris_sparc:$LD_LIBRARY_P
ATH

export LD_LIBRARY_PATH
```

For AIX:

```
JMONITOR_HOME=<install directory>

CLASSPATH=$JMONITOR_HOME/lib:$JMONITOR_HOME/lib/miniwebserver.j
ar:$JMONITOR_HOME/lib/jlrutils.jar:$JMONITOR_HOME/lib/jmonitor.jar:$CLA
SSPATH

LIBPATH=$JMONITOR_HOME/bin/aix_ppc:$LIBPATH

export LD_LIBRARY_PATH
```

- 4** In the same section of the file, modify the java command-line.

For Windows 2000/NT, Solaris, and AIX, add the following parameter:

```
-Xrunjmonitor:load_mws,proxy
```

The command line will look similar to the following (paths shown are for Windows, and are for a specific installation of WebLogic):

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m
-Xrunjmonitor:load_mws,proxy -classpath %CLASSPATH%
-Dweblogic.Domain=mydomain -Dweblogic.Name=myserver
"-Dbea.home=f:\bea"
"-Djava.security.policy==f:\bea\wlsrver6.0/lib/weblogic.policy"
-Dweblogic.management.password=%WLS_PW% weblogic.Server
```

- 5** Run the *startWebLogicJMonitor.cmd* or *startWebLogicJMonitor.sh* file to start the WebLogic instance with JMonitor enabled.

**To configure the WebSphere 4.0 server:**

- 1** Ensure that the WebSphere Administrative Server is running, and start the Administrator Console.
- 2** Expand the WebSphere Administrative Domain tree. Expand the Nodes, hostname, and Application Servers subtrees, and select the Default Server (or the Application Server you wish to use with JMonitor).
- 3** For Windows 2000/NT and Solaris, click the General tab, and add the following variables to the Environment Editor box:

---

**Note:** Replace any instances of *<install directory>* with the JMonitor installation directory.

---

For Windows 2000/NT:

```
name=PATH
```

```
value=<install directory>\bin\win32_ia32
```

For Solaris:

```
name=LD_LIBRARY_PATH
```

```
value=<install directory>/bin/solaris_sparc
```

Click **OK** to close the Environment Editor.

For AIX:

If the LIBPATH environment variable has been adjusted, you need to link the Sitraka JMonitor to the /usr/lib directory.

Add the following command:

```
#ln -s <install directory>/bin/aix_ppc/libjmonitor.so /usr/lib/libjmonitor.so
```

---

**Note:** you will likely require root permissions in order to create the link. Alternatively, the link can be placed in WebSphere's /bin directory (usually /usr/WebSphere/AppServer/bin).

---

- 4 Click the **JVM Settings** tab in the WebSphere Administrative Console, and add the following values to the classpath:

---

**Note:** Replace any instances of <install directory> with the JMonitor installation directory.

---

For Windows 2000/NT

```
<install directory>\lib
```

```
<install directory>\lib\jlrutils.jar
```

```
<install directory>\lib\miniwebserver.jar
```

```
<install directory>\lib\jmonitor.jar
```

For Solaris or AIX:

*<install directory>\lib*

*<install directory>\lib\jlrutils.jar*

*<install directory>\lib\miniwebserver.jar*

*<install directory>\lib\jmonitor.jar*

- 5** Click the **Advanced JVM Settings** button. In the Command line arguments field, add the following value for Windows 2000/NT, Solaris, and AIX:

*-Xrunjdkhook:jmonitor:load\_mws,proxy*

- 6** Click the **OK** and **Apply** buttons to save the changes for the Application server. You can now start and stop your WebSphere server using the Sitraka JMonitor.
- 7** To start your application server without using the Sitraka JMonitor, remove the changes made in step 5.

**To configure the Tomcat 3.2.3 server:**

- 1** Make a backup copy of the Tomcat startup script.

For Windows 2000/NT, this will be:

*<Tomcat Home>\bin\tomcat.bat*

Name the new file *tomcat\_jmonitor.bat*

For Solaris or AIX, this will be in:

*<Tomcat home>/bin/tomcat.sh*

Name the new file *tomcat\_jmonitor.sh*

- 2** Open the *tomcat\_jmonitor.bat* or *tomcat\_jmonitor.sh* file.

**3** Set the following environment variables:

For Windows 2000/NT, the following additions should be added prior to the line invoking the JVM (in the 'startServer' section of the batch file):

```
set JMONITOR_HOME=<install directory>

set CLASSPATH=%JMONITOR_HOME%\lib;%JMONITOR_HOME%\lib\
miniwebserver.jar;%JMONITOR_HOME%\lib\jlrutils.jar;%JMONITOR_HOME%\
lib\jmonitor.jar;%CLASSPATH%

set PATH=%PATH%;%JMONITOR_HOME%\bin\win32_ia32
```

For Solaris, the following lines should be added prior to invoking the JVM (just after the first export CLASSPATH found in the file):

```
JMONITOR_HOME=<install directory>

CLASSPATH=$JMONITOR_HOME/lib:$JMONITOR_HOME/lib/miniwebserver.
jar:$JMONITOR_HOME/lib/jlrutils.jar:$JMONITOR_HOME/lib/jmonitor.jar:
$CLASSPATH

export CLASSPATH

LD_LIBRARY_PATH=$JMONITOR_HOME/bin/solaris_sparc:$LD_LIBRARY_
PATH

export LD_LIBRARY_PATH
```

For AIX, the following lines should be added prior to invoking the JVM (just after the first export CLASSPATH found in the file):

```
JMONITOR_HOME=<install directory>

CLASSPATH=$JMONITOR_HOME/lib:$JMONITOR_HOME/lib/miniwebserver.j
ar:$JMONITOR_HOME/lib/jlrutils.jar:$JMONITOR_HOME/lib/jmonitor.jar:$CLA
SSPATH

export CLASSPATH

LIBPATH=$JMONITOR_HOME/bin/aix_ppc:$LD_LIBRARY_PATH

export LIBPATH
```

- 4** In the same section of the file, modify the java command-line.

For Windows 2000/NT, Solaris, and AIX, add the following parameter:

```
-Xrunjmonitor:load_mws,proxy
```

- 5** To start the Tomcat server with JMonitor, use the following command:  
tomcat\_jmonitor start.

**To configure the Tomcat 4.0.3 server:**

- 1** Make a backup copy of the Tomcat startup script.

For Windows 2000/NT, this will be:

```
<Catalina Home>\bin\catalina.bat
```

Name the new file catalina\_jmonitor.bat

For Solaris or AIX, this will be in:

```
<Tomcat home>/bin/catalina.sh
```

Name the new file catalina\_jmonitor.sh

- 2** Open the *tomcat\_jmonitor.bat* or *tomcat\_jmonitor.sh* file.

- 3** Set the following environment variables:

For Windows 2000/NT, the following additions should be added prior to the line invoking the JVM (in the 'doStart' section of the batch file):

```
set JMONITOR_HOME=<install directory>
```

```
set CLASSPATH=%JMONITOR_HOME%\lib;%JMONITOR_HOME%\lib\
miniwebserver.jar;%JMONITOR_HOME%\lib\jlrutils.jar;%JMONITOR_
HOME%\lib\jmonitor.jar;%CLASSPATH%
```

```
set PATH=%PATH%;%JMONITOR_HOME%\bin\win32_ja32
```

For Solaris, the following lines should be added prior to invoking the JVM (at the beginning of the 'Execute The Requested Command' section):

```
JMONITOR_HOME=<install directory>

CLASSPATH=$JMONITOR_HOME/lib:$JMONITOR_HOME/lib/miniwebserver.jar:$JMONITOR_HOME/lib/jlrutils.jar:$JMONITOR_HOME/lib/jmonitor.jar:$CLASSPATH

export CLASSPATH

LD_LIBRARY_PATH=$JMONITOR_HOME/bin/solaris_sparc:$LD_LIBRARY_PATH

export LD_LIBRARY_PATH
```

For AIX, the following lines should be added prior to invoking the JVM (at the beginning of the 'Execute The Requested Command' section):

```
JMONITOR_HOME=<install directory>

CLASSPATH=$JMONITOR_HOME/lib:$JMONITOR_HOME/lib/miniwebserver.jar:$JMONITOR_HOME/lib/jlrutils.jar:$JMONITOR_HOME/lib/jmonitor.jar:$CLASSPATH

export CLASSPATH

LIBPATH=$JMONITOR_HOME/bin/aix_ppc:$LD_LIBRARY_PATH

export LIBPATH
```

- 4** In the same section of the file, modify the java command-line.

For Windows 2000/NT, Solaris, and AIX, add the following parameter:

```
-Xrunjmonitor:load_mws,proxy
```

For Windows, the Java command line is found in the 'doneSetArgs' section. For Solaris and AIX, the Java call to start Tomcat is in the 'Execute The Requested Command' section.

- 5** To start the Tomcat server with JMonitor, use the following command:  
catalina\_jmonitor start

To stop the Tomcat server with JMonitor, use the regular *catalina.bat* or *catalina.sh* file.



## Configuring the EJB Monitor and the Sitraka JMonitor on the Application Server

If you want to monitor EJB objects using the Sitraka JMonitor and the EJB Monitor together, you must first install the EJB Monitor and run the monitor detector. For more information, see “Installing the EJB Monitor and Running the Monitor Detector,” on page 452.

### Installing the Sitraka JMonitor on the Application Server

To install the Sitraka JMonitor, create a home directory for the Sitraka JMonitor support files—for example, `MERC_MONITOR_HOME`—and unzip the installation file `<ProTune_Installation>\Add-ins\J2EE\Sitraka\jmonitor_<platform>.jar` file into that directory.

On UNIX platforms, use the `jar` utility to extract the installation jar.

Change to the Sitraka JMonitor installation directory and type the following command:

```
jar -xvf <path to your Sitraka JMonitor installation jar>
```

### Support Matrix:

| Application Server | Version  | Platform              |
|--------------------|----------|-----------------------|
| WebLogic           | 6.0; 6.1 | Windows; Solaris; AIX |
| WebSphere          | 4.0      | Windows; Solaris; AIX |

After you have installed the EJB and Sitraka JMonitor support files on your WebLogic or WebSphere machine, you must configure the application server to run with EJB Monitor and Sitraka JMonitor support.

---

**Note:** It is important to set the environment variables in the order in which they appear below.

---

**To configure the WebLogic 6.0/6.1 server:**

- 1** Make a backup copy of the `<WebLogic Home>\config\<domain name>\startWeblogic.cmd` file.
- 2** Open the `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` file.
- 3** In the 'runWebLogic' section of the file, just before the Java command line used to start the server, set the following environment variables:

---

**Note:** For IBM AIX platform replace LD\_LIBRARY\_PATH with LIBPATH. Replace `<EJB Monitor Home Directory>` with the EJB monitor installation root directory. Replace `<Sitraka JMonitor Home Directory>` with the Sitraka JMonitor installation root directory. On UNIX platforms you may have to export the library path variables.

---

For Windows platforms:

```
set MERC_MONITOR_HOME=<EJB Monitor Home Directory>
set CLASSPATH=%MERC_MONITOR_HOME%\dat;%MERC_MONITOR_HOME%\classes;%MERC_MONITOR_HOME%\classes\xerces.jar;%CLASSPATH%
set PATH=%PATH%;%MERC_MONITOR_HOME%\bin
set JMONITOR_HOME=<Sitraka Jmonitor Home Directory>
set CLASSPATH=%JMONITOR_HOME%\lib;%JMONITOR_HOME%\lib\miniwebserver.jar;%JMONITOR_HOME%\lib\jlrutils.jar;%JMONITOR_HOME%\lib\jmonitor.jar;%CLASSPATH%
set PATH=%PATH%;%JMONITOR_HOME%\bin\win32_ia32
```

For UNIX platforms:

MERC\_MONITOR\_HOME <EJB Monitor Home Directory>

```
CLASSPATH${MERC_MONITOR_HOME}/
dat:${MERC_MONITOR_HOME}/classes:${MERC_MONITOR_HOME}/classes/
xerces.jar:${CLASSPATH}
```

LD\_LIBRARY\_PATH\${LD\_LIBRARY\_PATH}:\${MERC\_MONITOR\_HOME}/bin

JMONITOR\_HOME <Sitraka Jmonitor Home Directory>

```
CLASSPATH${JMONITOR_HOME}/lib:${JMONITOR_HOME}/lib/
miniwebserver.jar:${JMONITOR_HOME}/lib/jlrutils.jar:${JMONITOR_HOME}/lib/
jmonitor.jar:${CLASSPATH}
```

```
setenv LD_LIBRARY_PATH${LD_LIBRARY_PATH}:${JMONITOR_HOME}/bin/
win32_ia32
```

**4** In the same section of the file, add the following parameter:

```
-Xrunjdkhook:jmonitor:load_mws,proxy
```

For Windows platforms:

```
%JAVA_HOME%\bin\java -hotspot -ms64m -mx64m -
Xrunjdkhook:jmonitor:load_mws,proxy -classpath
%CLASSPATH%
-Dweblogic.Domain=mydomain
-Dweblogic.Name=myserver
"-Dbea.home=f:\bea"
"-Djava.security.policy==f:\bea\wlsrver6.0/lib/weblogic.policy"
-Dweblogic.management.password=%WLS_PW% weblogic.Server
```

For UNIX platforms:

```
${JAVA_HOME}/bin/java -hotspot -ms64m -mx64m
-Xrunjdkhook:jmonitor:load_mws,proxy -classpath ${CLASSPATH}
-Dweblogic.Domain=mydomain -Dweblogic.Name=myserver "
-Dbea.home=usr/bea"
"-Djava.security.policy==usr/bea/wlsrver6.0/lib/weblogic.policy"
-Dweblogic.management.password=${WLS_PW} weblogic.Server
```

- 5 Run the `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` file to start the server with JMonitor and ProTune EJB monitor support.

**To configure the WebSphere 4.0 server:**

- 1 Ensure that the WebSphere Administrative Server is running, and start the Administrator Console.
- 2 In the WebSphere Administrative Domain tree, expand the Nodes, hostname, and Application Servers subtrees, and select the Default Server (or the Application Server you wish to use with JMonitor).
- 3 For Windows 2000/NT and Solaris, click the General tab, and add the following variables to the Environment Editor box:

---

**Note:** Replace `<EJB Monitor Home Directory>` with the EJB monitor installation root directory, and `<Sitraka Jmonitor Home Directory>` with the JMonitor installation root directory.

---

For Windows 2000/NT:

name=PATH

value=<EJB Monitor Home Directory>\bin;<Sitraka Jmonitor Home Directory>\bin\win32\_ia32

For Solaris:

name=LD\_LIBRARY\_PATH

value=<EJB Monitor Home Directory>/bin:<Sitraka Jmonitor Home Directory>/bin/solaris\_sparc

Click **OK** to close the Environment Editor.

For AIX:

If the LIBPATH environment variable has been changed, you need to link the Sitraka JMonitor and the EJB monitor libraries in the /usr/lib directory.

Add the following command:

```
#ln -s <Sitraka Jmonitor Home Directory>/bin/aix_ppc/libjmonitor.so
/usr/lib/libjmonitor.so
```

```
#ln -s <EJB Monitor Home Directory>/bin/libcjhock_mon.so
/usr/lib/libcjhock_mon.so
```

```
#ln -s <EJB Monitor Home Directory>/bin/libconfig.so /usr/lib/libconfig.so
```

```
#ln -s <EJB Monitor Home Directory>/bin/libjdkhook.so /usr/lib/libjdkhook.so
```

```
#ln -s <EJB Monitor Home Directory>/bin/libmlib_ds.so/usr/lib/
libcjhock_mon.so
```

```
#ln -s <EJB Monitor Home Directory>/bin/libmosifs.so /usr/lib/libmosifs.so
```

```
#ln -s <EJB Monitor Home Directory>/bin/libthrdutil.so/usr/lib/libthrdutil.so
```

---

**Note:** You will likely require root permissions in order to create the link. Alternatively, you can place the link in WebSphere's /bin directory (usually /usr/WebSphere/AppServer/bin).

---

- 4** Click the **JVM Settings** tab in the WebSphere Administrative Console, and add the following values to the classpath:

---

**Note:** Replace <EJB Monitor Home Directory> with the EJB monitor installation root directory, and <Sitraka Jmonitor Home Directory> with the JMonitor installation root directory.

---

For Windows 2000/NT:

```
<EJB Monitor Home Directory>\dat  
<EJB Monitor Home Directory>\classes  
<EJB Monitor Home Directory>\classes\xerces.jar  
<Sitraka Jmonitor Home Directory>\lib  
<Sitraka Jmonitor Home Directory>\lib\jlrutils.jar  
<Sitraka Jmonitor Home Directory>\lib\miniwebserver.jar  
<Sitraka Jmonitor Home Directory>\lib\jmonitor.jar
```

For Solaris or AIX:

```
<EJB Monitor Home Directory>/dat  
<EJB Monitor Home Directory>/classes  
<EJB Monitor Home Directory>/classes/xerces.jar  
<Sitraka Jmonitor Home Directory>/lib  
<Sitraka Jmonitor Home Directory>/lib/jlrutils.jar  
<Sitraka Jmonitor Home Directory>/lib/miniwebserver.jar  
<Sitraka Jmonitor Home Directory>/lib/jmonitor.jar
```

- 5 Click the **Advanced JVM Settings** button. In the Command line arguments field, add the following value for Windows 2000/NT, Solaris, and AIX:  
-Xrunjdkhook:jmonitor:load\_mws,proxy
- 6 Click the **OK** and **Apply** buttons to save the changes for the Application server. You can now start and stop your WebSphere server using the Sitraka JMonitor.

## Configuring the Sitraka JMonitor in the Console

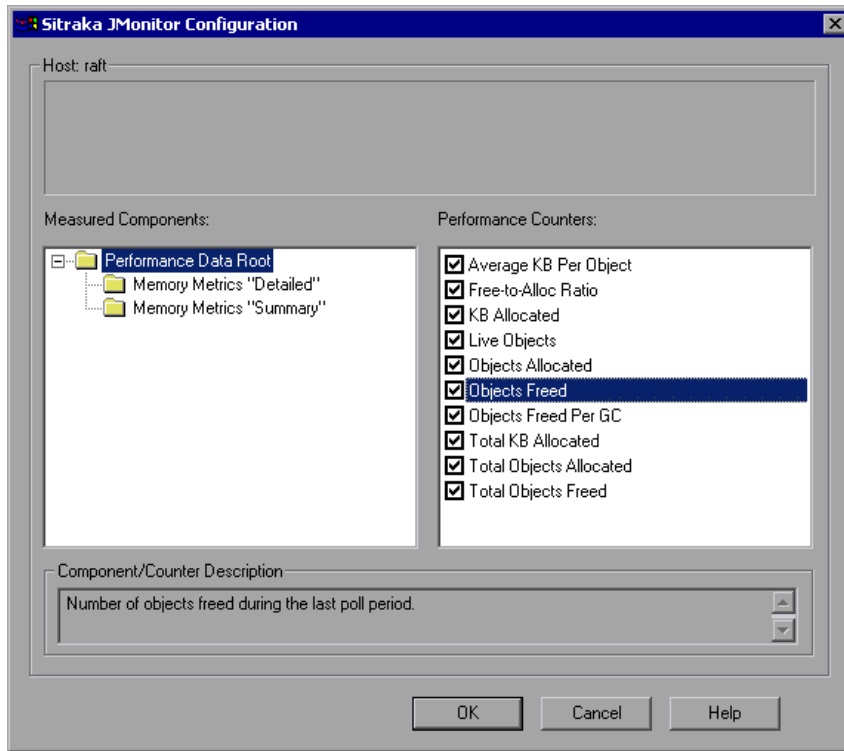
To monitor a Java-based application, you must select the counters you want the Sitraka JMonitor to measure. You select these counters using the Console's Sitraka JMonitor Configuration dialog box.

### To configure the Sitraka JMonitor monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select Sitraka JMonitor (in the Java Technology category) and then click **Add**.

The Sitraka JMonitor Configuration dialog box opens, displaying the available counters.



- 5 Expand the Measured Components tree and select the methods and counters you want to monitor.



The following counters are available for the Sitraka JMonitor:

### SummaryMemoryMetrics

| Measurement                 | Description  |
|-----------------------------|--|
| <b>% Free Heap Space</b>    | The percentage of free heap space since the last report.   |
| <b>% GC-To-Elapsed Time</b> | The percentage of garbage collection to elapsed time.  |
| <b>% GC-To-Poll Time</b>    | The percentage of garbage collection to poll time.   |
| <b>% Used Heap</b>          | The percentage of used heap space since the last report.   |
| <b>Average GC Time (ms)</b> | The average time, in milliseconds, spent performing garbage collections since the metric was enabled. (Disabling metric resets value to zero).       |
| <b>GC Time (ms)</b>         | The time, in milliseconds, spent performing garbage collections during the last poll period.   |
| <b>Heap Size (KB)</b>       | Total heap size, in kilobytes.   |
| <b>KB Freed</b>             | The number of kilobytes freed in the last poll period.   |
| <b>KB Freed Per GC</b>      | Average number of kilobytes freed per garbage collection since the metric was enabled (Disabling the metric resets value to zero).                   |
| <b>Number of GCs</b>        | The number of garbage collections during the last poll period.   |
| <b>Total GC Time (ms)</b>   | The total time, in milliseconds, spent performing garbage collections since the metric was enabled. (Disabling the metric resets the value to zero). |
| <b>Total GCs</b>            | The total number of garbage collections since the metric was enabled. (Disabling the metric resets value to zero).                                   |

| Measurement           | Description  |
|-----------------------|--|
| <b>Total KB Freed</b> | The total number of kilobytes freed since the metric was enabled. (Disabling the metric resets value to zero). |
| <b>Used Heap (KB)</b> | Used heap size, in kilobytes.  |

### DetailedMemoryMetrics

| Measurement                  | Description   |
|------------------------------|---|
| <b>Average KB Per Object</b> | The average number of kilobytes per object since the metric was enabled. (Disabling the metric resets value to zero).                 |
| <b>Free-to-Alloc Ratio</b>   | The objects freed to objects allocated ratio since the metric was enabled (Disabling the metric resets value to zero).                |
| <b>KB Allocated</b>          | The number of kilobytes allocated since the metric was enabled. (Disabling the metric resets value to zero).                          |
| <b>Live Objects</b>          | The change in number of live objects during the last poll period.   |
| <b>Objects Allocated</b>     | The number of objects allocated in the last poll period.  |
| <b>Objects Freed</b>         | The number of objects freed during the last poll period.  |
| <b>Objects Freed Per GC</b>  | The average number of objects freed per garbage collection since the metric was enabled. (Disabling the metric resets value to zero). |
| <b>Total KB Allocated</b>    | The kilobytes allocated since metric was enabled. (Disabling the metric resets value to zero).  |

| Measurement                    | Description  |
|--------------------------------|--|
| <b>Total Objects Allocated</b> | The number of objects allocated since the metric was enabled. (Disabling the metric resets value to zero). |
| <b>Total Objects Freed</b>     | The number of objects freed since the metric was enabled. (Disabling the metric resets value to zero).     |

- 6 Click **OK** in the Sitraka JMonitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the Sitraka JMonitor monitor.



# 27

---

## J2EE Performance Monitoring

The J2EE performance monitor provides complete insight into the J2EE components on the application server (Servlets, JSP's, EJB's, JNDI, JDBC, and DB SQL calls).

This chapter describes:

- ▶ About J2EE Performance Monitoring
- ▶ Installing the J2EE Monitor on the Application Server
- ▶ Initial J2EE Monitor Configuration Settings
- ▶ Activating the J2EE Monitor on the Client Machine
- ▶ Examples of Modifying Application Server Configurations
- ▶ Troubleshooting the J2EE Monitor

## About J2EE Performance Monitoring

The J2EE monitor provides the following information for each J2EE component:

- ▶ Average response time per method/query
- ▶ Number of method calls per second

With such coverage of the J2EE architecture, users can get an overview of the entire activity within the system. They can very easily correlate the end user response time with the Web server activity (Servlets and JSPs data), application server activity (JNDI and EJB's), and back-end activity of database requests (JDBC methods and SQL queries).

The J2EE Monitor allows ProTune users to analyze J2EE component metrics during a session step run by using an agent which is installed on the application server to collect information on the J2EE components. These measurements are sent from the application server back to the ProTune Console through a Web server contained in the J2EE monitor. The J2EE Monitor supports the leading applications servers, such as: IBM WebSphere, BEA WebLogic, Oracle 9iAS and JBoss. For information about the supported application servers, refer to the "Support Matrix" on page 493.

---

**Note:** The J2EE Monitor requires MSXML 3.0 and later (this is included in Internet Explorer 6.0). You can install MSXML 3.0 from the Microsoft MSDN Web site (<http://msdn.microsoft.com/downloads/default.asp?url=/downloads/sample.asp?url=/msdn-files/027/001/772/msdncompositedoc.xml>).

---

## Installing the J2EE Monitor on the Application Server

In order to monitor J2EE objects, you must first install and activate the J2EE monitor on the application server machine. You then configure the J2EE monitor on the client machine by selecting the counters you want the monitor to measure.

You can monitor Java 2 Platform, Enterprise Edition (J2EE) objects on a WebLogic, WebSphere, Oracle 9iAS, or JBoss application server during a session step run using the J2EE performance monitor.

### Support Matrix

| Application Server | Version            | Platform              |
|--------------------|--------------------|-----------------------|
| WebLogic           | 4.x; 5.x; 6.x; 7.0 | Windows; Solaris; AIX |
| WebSphere          | 3.x; 4.x           | Windows; Solaris; AIX |
| Oracle 9iAS        | 1.0.2.2            | Windows; Solaris; AIX |
| JBoss              | 2.4.x              | Windows; Solaris; AIX |

### To install the J2EE monitor on the application server:

- 1 Create a home directory on the application server machine—for example, *J2EEMonitor*, and unzip the installation file `<ProTune CD>\Add-ins\J2EE\jmonitor_<platform>.jar` file into that directory.

If you do not have WinZip to unzip the installation file, use the following command line to extract the installation file:

```
<JDK>\bin\jar.exe -xf <installation file>
```

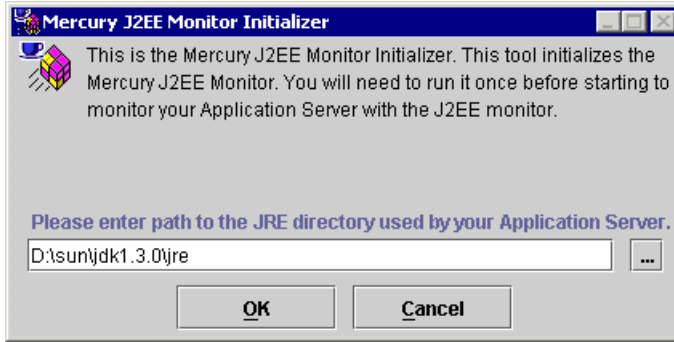
If you are working on a UNIX platform, UNIX scripts extracted from the jar file may lose their execute permissions. To fix this, open the *J2EEMonitor Home Directory*, and change the permissions using the command line:  
`chmod +x *.sh.`

- 2 Double-click the *sipatool.jar* file located in `<J2EEMonitor Home Directory>\classes` to open the Mercury J2EE Monitor Initializer.

---

**Note:** If you are working on a UNIX platform, or if the *.jar* extension in your system is not associated with the Java runtime environment, go to the *<J2EEMonitor Home Directory>\classes* directory, and type `java -jar sipatool.jar`.

---



- 3** In the Mercury J2EE Monitor Initializer, enter the path to the application server Java home directory, and click **OK** to run the tool.
- 4** Add `-Xbootclasspath/p:<J2EEMonitor Home Directory>\classes\boot` to the application server command line arguments.

Refer to “Examples of Modifying Application Server Configurations”, on page 500 to see syntax for WebLogic, WebSphere, Oracle 9iAS, or JBoss application servers.



## Initial J2EE Monitor Configuration Settings

The J2EE monitor application server installation configured the hooking mechanism, operation mode, JDBC, and EJB information retrieval.

**Hooking mechanism:** The J2EE monitor uses the Mercury J2EE Monitor Initializer and Java hooking library.

**Operation mode:** The J2EE monitor uses the Auto Discovery operating mode. In this mode, the system automatically discovers the J2EE components (Servlet, JSP, JNDI, EJB and JDBC) that actually participate in the business process.

**JDBC information retrieval:** The JDBC information retrieval setting determines which data to return from the JDBC call. By default, the J2EE monitor aggregates the measured data according to the JDBC operation, for example: SELECT, UPDATE, CREATE. To modify this configuration, refer to “Configuring JDBC Information Retrieval:” on page 496.

**EJB information retrieval:** The EJB information retrieval setting determines which data to return from the EJB call. By default, the J2EE monitor is not configured to measure container methods, (e.g., `ejbPassivate()`, `ejbCreate()`). To modify this configuration, refer to “Configuring the EJB Information Retrieval” on page 496.

---

**Note:** For information about alternative configuration settings, please contact Mercury Interactive Customer Support.

---

### **Configuring JDBC Information Retrieval:**

- 1** Open `<J2EEMonitor Home Directory>\dat\monitor.properties`.
- 2** In the property `monitor.jdbc.mode`, enter one of the following:
  - "1" to measure the JDBC the method calls, like any other (non-JDBC) measured method calls.
  - "2" to aggregate the measured data according to the JDBC operation, for example: SELECT,UPDATE,CREATE.
  - "3" to aggregate the measured data according to specific SQL statement (including the operation, the table(s) it acted on, and other parameters of this statement).

---

**Note:** SQL Statements that exceed 3000 characters in length are not supported.

---

### **Configuring the EJB Information Retrieval**

- 1** Open `<J2EEMonitor Home Directory>\dat\java_monitor.ini`.
- 2** In the "EJB\_CONFIG" section of the file, set the following property:  
`hook_files` to `hook_files=auto_detect_container`.

## Activating the J2EE Monitor on the Client Machine

To monitor J2EE performance, you must select the counters you want the J2EE monitor to measure. You select these counters using the Console's J2EE Monitor Configuration dialog box.

### Before configuring the J2EE monitor:

In Auto Discovery mode (the J2EE monitor's default operating mode), the system discovers which methods of the components (Servlet, JSP, JNDI, EJB and JDBC) are participating in your business process and measures those objects only.

To start the Auto Discovery process, start the application server, and run the Vuser script that you intend to use in your load test against the application server. This provides the Console with a list of measurements that will be available for monitoring.

---

**Note:** The next time you run the same script, you don't need to run a Vuser before selecting the methods and counters you want to monitor.

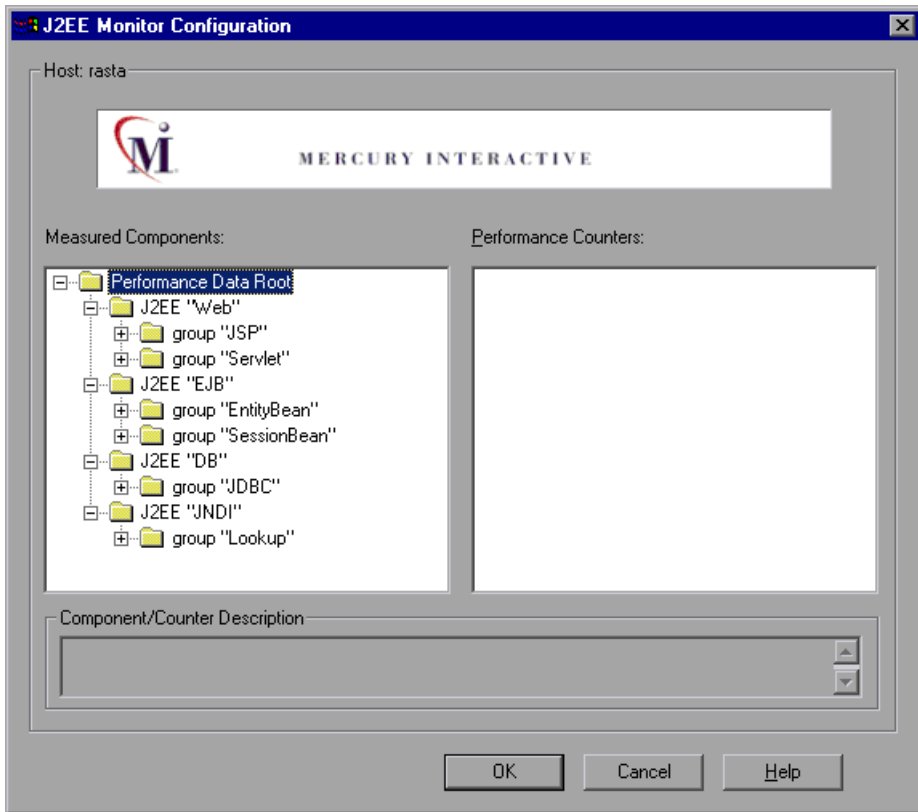
---

### To configure the J2EE monitor:



- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select J2EE (in the Java Technology category) and then click **Add**.

The J2EE Monitor Configuration dialog box opens, displaying the available J2EE counters.



- Expand the Measured Components tree and select the methods and counters you want to monitor. The following counters can be monitored for each method:

| Measurement             | Description   |
|-------------------------|---|
| Average Response Time   | The average response time, in milliseconds, of the J2EE object being monitored. |
| Method Calls per Second | The number of J2EE object method calls per second.                              |

---

**Note:** The size of a measurement name that can be displayed in the Analysis is limited to 255 characters. If a measurement name exceeds this limit, the counter name is truncated, and given a unique ID (UID). If you monitor different events or make cross result graphs on the same counter, the UID will stay the same.

The measurement name is truncated as follows:

standard prefix/**counter truncated name<UID>**/monitored event

For example:

```
/DB/JDBC/weblogic.jdbc.rmi.SerialPreparedStatement/int  
executeUpdate()/INSERT INTO orders ( orderid _ userid _ orderdate _  
shipaddr1 _ shipaddr2 _ shipcity _ shipstate _ shipzip _ shipcountry _  
billaddr1 _ billaddr2 _ b <1> /Average Response Time
```

The full measurement name appears in the Measurement Description box.

---

- 6 Click **OK** in the J2EE Monitor Configuration dialog box, and in the Select Measurements to Monitor dialog box, to activate the J2EE monitor.

## Examples of Modifying Application Server Configurations

When you installed Mercury Interactive's J2EE monitor files on your application server, you already configured it to run with J2EE monitor support. This section provides examples modifying the configuration of the following application servers:

- ▶ WebLogic Server
- ▶ WebSphere Server - Version 3.x
- ▶ WebSphere Server - Version 4.x
- ▶ Oracle 9iAS Server
- ▶ JBoss 2.4.x Server

---

**Note:** It is important to set the environment variables in the order in which they appear below.

---

### WebLogic Server

The WebLogic 4.x-5.x server, WebLogic 6.x server, and the WebLogic 7.x server are configured differently.

#### To configure the WebLogic 4.x-5.x server:

- 1** Copy the `<WebLogic Home>\startWeblogic.cmd` file into `<WebLogic Home>\startWeblogicMercury.cmd` so that the file is backed up.
- 2** Open the `<WebLogic Home>\startWeblogicMercury.cmd` file.
- 3** Just before the Java command line used to start the server add the following variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<J2EEMonitor Home Directory>
set JAVA_CLASSPATH=%JAVA_CLASSPATH%;%MERC_MONITOR_HOME
%\dat
```

For UNIX platforms (csh):

```
MERC_MONITOR_HOME <J2EEMonitor Home Directory>
JAVACLASSPATH=$JAVACLASSPATH:$MERC_MONITOR_HOME/dat
```

- 4 In the same section of the file, add the following parameter to the Java command line: `-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot`

**Example:**

```
%JAVA_HOME%\bin\java -ms64m -mx64m -
-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot
-Dweblogic.class.path=%WEBLOGIC_CLASSPATH% -Dweblogic.home=.
-Djava.security.manager
-Djava.security.policy==.\weblogic.policy weblogic.Server
```

- 5 Run the `<WebLogic Home>\startWeblogicMercury.cmd` file.

**To configure the WebLogic 6.x server:**

- 1 Copy the `<WebLogic Home>\config\<domain name>\startWeblogic.cmd` file into `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` so that the file is backed up.
- 2 Open the `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` file.
- 3 Just before the java command line used to start the server add the following variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<J2EEMonitor Home Directory>
set CLASSPATH=%CLASSPATH%;%MERC_MONITOR_HOME%\dat
```

For UNIX platforms:

```
MERC_MONITOR_HOME=<J2EEMonitor Home Directory>
CLASSPATH=$CLASSPATH:$MERC_MONITOR_HOME/dat
```

- 4 In the same section of the file add a parameter to the command line:

```
-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot
```

**Example:**

- 5 Run the `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` file.

**To configure the WebLogic 7.x server:**

- 1 Copy the `<WebLogic Home>\server\bin\startwls.cmd` file into `<WebLogic Home>\server\bin\startwlsMercury.cmd` so that the file is backed up.
- 2 Open the `<WebLogic Home>\server\bin\startwlsMercury.cmd` file.
- 3 Just before the java command line used to start the server add the following variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<J2EEMonitor Home Directory>  
set CLASSPATH=%CLASSPATH%;%MERC_MONITOR_HOME%\dat
```

For UNIX platforms:

```
MERC_MONITOR_HOME=<J2EEMonitor Home Directory>  
CLASSPATH=$CLASSPATH:$MERC_MONITOR_HOME/dat
```

- 4 In the same section of the file add a parameter to the command line:

```
-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot
```

**Example:**

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m  
-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot  
-classpath %CLASSPATH% -Dweblogic.Domain=mydomain  
-Dweblogic.Name=myserver "-Dbea.home=f:\bea" "  
-Djava.security.policy==f:\bea\wlserver6.0/lib/weblogic.policy"  
-Dweblogic.management.password=%WLS_PW% weblogic.Server
```



- 5 Copy the `<domain name>\startWeblogic.cmd` file into `<domain name>\startWeblogicMercury.cmd` so that the file is backed up.
- 6 Open the `<domain name>\startWeblogicMercury.cmd` file.
- 7 Find the call to the Weblogic server. For example, call:  
D:\bea\weblogic700\server\bin\startWLS.cmd
- 8 Change the call from `startWLS.cmd` to `startWLSMercury.cmd`, and save the file.

### WebSphere Server - Version 3.x

By default, the WebSphere 3.x application server runs on Windows as an automatic service, upon machine startup. Since Mercury Interactive does not currently support ProTune J2EE monitoring on a WebSphere server run as an automatic service, you must change the default WebSphere server startup to *manual*.

#### To change the default WebSphere 3.x server startup:

- 1 Select **Start > Settings > Control Panel**.
- 2 Double-click **Services**.
- 3 Select **IBM WS AdminServer**, and click the **Stop** button.
- 4 Double-click **IBM WS AdminServer**, and select the **Manual** Startup Type.
- 5 Click **OK** to save your settings and close the dialog box.

You can now start the WebSphere Server from `<WebSphere Home>\AppServer\bin\debug\adminserver.bat`, instead of using the automatic service.

#### To add ProTune J2EE monitor support to the WebSphere 3.x server:

- 1 Make a backup copy of the `<WebSphere Home>\AppServer\bin\debug\adminserver.bat` file.
- 2 Open the `<WebSphere Home>\AppServer\bin\debug\adminserver.bat` file.
- 3 Add the following environment variables at the end of the 'SET\_CP' section:

For Windows platforms:

```
set MERC_MONITOR_HOME=<J2EEMonitor Home Directory>
```

For UNIX platforms:

```
MERC_MONITOR_HOME=<J2EEMonitor Home Directory>  
export MERC_MONITOR_HOME
```

- 4 Run the *adminserver.bat* file.
- 5 Open the WebSphere Advanced Administrative Console, and select **View > Topology**.
- 6 Expand the WebSphere Administrative Domain tree by selecting **<server machine name> > Default Server**.
- 7 Select the **General** tab in the Application Server:Default Server window.
- 8 Add `-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot` to the command line Arguments box, and click **Apply**.

If you are working with a WebSphere 3.0 Server with JDK1.1.7 IBM, double-click on **Environment**. Type `_CLASSLOAD_HOOK` in the Variable Name box, and `jdkhook` in the Value box. Click the **Add**, **OK**, and **Apply** buttons.

- 9 For Windows 2000/NT or Solaris, open the Environment Editor dialog box from the General tab, and add the following variables to the Environment box:

For Windows 2000/NT:

```
name=CLASSPATH  
value=<J2EEMonitor Home Directory>\dat
```

For Solaris:

```
name=CLASSPATH  
value=<J2EEMonitor Home Directory>\dat
```

Click **OK** to close the Environment Editor dialog box.

- 10 Close the WebSphere Advanced Administrative Console.
- 11 Close and restart the *adminserver.bat* file.

## WebSphere Server - Version 4.x

You can start the WebSphere 4.x server using the startServerBasic.bat file or the startServer.bat file.

### To configure the WebSphere 4.x server:

- 1** Ensure that the WebSphere Administrative Server is running, and start the Administrator Console.
- 2** In the WebSphere Administrative Domain tree, expand the Nodes, Hostname, and Application Servers subtrees, and select the Default Server (or the application server you wish to use with J2EE monitor).
- 3** Right-click the Default Server, select Properties from the menu, and click the General tab.
- 4** For Windows 2000/NT or Solaris, open the Environment Editor dialog box from the General tab, and add the following variables to the Environment box:

For Windows 2000/NT:

```
name=CLASSPATH
value=<J2EEMonitor Home Directory>\dat
```

For Solaris:

```
name=CLASSPATH
value=<J2EEMonitor Home Directory>\dat
```

Click **OK** to close the Environment Editor dialog box.

- 5** Click the Advanced JVM Settings tab and select Advanced JVM settings. In the Command line arguments field, add the following value for Windows 2000/NT, Solaris, and AIX:
 

```
-Xbootclasspath/p:%MERC_MONITOR_HOME%\classes\boot
```
- 6** Click the **OK** and **Apply** buttons to save the changes for the Application server. You can now start and stop your WebSphere server using the ProTune J2EE Monitor.

### Oracle 9iAS Server

- 1 Edit the file *env.cmd* (*env.sh* in Unix platforms) as follows:
  - the JAVA\_HOME environment variable should point to the location of the Java Virtual machine used to run the application server.
  - the DETECTOR\_INS\_DIR environment variable should point to the location of the monitor installation.
  - the APP\_SERVER\_DRIVE environment variable should specify the drive hosting the application server installation (e.g., D:). Do not modify this variable on Unix Platforms.
  - the APP\_SERVER\_ROOT environment variable should specify the application server root directory.
- 2 Run the *oc4jMonitor.cmd* (*oc4jMonitor.sh* in Unix platforms).

### JBoss 2.4.x Server

- 1 Make a backup copy of *<JBoss Home>\run.bat* (*run.sh* on Unix platforms) file into *<JBoss Home>\runMercury.bat* (*runMercury.sh* for Unix).
- 2 Open the *<JBoss Home>\runMercury.bat* file (*runMercury.sh* on Unix).

Just before the Java command line used to start the server add the following variables:

For Windows platforms:

```
set MERC_MONITOR_HOME=<J2EE Monitor Home Directory>
```

For UNIX platforms:

```
MERC_MONITOR_HOME=<J2EE Monitor Home Directory>
```

- 3 In the same section of the file add the following parameter to the command line:
  - Xbootclasspath/p:%MERC\_MONITOR\_HOME%\classes\boot

**Example:**

```
%JAVA_HOME%\bin\java -ms64m -mx64m -Xbootclass-
path/p:%MERC_MONITOR_HOME%\classes\boot
-Dweblogic.class.path=%WEBLOGIC_CLASSPATH% -Dweblogic.home=.
-Djava.security.manager
-Djava.security.policy==.\weblogic.policy weblogic.Server
```

- 4 Run the `<JBoss Home>\runMercury.bat` (`<JBoss Home>\runMercury.sh`) file.

## Troubleshooting the J2EE Monitor

### Changing the Default Port

The J2EE monitor communicates with ProTune, by default, using port 2004. If this port has already been taken, you can select another port as follows:

- 1 On the application server machine, open `<J2EEMonitor Home Directory>\dat\monitor.properties` and change the port number specified in the property: `webserver.monitor.port`
- 2 On the ProTune machine, open `<ProTune Home Directory>\dat\monitors\xmlmonitorshared.ini` and change the port number specified in section “`mon_j2ee`” under the key “`DefaultPort`”.

### Initialization Errors

If you are getting application server initialization errors such as: “`UnsupportedClassVersionError`”, “`NoSuchMethodError`” or “`NoClassDefFoundError`”, there might be a conflict between the JDK version specified using the Mercury J2EE Monitor\_INITIALIZER, and the actual JDK version used in application server launch.

Make sure that you selected the correct JDK that is currently being used by the application server. Note that if you switched the application server to work with a different JDK, you must run the Mercury J2EE Monitor\_INITIALIZER again.



# 28

---

## Application Deployment Solution Monitoring

During a session step run, you can monitor the Citrix MetaFrame XP server in order to isolate server performance bottlenecks.

This chapter describes:

- ▶ Monitoring Citrix MetaFrame Servers
- ▶ Configuring the Citrix MetaFrame Server Monitor

### About Application Deployment Solution Monitoring

ProTune's Citrix MetaFrame XP monitor provides you with information about the application deployment usage of the Citrix MetaFrame XP server during a session step execution. The Citrix Monitor allows you to monitor the server performance statistics from Citrix Metaframe Servers. You can monitor multiple parameters (counters) with a single monitor instance. This allows you to watch server loading for performance, availability, and capacity planning.

### Monitoring Citrix MetaFrame Servers

The Citrix MetaFrame Resource Monitor is an Application Deployment Solution monitor that provides performance information for the Citrix MetaFrame XP server. In order to obtain this data, you need to activate the Citrix MetaFrame server monitor before executing the session step, enable the counters you want to monitor on the Citrix server, and specify which measurements and resources you want the Citrix monitor to measure.

## Configuring the Citrix MetaFrame Server Monitor

You select measurements to monitor the Citrix MetaFrame server using the Citrix MetaFrame XP dialog box.

---

**Note:** The port you use to monitor a Citrix MetaFrame server through a firewall depends on the configuration of your server.

---

### Before setting up the Citrix MetaFrame Server monitor:

- 1** Make sure that Citrix MetaFrame Server has been installed and is running on a computer. If the computer running Citrix MetaFrame Server is running Windows 2000, make sure that the Remote Registry service is running on it.
- 2** Make sure that the computer on which you are running ProTune has administrator privileges on the machine running Citrix.
- 3** From the Console machine, map a network drive to the Citrix server machine. This ensures that the required authentication is provided to the Console to access the resource counters.
- 4** Launch PerfMon from the Console machine to enable the counters on the Citrix server. This allows you to monitor the same counters for the ICA Session object on the Citrix monitor.
- 5** In the System Topology window, add an Application Deployment element representing the server running Citrix. Assign the following values to the element's settings:
  - ▶ **O/S:** The version of Windows running on the Citrix machine
  - ▶ **Product:** Citrix MetaFrame
  - ▶ **Version:** XP



- 6 You can configure the Citrix monitor to view ICA Session object counters only if at least one session is being run on the Citrix server. If no “real” user has opened a connection with the Citrix server, you need to first initialize or run a Citrix Vuser against the server, and only then configure the Citrix Monitor and add the ICA Session counters. If you configure the Citrix monitor without first initializing or running a Citrix Vuser (or connecting to the Citrix server as a “real” user), you will not be able to view the ICA Session object.

---

**Note:** Measurements that monitor instances are valid for the currently running Citrix session only. If you run this session step again, you will need to reconfigure the measurements that are instance-oriented.

---

---

**Note:** In order to monitor the different instances, ensure that the server login and logout procedures are recorded in the Vuser\_init and Vuser\_end sections respectively, and not in the Action section of the script. For more information, refer to the *Creating Vuser Scripts* guide.

---

**To configure the Citrix MetaFrame Server monitor:**

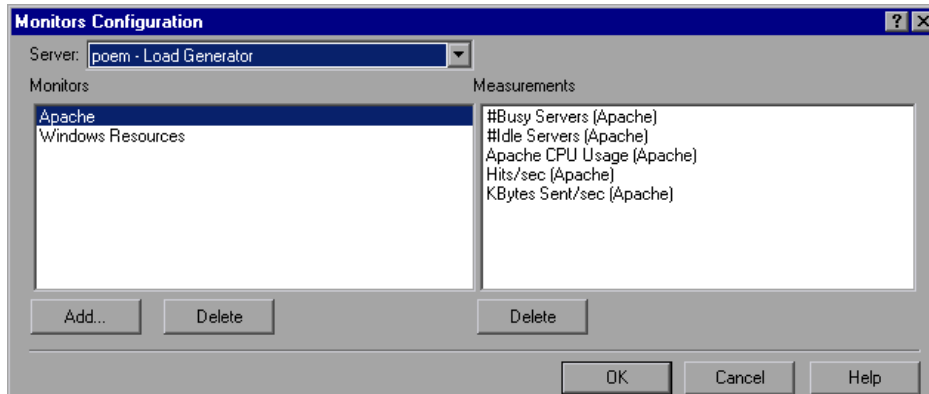
- 1 Click the Citrix MetaFrame graph in the Available Graphs section of the **Execute** tab, and drag it into the right pane of the tab.

---

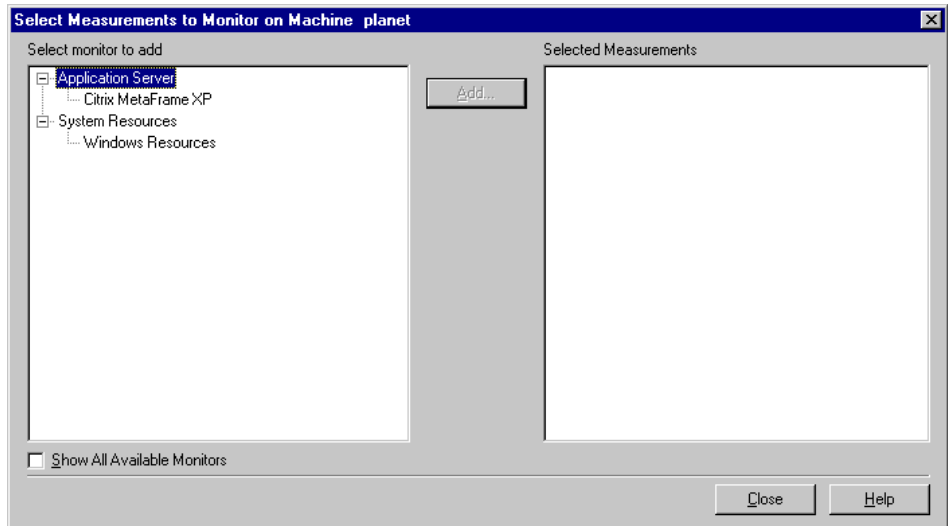
**Tip:** You’ll find the Citrix MetaFrame graph in the Application Deployment Solutions category.

---

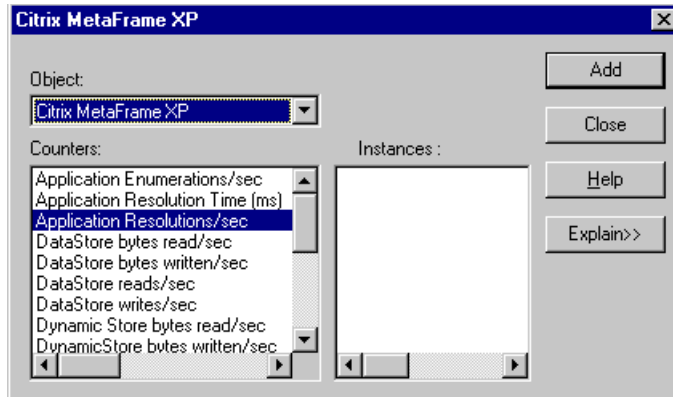
- 2 Right-click the graph and choose **Monitors Configuration**, or click the **Monitors** button on the toolbar. The **Monitors Configuration** dialog box appears.



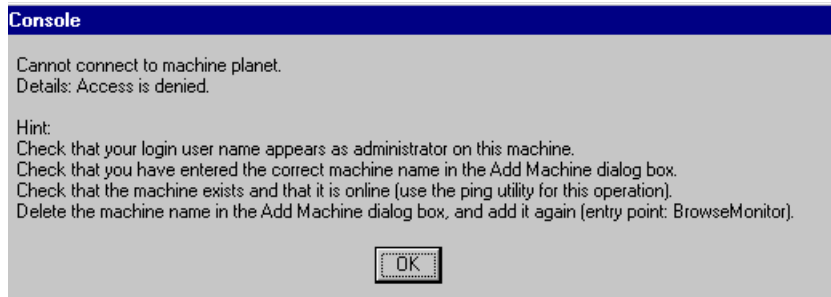
- 3 From the Server list, choose the server running Citrix and click **Add**. The **Select Measurements to Monitor on Machine planet** dialog box appears.



- 4 In the left pane, expand the Application Server element, click Citrix MetaFrame XP, and click **Add**. The Citrix MetaFrame XP dialog box is displayed.



**Note:** If, instead of the Citrix MetaFrame XP dialog box, ProTune displays the following error message,



it means that the computer on which you are running ProTune does not have administrator privileges on the Citrix machine. To remedy this, close the ProTune session, acquire administrator privileges on the Citrix machine, and reopen the ProTune session. Then repeat this procedure, starting again with step 1 above.

---

**Note:** If the dialog box freezes after clicking Add, you may need to rebuild the localhost cache on the Citrix server machine.

For more information, refer to Documents IDs CTX003648 and CTX759510 in the Citrix Knowledge Base (<http://knowledgebase.citrix.com/cgi-bin/webcgi.exe?New,KB=CitrixKB>).

---

- 5** The Object listbox allows you to display three types of counters: Citrix MetaFrame XP, Citrix IMA Networking, and ICA Session. Choose the object whose counters you want to display. ProTune displays the object's counters in the Counters pane.

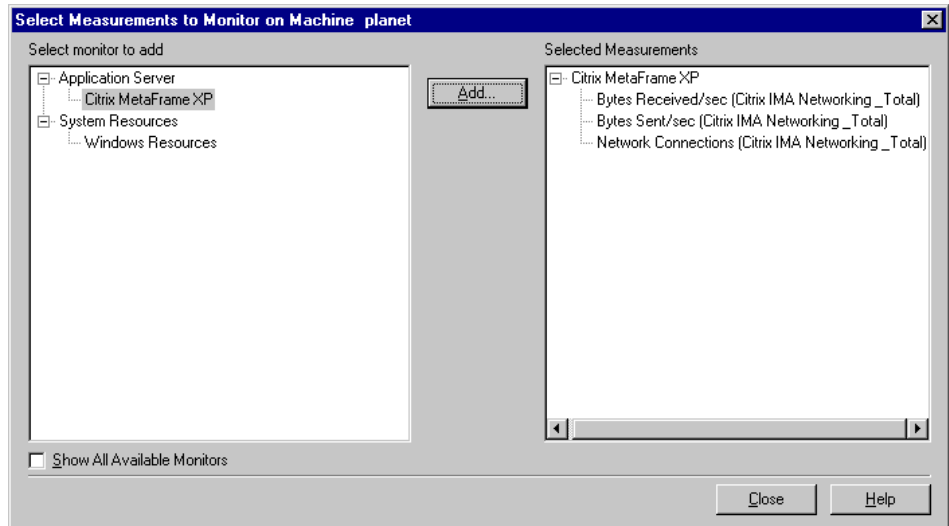
The following table lists some of the counters that can be measured:

| Measurement                     | Description  |
|---------------------------------|--|
| <b>% Disk Time</b>              | The percentage of elapsed time that the selected disk drive is busy servicing read or write requests.  |
| <b>% Processor Time</b>         | The percentage of time that the processor is executing a non-Idle thread. This counter is a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the Idle process in each sample interval, and subtracting that value from 100%. (Each processor has an Idle thread which consumes cycles when no other threads are ready to run). It can be viewed as the percentage of the sample interval spent doing useful work. This counter displays the average percentage of busy time observed during the sample interval. It is calculated by monitoring the time the service was inactive, and then subtracting that value from 100%. |
| <b>File data Operations/sec</b> | The rate that the computer is issuing Read and Write operations to file system devices. It does not include File Control Operations.   |

| Measurement                      | Description  |
|----------------------------------|--|
| <b>Input Session Bandwidth</b>   | This value represents the bandwidth from client to server traffic for a session in bps   |
| <b>Interrupts/sec</b>            | The average number of hardware interrupts the processor is receiving and servicing in each second. It does not include DPCs, which are counted separately. This value is an indirect indicator of the activity of devices that generate interrupts, such as the system clock, the mouse, disk drivers, data communication lines, network interface cards and other peripheral devices. These devices normally interrupt the processor when they have completed a task or require attention. Normal thread execution is suspended during interrupts. Most system clocks interrupt the processor every 10 milliseconds, creating a background of interrupt activity. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval. |
| <b>Latency – Session Average</b> | This value represents the average client latency over the life of a session.   |
| <b>Output Seamless Bandwidth</b> | This value represents the bandwidth from server to client traffic on this virtual channel. This is measured in bps   |
| <b>Output Session Bandwidth</b>  | This value represents the bandwidth from server to client traffic for a session in bps   |
| <b>Page Faults/sec</b>           | A count of the Page Faults in the processor. A page fault occurs when a process refers to a virtual memory page that is not in its Working Set in main memory. A Page Fault will not cause the page to be fetched from disk if that page is on the standby list, and hence already in main memory, or if it is in use by another process with whom the page is shared.   |

| Measurement                   | Description  |
|-------------------------------|--|
| <b>Pages/sec</b>              | The number of pages read from the disk or written to the disk to resolve memory references to pages that were not in memory at the time of the reference. This is the sum of Pages Input/sec and Pages Output/sec. This counter includes paging traffic on behalf of the system Cache to access file data for applications. This value also includes the pages to/from non-cached mapped memory files. This is the primary counter to observe if you are concerned about excessive memory pressure (that is, thrashing), and the excessive paging that may result. |
| <b>Pool Nonpaged Bytes</b>    | The number of bytes in the Nonpaged Pool, a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. Nonpaged Pool pages cannot be paged out to the paging file, but instead remain in main memory as long as they are allocated.   |
| <b>Private Bytes</b>          | The current number of bytes this process has allocated that cannot be shared with other processes.   |
| <b>Processor Queue Length</b> | The instantaneous length of the processor queue in units of threads. This counter is always 0 unless you are also monitoring a thread counter. All processors use a single queue in which threads wait for processor cycles. This length does not include the threads that are currently executing. A sustained processor queue length greater than two generally indicates processor congestion. This is an instantaneous count, not an average over the time interval.   |
| <b>Threads</b>                | The number of threads in the computer at the time of data collection. Notice that this is an instantaneous count, not an average over the time interval. A thread is the basic executable entity that can execute instructions in a processor.   |

- 6 Select a counter and instance. You can select multiple counters using the **Ctrl** key. The instance is relevant only if multiple instances of the highlighted counter are running. For a description of each counter, click **Explain>>** to expand the dialog box.
- 7 Click **Add** to add the counter to the list of counters that you want to monitor.
- 8 Repeat steps 6 through 7 until you've added all the Citrix counters that you want to monitor.
- 9 Click **Close**. The counters you selected are displayed in the Selected Measurements pane of the Select Measurements to Monitor dialog box.



- 10 Click **Close**.

---

**Note:** For troubleshooting tips and limitations, see the section on “Troubleshooting Server Resource Monitors,” on page 541.

---





# 29

---

## Middleware Performance Monitoring

Using ProTune's Middleware Performance monitors, you can monitor the TUXEDO and the IBM WebSphere MQ servers during a session step run and isolate server performance bottlenecks.

This chapter describes:

- ▶ Configuring the IBM WebSphere MQ Monitor
- ▶ Configuring the TUXEDO Monitor

### About Middleware Performance Monitoring

A primary factor in a transaction's response time is the middleware performance usage. ProTune's Middleware Performance monitors provide you with information about the middleware performance usage of the TUXEDO and IBM WebSphere MQ servers during a session step execution. In order to obtain performance data, you need to activate the online monitor for the server and specify which resources you want to measure before executing the session step.

The IBM WebSphere MQ monitor is used to monitor channel and queue performance counters on an IBM WebSphere MQ (version 5.x) Server.

The TUXEDO monitor can monitor the server, load generator machine, workstation handler, and queue in a TUXEDO system. Note that in order to run the TUXEDO monitor, you must install the TUXEDO client libraries on the machine you want to monitor.

The procedures for selecting monitor measurements and configuring the monitors vary according to server type. The following sections contain specific configuration instructions for each server type.

## Configuring the IBM WebSphere MQ Monitor

To use the IBM WebSphere MQ monitor you must first install the IBM WebSphere MQ client, configure the MQ server environment to monitor events, and then select the measurements you want to monitor using the IBM WebSphere MQ Add Measurements dialog box.

---

**Note:** To monitor the MQ middleware performance monitor, the Windows user must be part of the Administration Group of the IBM WebSphere MQ server.

---

### Connecting to the IBM WebSphere MQ Server

The IBM WebSphere MQ monitor connects to the IBM WebSphere MQ server (via the MQ Client Connection installed on the Console machine). In MQ Client environments, MQ does not run on the client machine. Instead, the client machine connects to an MQ Server instance, and uses the Server's resources as if they were local to the client machine.

---

**Note:** Note: The IBM WebSphere MQ monitor provides resource usage information for machines running the IBM MQ Server (version 5.2) for Windows monitoring.

---

#### Before you set up the monitor:

Ensure that an IBM WebSphere MQ Client Connection (version 5.21 only) is installed on the Console machine.

---

**Note:** For additional information on the IBM WebSphere MQ Server/Client, please refer to the IBM MQSeries Web site (<http://www-3.ibm.com/software/ts/mqseries/library/manuals/index.htm>).

---

## Configuring the Server Environment to Monitor Events

The ProTune MQ Monitor retrieves event messages from two standard MQSeries queues only:

- SYSTEM.ADMIN.PERFM.EVENT – performance events, such as “queue depth high”
- SYSTEM.ADMIN.CHANNEL.EVENT – channel events, such as “channel stopped”

Events must be enabled for the queue manager (and in many cases, on the applicable object, as well). Performance events are enabled by setting attributes for the queue on the MQ Server. Channel events are enabled by default, and cannot be disabled.

---

**Note:** The IBM WebSphere MQ monitor does not retrieve data from a queue manager after the queue manager has been restarted.

---

**To enable performance events for the Queue Manager:**

- 1** Use the following MQSC command: ALTER QMGR PERFMEEV(ENABLED).
- 2** Set the following attributes for the queue:

| Measurement              | Set Event Attributes  |
|--------------------------|---|
| Event - Queue Depth High | <ul style="list-style-type: none"> <li>• QDEPTHHI(integer) – where integer is a value expressed as a percentage of maximum messages allowed, and is in the range of 0 to 100 inclusive.</li> <li>• QDPHIEV(action) – where action is the word “ENABLED” or “DISABLED”, enabling or disabling the generation of the event, respectively.</li> </ul>  |
| Event - Queue Depth Low  | <p>To enable the event for a queue, the following attributes of the queue must be set:</p> <ul style="list-style-type: none"> <li>• QDEPTHLO(integer) – where integer is a value expressed as a percentage of maximum messages allowed, and is in the range of 0 to 100 inclusive.</li> <li>• QDPLOEV(action) – where action is the word “ENABLED” or “DISABLED”, enabling or disabling the generation of the event, respectively.</li> </ul> |
| Event - Queue Full       | <ul style="list-style-type: none"> <li>• QDEPTHHI(integer) – where integer is a value expressed as a percentage of maximum messages allowed, and is in the range of 0 to 100 inclusive.</li> <li>• QDPMAXEV(action) – where action is the word “ENABLED” or “DISABLED”, enabling or disabling the generation of the event, respectively.</li> </ul>   |

| Measurement                         | Set Event Attributes  |
|-------------------------------------|---|
| Event - Queue Service Interval High | <ul style="list-style-type: none"> <li>• QSVCINT(integer) – where integer is a value expressed as milliseconds, in the range of 0 and 999,999,999, inclusive. Note: this value is shared with Queue Service Interval OK.</li> <li>• QSVCI EV(type) – where type is the word “HIGH”, “OK”, or “NONE”, enabling service interval high events, enabling service interval ok events, or disabling the generation of the event, respectively.</li> </ul>   |
| Event - Queue Service Interval OK   | <ul style="list-style-type: none"> <li>• QSVCINT(integer) – where integer is a value expressed as milliseconds, in the range of 0 and 999,999,999, inclusive. Note: this value is shared with Queue Service Interval High.</li> <li>• QSVCI EV(type) – where type is the word “HIGH”, “OK”, or “NONE”, enabling service interval high events, enabling service interval ok events, or disabling the generation of the event, respectively.</li> </ul> |

---

**Note:** If you encounter an MQ Server error message (starting with the characters MQRC\_), please refer to the Reason Codes section of the IBM MQSeries Web site (<http://www-3.ibm.com/software/ts/mqseries/library/manuals/mqw20/AMQ43M32.HTM#HDRMQSCRN>).

---

## Configuring the IBM WebSphere MQ Monitor in the Console

After you have installed the MQ Client on the Console, and configured the server environment to monitor events, you can specify which resources you want to measure.

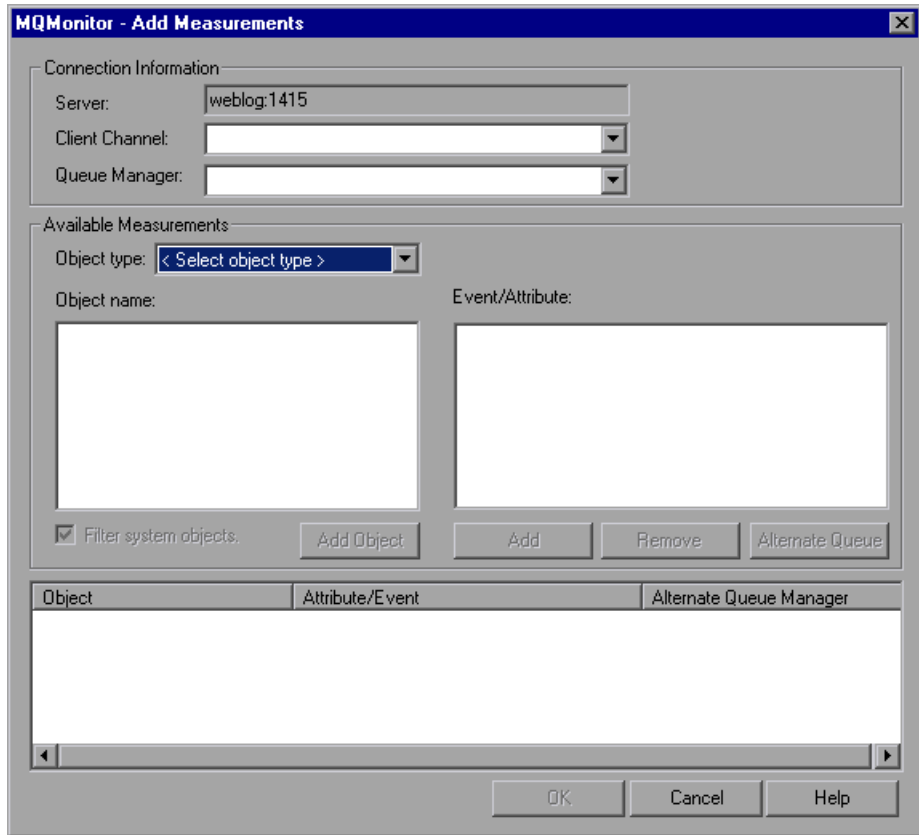
**To configure the IBM WebSphere MQ monitor:**



- 1 Click **Monitors** to open the Monitors Configuration dialog box.
- 2 Select the server whose monitors you want to configure from the Server list box.

- 3 Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4 In the left section of the dialog box, select IBM WebSphere MQ (in the Middleware Component category) and then click **Add**.

The IBM WebSphere MQ Add Measurements dialog box opens.



- 5** In the Connection Information section, enter the name of the channel through which a client connection is made to an MQ Server in the Client Channel box.

You can set up a specific channel on an MQ Server instance, or use the default “SYSTEM.DEF.SVRCONN” channel. If the client channel is undefined, the MQ Server will be inaccessible via client connections (the MQ Monitor will not work, as it will not be able to connect to the queue manager which it is supposed to monitor).

---

**Note:** User entries for any text box are limited to 48 characters.

---

- 6** Enter the name of the queue manager to be monitored in the Queue Manager box.

The monitor is not restricted to monitoring only the queue manager to which it is connected. You can configure multiple queue managers to write to the event queue of a central queue manager for centralized monitoring (this applies to Events only, not polled object attributes). All events contain a queue manager attribute identifying their source.

---

**Note:** A queue manager can only be accessed by one Console or monitoring application at any one time.

---

- 7** In the Available Measurements section, select an object type.

A list of previously added objects of the selected object type appear in the Object name list. A list of attributes or events applicable to the selected object type appear in the Events/Attributes list.

The names of monitored objects, event/attribute selected, and alternate queue managers, are listed in the monitored objects pane.

- 8** By default, only user-defined objects are displayed in the Object name list. To show all objects, clear the **Filter System Objects** check box. You can modify the filter settings, in the  
`<ProTune_installation>\dat\monitors\mqseries.cfg` file.

- 9 Select an object or add a new object to the Object name list. To add a new object name, click the **Add Object** button. In the Add Object Name dialog box, enter the name of an object to be monitored and click **OK**. The dialog box closes and the name of the object appears in the Object name list.
- 10 Select the attributes or events to be measured from the Attribute/Event box. The list of attributes or events is applicable to the selected object type.

The following tables list the available IBM WebSphere MQ monitor measurements:

### Queue Performance Counters

| Measurement   | Description   |
|---|---|
| Event - Queue Depth High (events per second)            | An event triggered when the queue depth reaches the configured maximum depth.   |
| Event - Queue Depth Low (events per second)             | An event triggered when the queue depth reaches the configured minimum depth.   |
| Event - Queue Full (events per second)                  | An event triggered when an attempt is made to put a message on a queue that is full.                                      |
| Event - Queue Service Interval High (events per second) | An event triggered when no messages are put to or retrieved from a queue within the timeout threshold.                    |
| Event - Queue Service Interval OK (events per second)   | An event triggered when a message has been put to or retrieved from a queue within the timeout threshold.                 |
| Status - Current Depth                                  | Current count of messages on a local queue. This measurement applies only to local queues of the monitored queue manager. |
| Status - Open Input Count                               | Current count of open input handles. Input handles are opened so that an application may "put" messages to a queue.       |
| Status - Open Output Count                              | Current count of open output handles. Output handles are opened so that an application may "get" messages from a queue.   |



## Channel Performance Counters

| Measurement  | Description   |
|--|---|
| <b>Event - Channel Activated (events per second)</b>       | Event generated when a channel, waiting to become active but inhibited from doing so due to a shortage of queue manager channel slots, becomes active due to the sudden availability of a channel slot.   |
| <b>Event - Channel Not Activated (events per second)</b>   | Event generated when a channel, attempts to become active but inhibited from doing so due to a shortage of queue manager channel slots.   |
| <b>Event - Channel Started (events per second)</b>         | Event generated when a channel is started.  |
| <b>Event - Channel Stopped (events per second)</b>         | Event generated when a channel is stopped, regardless of source of stoppage.  |
| <b>Event - Channel Stopped by User (events per second)</b> | Event generated when a channel is stopped by a user.  |
| <b>Status - Channel State</b>                              | The current state of a channel. Channels pass through several states from STOPPED (inactive state) to RUNNING (fully active state). Channel states range from 0 (STOPPED) to 6 (RUNNING).   |
| <b>Status - Messages Transferred</b>                       | The count of messages that have been sent over the channel. If no traffic is occurring over the channel, this measurement will be zero. If the channel has not been started since the queue manager was started, no measurement will be available.    |
| <b>Status - Buffer Received</b>                            | The count of buffers that have been received over the channel. If no traffic is occurring over the channel, this measurement will be zero. If the channel has not been started since the queue manager was started, no measurement will be available. |

| Measurement                    | Description  |
|--------------------------------|--|
| <b>Status - Buffer Sent</b>    | The count of buffers that have been sent over the channel. If no traffic is occurring over the channel, this measurement will be zero. If the channel has not been started since the queue manager was started, no measurement will be available.          |
| <b>Status - Bytes Received</b> | The count of bytes that have been received over the channel. If no traffic is occurring over the channel, this measurement will appear as zero. If the channel has not been started since the queue manager was started, no measurement will be available. |
| <b>Status - Bytes Sent</b>     | The count of bytes that have been sent over the channel. If no traffic is occurring over the channel, this measurement will appear as zero. If the channel has not been started since the queue manager was started, no measurement will be available.     |

---

**Note:** In order to enable the event for a queue, ensure that the attributes for the queue have been set. For more information, refer to “Configuring the Server Environment to Monitor Events,” on page 521.

---

- 11** If the event configured for monitoring is from a remote queue manager (other than the one identified in the queue manager field of the IBM WebSphere MQ Add Measurements dialog box), click the **Alternate Queue** button. Enter the name of an alternate queue manager in the Alternate Queue dialog box, and click **OK**.

**Note:** When you add an alternate queue manager, this becomes the default queue manager for any events that you subsequently add. To return to the queue manager to which you are connected, enter that name in the Alternate Queue Manager dialog box.

---

- 12** Click **Add** to add the object measurements to the monitored objects list. The name of object, its events and attributes, and queue managers, are listed in the monitored objects pane.
- 13** To remove a monitored object event or attribute, select the object measurement in the monitored objects pane, and click **Remove**. The entry is deleted from the monitored objects list.
- 14** Add all the desired counters to the monitored objects list, and click **OK**.
- 15** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.

## Configuring the TUXEDO Monitor

The TUXEDO monitor allows you to measure and view your TUXEDO client's performance.

---

**Note:** If TUXEDO 7.1 or higher is installed on the Console machine, more than one TUXEDO application server can be monitored at a time. However, if TUXEDO 6.5 or below is installed on the Console machine, only one TUXEDO application server can be monitored at a time. Use a TUXEDO 6.x client if a TUXEDO 6.x server is used, and TUXEDO 7.1 or above if a TUXEDO 7.1 or above server is used.

---

### Before you set up the monitor:

- 1** Ensure that a TUXEDO workstation client (not a native client) is installed on the Console machine.
- 

**Note:** A TUXEDO workstation client communicates with the application server over the network, and is not required to run the TUXEDO application server on the same machine. A native client can only communicate with the TUXEDO application server if it is part of the relevant TUXEDO domain.

---

- 2** Define the TUXEDO environment variables on the Console machine—set the TUXDIR variable to the TUXEDO installation directory, and add the TUXEDO bin directory to the PATH variable.
- 3** Configure the TUXEDO application server so that the workstation listener (WSL) process is running. This enables the application server to accept requests from workstation clients. Note that the address and port number used to connect to the application server must match those dedicated to the WSL process.

**To configure the TUXEDO monitor:**

- 1** Click **Monitors** to open the Monitors Configuration dialog box.
- 2** Select the server whose monitors you want to configure from the Server list box.
- 3** Click **Add Monitor**. The Select Measurements to Monitor dialog box is displayed.
- 4** In the left section of the dialog box, select TUXEDO (in the Middleware Component category) and then click **Add**.

The TUXEDO Logon dialog box appears, prompting you to enter information about the TUXEDO server: Login Name, Password, Server Name, Client Name.

---

**Note:** This information is located in the Logon section of the *tpinit.ini* file in the recorded script's directory. It is recommended that you use the Browse button and select the *tpinit.ini* file from a recorded script, rather than enter the values manually.

---

To obtain the correct settings for the TUXEDO monitor using the *tpinit.ini* file, click the **Browse** button and navigate to the *tpinit.ini* file of that ProTune script. You can also determine the client name from the **Irt\_tpinitialize** statement in the recorded script.

In the following example of a *tpinit.ini* file, the TUXEDO monitor was configured for a server named URANUS using port 65535, and a client named bankapp. The logon user name was Smith and the password was mypasswd.

```
[Logon]
LogonServername=//URANUS:65535
LogonUsrName=Smith
LogonCltName=bankapp
LogonGrpName=
LogonPasswd=myspasswd
LogonData=
```

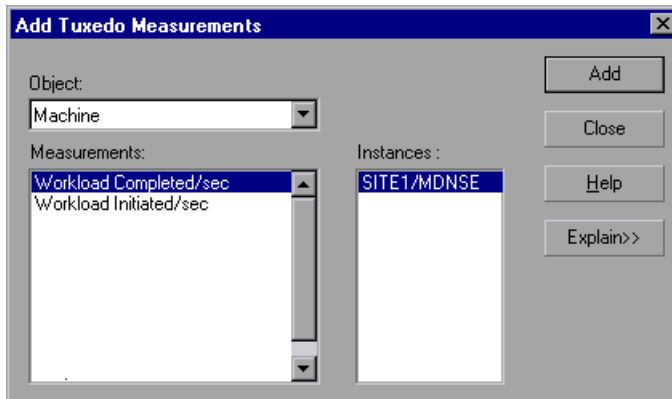
If you already know the required values, you can manually type them into the dialog box. The format of the server name is `//<machine name>:<port number>`. Alternatively, you can specify the IP address instead of the machine name. The hexadecimal format used by old versions of TUXEDO is also supported. Note that quotation marks should not be used.

---

**Note:** If you are using TUXEDO 6.5 or below, the monitor can only connect to one application server during a Console session. Once it connects to an application server, that server is the only one used by the monitor until the Console is closed. This applies even when all of the counters are deleted from the monitor.

---

- 5 Click **OK**. The Add TUXEDO Measurements dialog box opens.



- 6 Select a TUXEDO object from the Object list. Select the measurements and instances you want to monitor. The following table lists the available TUXEDO monitor measurements:

| Monitor | Measurements  |
|---------|---|
| Server  | <b>Requests per second</b> - How many server requests were handled per second   |
|         | <b>Workload per second</b> -The workload is a weighted measure of the server requests. Some requests could have a different weight than others. By default, the workload is always 50 times the number of requests. |
| Machine | <b>Workload completed per second</b> - The total workload on all the servers for the machine that was completed, per unit time  |
|         | <b>Workload initiated per second</b> - The total workload on all the servers for the machine that was initiated, per unit time  |
|         | <b>Current Accessers</b> - Number of clients and servers currently accessing the application either directly on this machine or through a workstation handler on this machine.                                      |
|         | <b>Current Clients</b> - Number of clients, both native and workstation, currently logged in to this machine.   |
|         | <b>Current Transactions</b> - Number of in use transaction table entries on this machine.   |
| Queue   | <b>Bytes on queue</b> - The total number of bytes for all the messages waiting in the queue   |
|         | <b>Messages on queue</b> - The total number of requests that are waiting on queue. By default this is 0.  |

| Monitor                          | Measurements   |
|----------------------------------|--|
| <b>Workstation Handler (WSH)</b> | <b>Bytes received per second</b> - The total number of bytes received by the workstation handler, per unit time  |
|                                  | <b>Bytes sent per second</b> - The total number of bytes sent back to the clients by the workstation handler, per unit time  |
|                                  | <b>Messages received per second</b> - The number of messages received by the workstation handler, per unit time  |
|                                  | <b>Messages sent per second</b> - The number of messages sent back to the clients by the workstation handler, per unit time  |
|                                  | <b>Number of queue blocks per second</b> - The number of times the queue for the workstation handler blocked, per unit time. This gives an idea of how often the workstation handler was overloaded. |

- 7** Click **Add** to place the selected object on the resource list. Add all the desired objects to the list, and click **Close**.
- 8** Click **OK** in the Select Measurements to Monitor dialog box to activate the monitor.



# 30

---

## Application Traffic Management

Using ProTune's Application Traffic Management monitor, you can monitor the BIG-IP load balancing device during a session step run and isolate server performance bottlenecks.

This chapter describes:

- Configuring the F5 BIG-IP Monitor

### About Application Traffic Management Monitoring

ProTune's F5 BIG-IP monitor provides you with information about the content of event logs and other data from F5 BIG-IP load balancing device using SNMP during a session step execution. In order to obtain performance data, you need to activate the online monitor for the server and specify which resources you want to measure before executing the session step.

## Configuring the F5 BIG-IP Monitor

To use the F5 BIG-IP monitor, you must first configure the F5 BIG-IP monitor on the SiteScope machine, and then select the counters you want the F5 BIG-IP monitor to measure. You select these counters using the Console's F5 BIG-IP dialog box.

### Before setting up the F5 BIG-IP monitor:

- 1** Make sure that SiteScope has been installed on a server. SiteScope is the application that is used to monitor the F5 BIG-IP server. Although you can install SiteScope on the Console machine, we recommend installing it on a dedicated server.
- 2** On the machine where SiteScope is installed, configure SiteScope to monitor the required F5 BIG-IP machines. For more information on configuring the SiteScope server, refer to the SiteScope User Guide (<http://www.freshwater.com/SiteScope/UGtoc.htm>).

---

**Note:** When you assign a name to a monitor, include the server name in the monitor name. This avoids any confusion as to which host the monitor belongs.

---

- 3** Verify that SiteScope is collecting the required data from the servers it is monitoring. From the SiteScope Panel, select the monitor group polling the F5 BIG-IP server machines, and check that the monitor displays a list of server measurements in the Status column.

### Configuring the F5 BIG-IP monitor on the SiteScope machine:

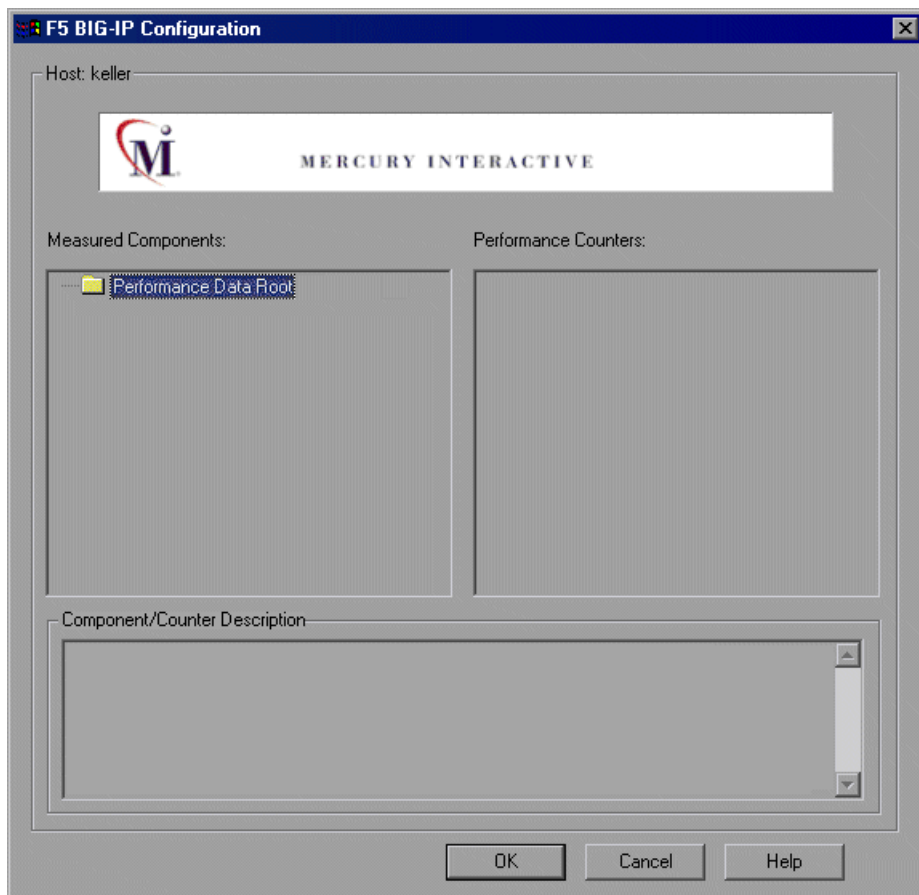
- 1** Open the SiteScope Add Monitors to Group page, and click **F5 BIG-IP Monitor**.
- 2** Click **Choose Server**.
- 3** In the server field, enter the name or IP address of the F5 BIG-IP server that you want to monitor.
- 4** Enter the community for the SNMP object. The default community is public.

- 5 Enter the appropriate value in the Retry Delay field. The default value is 1 second.
- 6 Enter the appropriate value in the Timeout field. The default value is 5 seconds.
- 7 Click **Browse Counters** and select your desired counters.
- 8 Click **Choose Counters**, and then click **Add Monitor**.

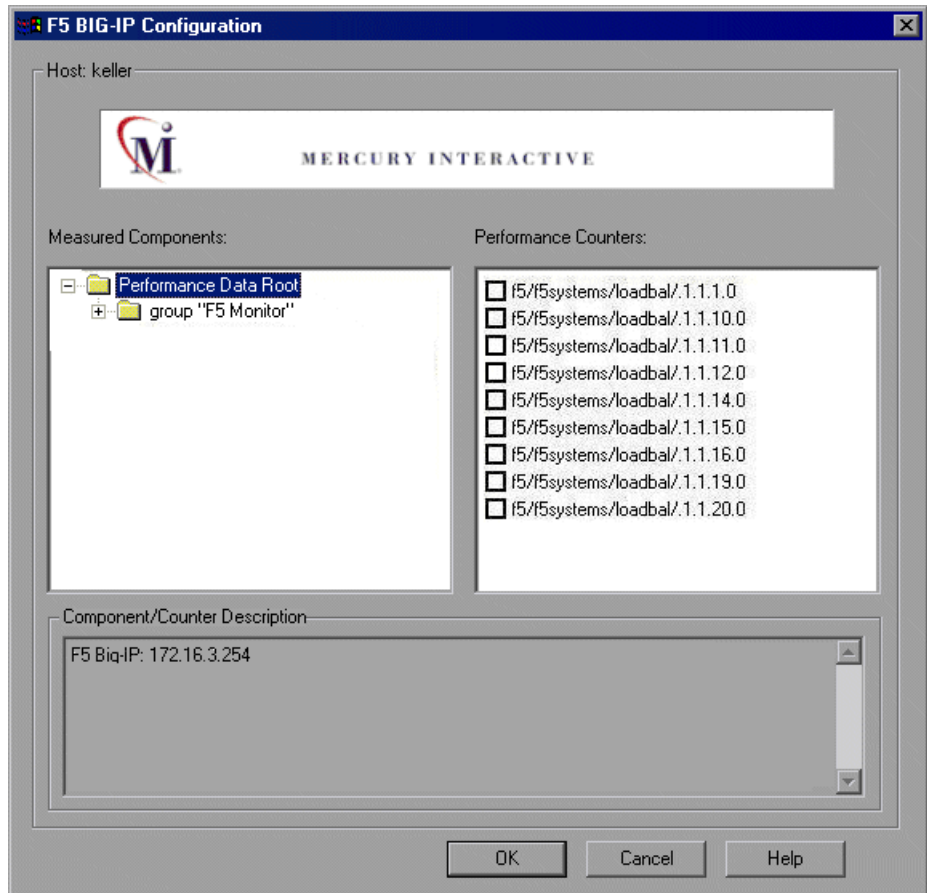
**Configuring the F5 BIG-IP monitor on the Console machine:**

- 1 Click the F5 BIG-IP graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Monitors Configuration**, or click the **Monitors** button on the toolbar. The **Monitors Configuration** dialog box appears.
- 3 From the Server list, choose the server running F5 BIG-IP and click **Add**. The **Select Measurements to Monitor** dialog box appears.

- 4 In the left pane, expand the Load Balancer category, click F5 BIG-IP, and click **Add**. The F5 BIG-IP Monitor Configuration dialog box is displayed.



- 5 In the Measured Components pane, locate the F5 BIG-IP measurement that you are monitoring and click it. The performance counters that F5 BIG-IP is monitoring on the selected component are displayed in the Performance Counters pane.



The following table shows the default counters that can be measured:

| Measurement |
|-------------|
| pktsin      |
| pkcout      |
| concur      |
| portdeny    |
| uptime      |
| droppedin   |
| droppedout  |
| MemoryUsed  |

- 6 Check the required performance counters in the Performance Counters pane.
- 7 When you have selected the performance counters for the F5 BIG-IP measurements you are monitoring, click **OK** to close the F5 BIG-IP Configuration dialog box. The **Select Measurements to Monitor** dialog box appears with the selected F5 BIG-IP measurements in the **Selected Measurements** pane.
- 8 Click **OK** in the **Select Measurements to Monitor** dialog box, and click **OK** in the **Monitors Configuration** dialog box, to activate the F5 BIG-IP monitor.

---

**Note:** For troubleshooting tips and limitations, see “Troubleshooting Server Resource Monitors,” on page 541.

---

# 31

---

## Troubleshooting Online Monitors

ProTune monitors allow you to view the performance of the session step during execution.

The following sections describe several tips and known issues relating to the online monitors.

- Troubleshooting Server Resource Monitors
- Troubleshooting the Network Delay Monitor
- Network Considerations

### Troubleshooting Server Resource Monitors

In order to monitor resources on a server machine, you must be able to connect to that machine. If monitoring is unsuccessful and ProTune cannot locate the specified server, make sure that the specified server is available. Perform a “ping” operation by typing `ping <server_name>` from the Console machine command line.

Once you verify that the machine is accessible, check this table for additional tips on troubleshooting the monitor.

| Problem   | Solution   |
|---|--|
| Cannot monitor a Windows machine on a different domain, or “access denied.” | To gain administrative privileges to the remote machine, perform the following from the command prompt:<br><code>net use \\&lt;MachineName&gt;/</code><br><code>user:[&lt;Domain&gt;\&lt;RemoteMachineUsername&gt;]</code><br>At the password prompt, enter the password for the remote machine. |

| Problem  | Solution   |
|--|--|
| <p>Cannot monitor an NT/Win 2000 machine (An error message is issued: "computer_name not found" or "Cannot connect to the host")</p> | <p>The NT/Win 2000 machine you want to monitor only enables monitoring for users with administrator privileges. In order to allow monitoring for non-admin users, you must grant read permission to certain files and registry entries (Microsoft tech-note number Q158438.) The required steps are:</p> <ol style="list-style-type: none"> <li>a. Using Explorer or File Manager, give the user READ access to: <ul style="list-style-type: none"> <li>%windir%\system32\PERFCxxx.DAT</li> <li>%windir%\system32\PERFHxxx.DAT</li> </ul>                     where xxx is the basic language ID for the system—for example, 009 for English. These files may be missing or corrupt. If you suspect this; expand these files off of the installation cd.                 </li> <li>b. Using REGEDT32, give the user READ access to: <ul style="list-style-type: none"> <li>HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Perflib</li> </ul>                     and all sub keys of that key.                 </li> <li>c. Using REGEDT32, give the user at least READ access to: <ul style="list-style-type: none"> <li>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurePipeServers\winreg</li> </ul> </li> </ol> |
| <p>Some Win 2000 counters cannot be monitored from an NT machine.</p>  | <p>Run the Console on a Win 2000 machine.</p>  |
| <p>Some Windows default counters are generating errors</p>   | <p>Remove the problematic counters and add the appropriate ones using the "Add Measurement" dialog box.</p>  |
| <p>You cannot get performance counters for the SQL server (version 6.5) on the monitored machine.</p>                                | <p>There is a bug in SQL server version 6.5. As a workaround, give read permission to the following registry key at the monitored machine (use regedt32):<br/>                     HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSSQLServer\MSSQLServer<br/>                     (Microsoft tech-note number Q170394)</p>   |



| Problem  | Solution  |
|--|---|
| The selected measurements are not displayed in the graph.                      | Ensure that the display file and online.exe are registered. To register the monitor dll's, without performing a full installation, run the <i>set_mon.bat</i> batch file located in ProTune/bin.                  |
| When monitoring a Windows machine, no measurements appear in the graph.        | Check the built-in Windows Performance Monitor. If it is not functional, there may be a problem with the communication setup.   |
| When monitoring a UNIX machine, no measurements appear in the graph.           | Ensure that an <i>rstatd</i> is running on the UNIX machine (Refer to Chapter 18, "System Resource Monitoring.").   |
| Cannot monitor one of the following Web servers: MS IIS, MS ASP, or ColdFusion | Refer to problem above, "Cannot monitor a Windows machine."   |
| Cannot monitor the WebLogic (JMX) server                                       | Open the <i>ProTuneroot</i> folder\ <i>dat\monitors\WebLogicMon.ini</i> file, and search for:<br>[WebLogicMonitor]<br>JVM=javaw.exe<br>Change javaw.exe to java.exe. A window containing trace information opens. |

## Troubleshooting the Network Delay Monitor

If monitoring is unsuccessful and ProTune cannot locate the source or destination machines, make sure that the specified machines are available to your machine. Perform a “ping” operation. At the command line prompt, type:

```
ping server_name
```

To check the entire network path, use the trace route utility to verify that the path is valid.

For Windows, type `tracert <server_name>`.

For UNIX, type `traceroute <server_name>`.

If the monitoring problem persists once you verify that the machines are accessible and that the network path is valid, perform the following procedures:

1) If you are using the TCP protocol, run `<ProTune root folder\bin\webtrace.exe` from the source machine to determine whether the problem is related to the Console, or the WebTrace technology on which the Network Delay monitor is based. If you are using the UDP or ICMP protocols, the problem must be related to the Console and not WebTrace, since these protocols are not WebTrace technology-based.

2) If you receive results by running `webtrace.exe`, the problem is related to the Console. Verify that the source machine is not a UNIX machine, and contact Mercury Interactive's Customer Support with the following information:

- the Console log file, `drv_log.txt`, located in the temp directory of the Console machine.
- the `traceroute_server` log file, located on the source machine.
- the debug information located in the `TRS_debug.txt` and `WT_debug.txt` files in the path directory. These files are generated by adding the following line to the [monitors\_server] section of the `<ProTune root folder\dat\mdrv.dat` file, and rerunning the Network monitor:

ExtCmdLine=-tracert debug path

3) If you do not receive results by running *webtrace.exe*, the problem is related to the WebTrace technology, on which the Network Delay monitor is based. Perform the following procedures on the source machine:

- Verify that the *packet.sys* file (the Webtrace driver) exists in the WINNT\system32\drivers directory.
- Check whether a driver (such as “Cloud” or “Sniffer”) is installed on top of the network card driver. If so, remove it and run WebTrace again.
- Verify that there are administrator permissions on the machine.
- Using `ipconfig /all`, check that only one IP address is assigned to the network card. WebTrace does not know how to handle multiple IP addresses assigned to the same card (IP spoofing).
- Check the number of network cards installed. Run `webtrace -devlist` to receive a list of the available network cards.
- If there is more than one card on the list, run `webtrace -dev <dev_name> <destination>`, where `<dev_name>` is one of the network card names shown in the list. If you discover that WebTrace is binding to the wrong card, you can use `webtrace set_device <dev_name>` in order to set a registry key that instructs WebTrace to use a specified card instead of the default one.
- Verify that the network card is of the Ethernet type.
- Contact Mercury Interactive’s Customer Support with the output of `webtrace.exe -debug` (for example, `webtrace.exe -debug www.merc-int.com`) and `ipconfig /all` on the machine.

## Network Considerations

If you notice extraordinary delays on the network, refer to one of the following sections to increase the performance:

- Network Bandwidth Utilization
- Ethernet-bus Based Networks
- Working on a WAN or Heavily Loaded LAN

### Network Bandwidth Utilization

In most load-testing session steps, the network card has little impact on session step performance. Network cards are manufactured to handle the bandwidth of the physical network layer. Packets are transferred over an Ethernet at a rate that complies with IEEE 803.x standards. If the network becomes a bottleneck, the issue is not the brand of the network card, but rather the bandwidth limitations on the physical layer (--i.e. Ethernet, FDDI, ATM, Ethernet Token-ring, etc.).

That is, instead of load testing over a T10 line, upgrade your line to DS3 (45Mbps), or T100 (100Mbps).

Below are a few tips that will help qualify the need to upgrade the network:

- 1) Run the performance monitor on the Vuser load generators. As the number of Vusers increases, check the network byte transfer rate for saturation. If a saturation point has been reached, do not run any more Vusers without upgrading the network—otherwise performance of Vusers will degrade. Degradation is exponential in networking environments.
- 2) Run the performance monitor on the server machine. Run many Vusers on several load generator machines. Check the kernel usage and network transfer rate for saturation. If saturation is reached with less than the desired Vuser load, upgrade the network.
- 3) Every network has a different Maximum Transmission Unit or MTU, which is set by the network administrator. The MTU is the largest physical packet size (in bytes) that a network can transmit. If a message is larger than the MTU, it is divided into smaller packets before being sent.

If clients and servers are passing large data sets back and forth, instruct the network administrator to increase the MTU in order to yield better bandwidth utilization. Ideally, you want the MTU to be the same as the smallest MTU of all the networks between your machine and a message's final destination.

If you send a message that is larger than one of the MTUs, it will be broken up into fragments, slowing transmission speeds. If the MTU is too high, it may cause unintended degradation. Trial and error is the only sure way of finding the optimal MTU, but there are some guidelines that can help. For example, most Ethernet networks have an MTU of 1500.

If the desired MTU reduces performance, upgrade the network or reduce the MTU to improve performance.

### **Ethernet-bus Based Networks**

The following guidelines apply to Ethernet-bus based networks:

Networks with only 2 active machines communicating yield a maximum of 90% bandwidth utilization.

Networks with 3 active machines communicating yield a maximum of approximately 85% bandwidth utilization.

As the number of active machines on the network increases, the total bandwidth utilization decreases.

### **Working on a WAN or Heavily Loaded LAN**

When you work with ProTune on a WAN or heavy loaded LAN, you may notice some unusual ProTune behavior, which indicates network problems. The Output window may contain messages about retries, lost packets, or message mismatch. This is because some of the messages from the Console may not be reaching the ProTune agent. To solve this problem, you should reduce the network traffic or improve the network bandwidth.

The following steps may help reduce network traffic:

- Click the **Run-Time Settings** button. In the Log tab, select **Disable logging**.
- Initialize all users before running them. Run them only after initialization is completed.

# Part VI

---

## Tuning Your System





# 32

---

## Tuning Your System from the Console

Once you've run session steps on your topology and analyzed the test results, you use the Console to administer and tune your hosts and services from a remote location. You continue this process until your system reaches optimal performance.

This chapter includes the following topics:

- Supported Operating Systems
- Applications That ProTune Can Tune
- Tuning Flow
- Configuring Host Connection Parameters
- Connecting to the Host Computer
- Viewing the Host Information
- Resynchronizing the Information Tab
- Using Expert Mode
- Changing Tuning Parameter Values
- Updating the Host or Service with Changes
- Configuring Special Tuner Agent Settings

## About Tuning Your System from the Console

ProTune's tuning features allow you to remotely tune hosts and services, from the Console machine. After ensuring that you have the necessary permissions and access rights on the host, you install a tuning agent—a small application—on the host. The tuning agent allows the Console to view the host's services, and to configure their settings.

ProTune's **Tune** tab contains the following:

- ▶ **Server Configuration:** A tree structure listing the hosts and services you can tune.

---

**Note:** If a host is used only as a load generator, it will not appear in the Tune tab.

---

- ▶ **Information tab:** Describes the properties of the hosts and services you select in the Server Configuration tree.
- ▶ **Tuning tab:** Allows you to change the configurable properties of the hosts and services.
- ▶ Additional buttons that help you tune your system.
- ▶ Links to help topics that explain and guide you through the performance tuning process.

## Supported Operating Systems

You can install tuning agents on computers running the following operating systems:

**Windows:** Windows NT, Windows 2000, Windows XP

**UNIX:** Solaris, HP, AIX, Linux

## Applications That ProTune Can Tune

ProTune allows you to tune the following applications:

- Apache Web Server 1.x/2.x
- BEA Weblogic 6.x/7.x
- IBM HTTP Server
- IBM Websphere Advanced 4.x
- IBM Websphere Single Server 4.x
- iPlanet Enterprise Server 6 and higher
- Microsoft IIS 4/5
- Microsoft Active Server Pages 2/3
- Operating System (Windows: NT, 2000 and XP; UNIX: Solaris, HP, AIX and Linux)
- Oracle Database
- Oracle 9iAS
- PeopleSoft 8.x
- SAP Enterprise Portals 5
- Siebel 7.x
- SQL Server 7.5/2000

## Tuning Flow

You tune each host in your system by following this procedure:

- 1** Configuring Host Connection Parameters. This involves specifying the settings that allow you to connect to the host machine and view its configuration information.
- 2** Connecting to the Host Computer. This also installs and starts a tuning agent. The tuning agent is an application that runs on the host machine, and communicates with the Console. You can install the tuning agent remotely from the Console machine or locally on the host.
- 3** Viewing the Host Information. After you connect to the host machine, you can view the information that is returned by the tuning agent.
- 4** Changing Tuning Parameter Values. You change the values of the relevant host settings to enhance the host machine's performance.
- 5** Updating the Host or Service with Changes. This applies the new settings to the host machine.

## Host Requirements

Before you install and run a tuning agent on the host machine, note the following requirements:

### Java

The tuning agent requires a Java-enabled environment. Ensure that JRE/JDK 1.3 (or later) is installed on the host before you attempt to start the tuning agent. Either JRE or JDK can be used.

### WMI Support

If the Console machine runs Windows NT, it must have Windows Management Instrumentation (WMI) support installed.

In addition, to enable remote installation of a tuning agent on a Windows NT host machine, you need to install WMI support on the host.

To install WMI, download it from

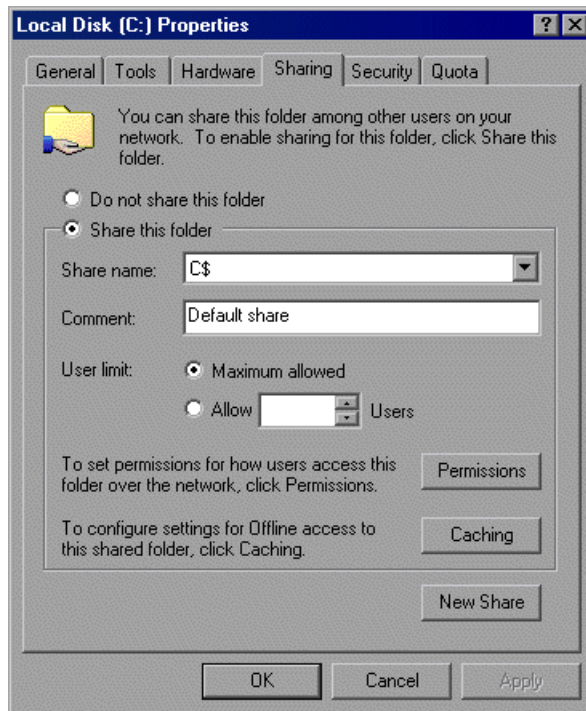
<http://msdn.microsoft.com/downloads/default.asp?url=/downloads/sample.asp?url=/msdn-files/027/001/576/msdncompositedoc.xml>

## Drive Sharing

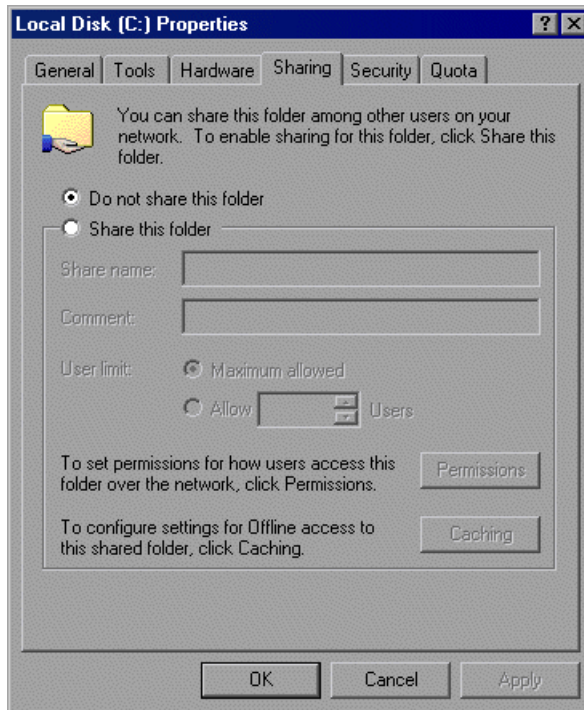
If the host that you are tuning is a Windows machine, ensure that its C: drive is shared with other users by default for administrative purposes.

### To share the C: drive by default:

- 1 In Windows Explorer, right-click the C: drive and choose Sharing. The following dialog box appears. If the C: disk is already shared, the Share Name and Comment fields appear as follows:



- 2 If the C: disk is not shared by default, the **Share this folder** section is disabled and the Share Name and Comment fields are empty.



If this is the case, click **Share this folder** and enter C\$ in the **Share name** list.

- 3 Click **OK**.

## Configuring Host Connection Parameters

Before you install the tuning agent on the host machine that you want to tune, you specify the host and configure its connection parameters.

You need to perform the following actions:

- 1** Choosing the Host Machine
- 2** Specifying Tuner Agent Settings
- 3** Specifying Operating System Settings

---

**Note:** In addition, you can optionally specify SNMP and F5 BIG-IP iControl settings.

---

- 4** Connecting to the Host Computer

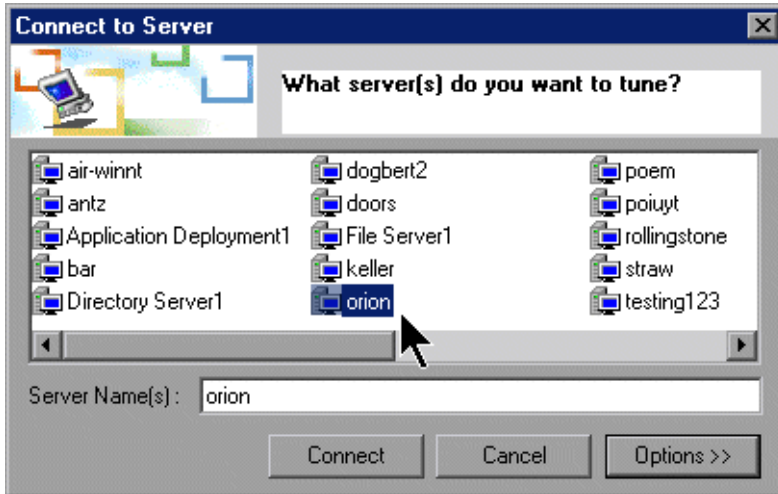
## Choosing the Host Machine

To choose the host that you want to tune:



- Click the **Connect to Host** button, or right-click the **Server Configurations** link on the left side of the window and click **Add New**.

The Connect to Server dialog box is displayed.



## Specifying Tuner Agent Settings

Before you can connect to the host machine, you need to specify the tuner agent's settings. This allows you to start and stop the tuning agent, and to use it to tune the host machine.



### To specify the tuner agent settings:

- 1 In the Connect to Server dialog box, click **Options**. Alternatively, for a host that you have already added to the Server Configurations tree (for example, if you added the host machine to your topology), click **View Host Properties**, or right-click the host's icon and choose Properties. The Host Settings dialog box is displayed.



**Host Settings** X

**poem**

Tuner Agent **OS** Operating System Integrations Snmp

Tuner Agent Options

Use following authentication information for agent(s)

Username :

Password :

Use Http proxy to establish connection to agent(s)

Proxy Address :

Username :

Password :

Use Secure Socket Layer (SSL)

Connect to tuning agent using user-defined port address :

Description :

OK Cancel

- 2** In the **Tuner Agent** tab, enter the tuner username and password. This determines the actions that the tuner allows you to perform on the host. ProTune encrypts the password. You can select a username and password defined by the system administrator, or choose one of the following predefined username/password pairs:

| Username | Password      | Authorization              | Comments   |
|----------|---------------|----------------------------|--|
| guest    | (no password) | viewing access only        | default  |
| mercury  | expert        | viewing and update access  |  |
| admin    | changeit      | full administration access | The administrator of the tuner can define and grant tuning privileges to a user. |

If you do not check the **Use following authentication information...** box, ProTune uses the *guest* username, allowing you only to query the host.

---

**Note:** For details on how to change passwords, see “Changing Tuning Agent Passwords,” on page 592.

---

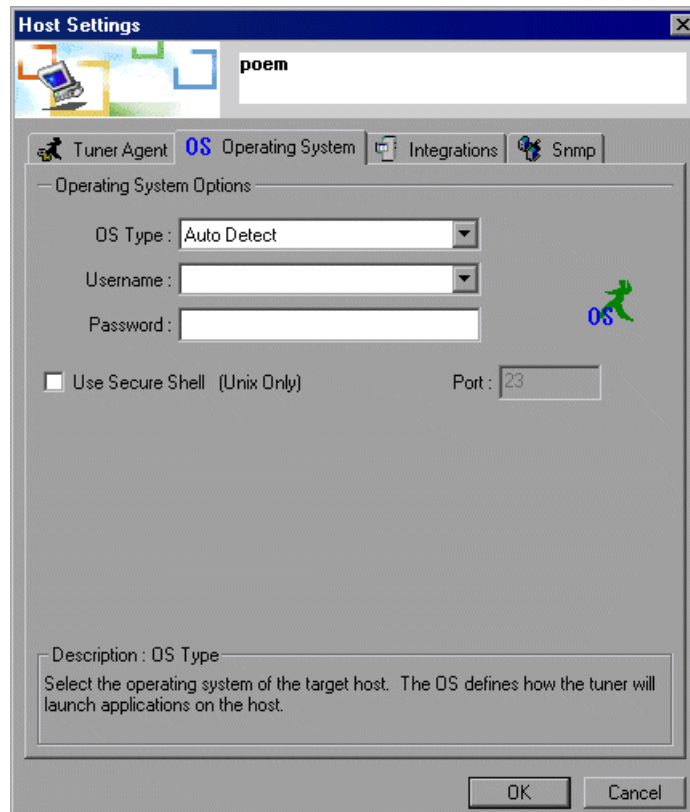
- 3** If the tuning agent on the host is accessible only through a proxy over a firewall, define the proxy settings in the **Use Http proxy...** dialog box.
- 4** If the agent is running (or will run) on a non-standard port (OTP-SSL 4863 or OTP 4862), you can specify a user-defined port. Check the **Connect to tuning agent** box and enter the port address in the relevant field.
- 5** Check the **Use Secure Socket Layer (SSL)** box to use SSL for all connections (recommended).

## Specifying Operating System Settings

You need to specify the host machine's operating system and username/password information, so that the host machine will allow you to access it via the tuner agent.

**To specify the operating system settings:**

- 1 In the Connect to Server dialog box, click **Options**. Alternatively, for a host that you have already added to the Server Configurations tree (for example, if you added the host machine to your topology), click **View Host Properties**, or right-click the host's icon and choose Properties. The Host Settings dialog box is displayed.
- 2 Click the **Operating System** tab.



- 3 Choose the host's operating system from the list in the OS Type box.

- 4** Enter the host's username (including the domain) in the Username field. For example, if your computer's name is **straw** and your username is **joe**, enter the string **straw\joe**.
- 5** Enter the host password in the Password field.
- 6** To use secure shell (for UNIX only), check the **Use Secure Shell** box and enter the port address in the Port field.

---

**Note:** After you have specified the tuning agent and operating system settings (and, optionally, F5 BIG-IP iControl and SNMP settings), you need to connect to the host computer. Follow the instructions in the next section (Connecting to the Host Computer).

---

## Connecting to the Host Computer

After configuring the tuning agent settings, you need to connect to the host computer.

**To connect to a newly-added host computer:**

- 1** In the Host Settings dialog box, click **OK** to save your settings and close the dialog box.
- 2** In the Connect to Server dialog box, ensure that the host appears in the Server Name field.
- 3** Click **Connect**.

This installs the tuning agent on the host machine and starts it.

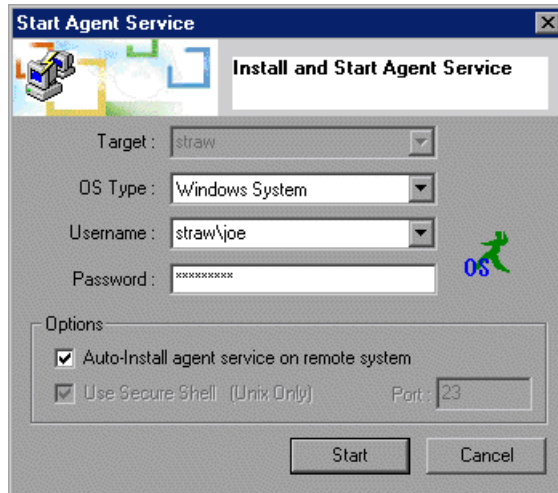
If you added a host machine to the Server Configurations tree but did not start the tuning agent (for example, if you added the host machine to your topology), use one of the following procedures to install and start the tuning agent.

**To remotely install a tuning agent from the Console machine:**

- 1 Click the host's icon in the Server Configurations tree, and then click the **Start Tuning Agent** button on the toolbar.



The Start Agent Service dialog box is displayed:



- 2 Specify the host's OS Type. To let ProTune automatically detect the operating system, choose Auto Detect.
- 3 Check the **Auto-Install...** box.
- 4 If you are installing the tuning agent on a UNIX machine, check the **Use Secure Shell** box and specify the port address, if applicable.
- 5 Click **Start**. If you've configured your settings correctly, ProTune installs the tuning agent on the host. On the host, the tuning agent opens a command window on which you can view the installation activity and the Performance Daemon.

Following is an example of the command window showing the tuning agent's activity on a Windows host when starting:

```
C:\WINNT\system32\cmd.exe
[AGENT]
[AGENT] Performance Daemon Lite (v. 1.1)
[AGENT] Date: Sun Jan 19 14:35:45 GMT+02:00 2003
[AGENT]
[AGENT] Agent Base Directory : C:\Program Files\Mercury Interactive\
Performance Expert\agent
[AGENT] Agent Log File      : agent-jan-19-2003-14-35-45.log
[AGENT] Loading Access Control List...
[AGENT] Secure Service started @ [SSL: ServerSocket[addr=0.0.0.0/0.0
.0.0,port=0,localport=4863]]
[AGENT] Ready!
```

On the Console machine, the host's icon in the Server Configurations tree changes to blue, indicating that the connection to the host is alive.

---

**Tip:** If the Console displays an error message telling you that it is unable to establish a connection, right-click the host's icon and choose Refresh.

---

The command window shows you the location of the tuning agent log file. (In the example above, the log file is located in c:\Program Files\Mercury Interactive\Performance Expert\agent\logs.) The log file keeps track of the following activity occurring during the tuning session:

- Client connections
- Client requests (get, set, start, stop)
- Parameters that are modified, and their new values
- Errors

Following is an example of the log file:

```

agent-jan-19-2003-14-35-45.log - Notepad
File Edit Format Help
[AGENT ] -----
[AGENT ] Performance Daemon Lite (v. 1.1)
[AGENT ] Date: Sun Jan 19 14:35:45 GMT+02:00 2003
[AGENT ] -----
[AGENT ] Agent Base Directory : C:\Program Files\Mercury Interactive\Perform
[AGENT ] Agent Log File      : agent-jan-19-2003-14-35-45.log
[AGENT ] Loading Access Control List...
[AGENT ] Configuring Secure Service...
[AGENT ] Secure Service started @ [SSL:
ServerSocket [addr=0.0.0.0/0.0.0.0,port=0,localport=4863]]
[AGENT ] Ready!
[CLIENT ] New client connection @ a6d51e[SSL_RSA_WITH_3DES_EDE_CBC_SHA:
Socket [addr=/192.168.80.72,port=38535,localport=4863]] [Sun Jan 19 14:35:54 GMT+02:00
[REQUEST ] Request performance.tuner.os.Handler::get (by: guest)
[REQUEST ] Request performance.tuner.Registry::list (by: guest)
[REQUEST ] Request performance.tuner.apache.Handler::get (by: guest)
[REQUEST ] Request performance.tuner.iisasp.Handler::get (by: guest)
[REQUEST ] Request performance.tuner.os.Handler::get (by: guest)
[REQUEST ] Request performance.service.Listener::shutdown (by guest) DENIED! I
privileges! Request denied!

```

If the tuning agent does not start, or doesn't show the expected information, you may need to configure its settings. See "Configuring Tuning Agents," on page 591.

**To install a tuning agent locally from the ProTune CD on a Windows host:**

- Choose the **Tuner Agent** action in the installation procedure.

---

**Note:** After you install the tuning agent locally from the CD, you can start the agent by choosing

**Start > Programs > Performance Expert > Tuning Agent** on the host machine.

---

**To install a tuning agent locally on a Windows host:**

- 1** Extract the `perfgent.tar` file (located in the `\ProTune\console\bin` directory) to a directory (for example, to `C:\Program Files\Mercury Interactive\Performance Expert`).
- 2** Set the `PE_HOME` environment variable (in this example, `PE_HOME = C:\Program Files\Mercury Interactive\Performance Expert`).

- 3** Launch the tuning agent batch file. In this example, you would run the following file:

```
C:\Program Files\Mercury Interactive\Performance Expert\agent\bin\pe_agent.bat
```

You can skip step 2 by passing the path to the installation directory to `pe_agent.bat` in the command line, as in the following example:

```
% pe_agent.bat 0 true C:\Program Files\Mercury Interactive\Performance Expert
```

**To install a tuning agent locally on a UNIX host:**

- 1** Copy the `perfagent.tar` file (from the `\ProTune\console\bin` directory on the Windows machine where the Console is installed) to a directory on the UNIX machine (for example, to `/usr/local/perfexpert`), and extract it to that directory.
- 2** Set `PE_HOME` environment variable (in this example, `setenv PE_HOME /usr/local/perfexpert -- for CSH, ...` )
- 3** Launch the tuning agent batch file. In this example, you would run the following file:

```
/usr/local/perfexpert/agent/bin/pe_agent
```

You can skip step 2 by passing the path to the installation directory to `pe_agent` in the command line, as in the following example:

```
% pe_agent 0 true /usr/local/perfexpert
```

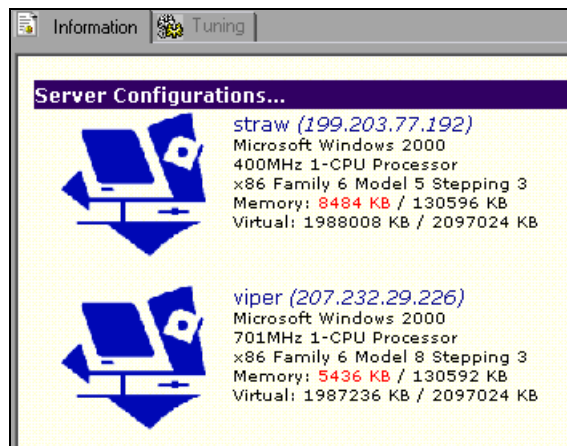


## Viewing the Host Information

The Tune window contains two tabs:

- ▶ The **Information** tab shows you the host property information.
- ▶ The **Tuning** tab allows you to change the values of those properties that are configurable.

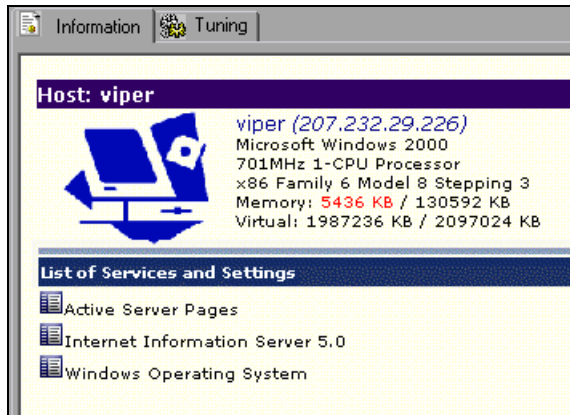
After you've connected to the host and the tuning agent is running, click the Server Configurations element. The **Information** tab shows summary information for each of the hosts, as in the following example:



To view detailed information about a host, click the host's icon in the Server Configurations tree. The **Information** tab then displays the following information about the host:

- ▶ Hostname and IP address
- ▶ CPU type, including number of processors
- ▶ Memory—available RAM over total RAM. The virtual memory information shows available memory over total virtual memory (pagefile).
- ▶ Services and settings

Following is an example of what the **Information** tab shows when you click a host icon:



Each host in the Server Configurations tree contains a list of the running services. Sub-elements of each service are divided into categories and sub-categories, as defined by the agent on the remote host machine.

**To view information about a host or host-related service:**

- 1** Click the host or service in the tree; ProTune displays the information in the **Information** tab.
- 2** Click a host to see the information about all of its services and settings.
- 3** Click a service to see only information about the selected host, its categories and its sub-categories.



4 To view a node's sub-elements, click the node's icon to expand it.

The screenshot shows a console window with a header bar for 'Host: viper'. Below the header, there is a blue icon of a laptop and a mouse. To the right of the icon, the following system information is displayed:

```
viper (207.232.29.226)
Microsoft Windows 2000
701MHz 1-CPU Processor
x86 Family 6 Model 8 Stepping 3
Memory: 5436 KB / 130592 KB
Virtual: 1987236 KB / 2097024 KB
```

Below this information is a section titled 'List of Services and Settings' with a blue header bar and a yellow background. It contains a list of services and their settings, each with a small blue icon to its left:

- Active Server Pages
- Internet Information Server 5.0
  - Listen Backlog 25
  - Bandwidth Throttle 0xFFFFFFFF
  - Memory Cache Size 13372620
  - Max Pool Threads 10
  - WWW Max Connections 1000
- Windows Operating System
  - File System Parameters
    - IO Page Lock Limit 0
    - Large System Cache OFF
    - Second Level Data Cache 0
    - System Pages 0
  - Networking Parameters
    - Tcp Timed Wait Delay 240
    - Keep Alive Interval 1000
    - Keep Alive Time 7200000
    - Maximum User Ports 5000
    - Tcp Window Size 8760
  - 3Com EtherLink XL 10/100 PCI TX NIC (3C905B-TX)

5 To collapse a node, click its icon again.



6 To expand all the nodes, click the **Expand All** button.



7 To collapse all the nodes, click the **Collapse All** button.

## Viewing Windows Services

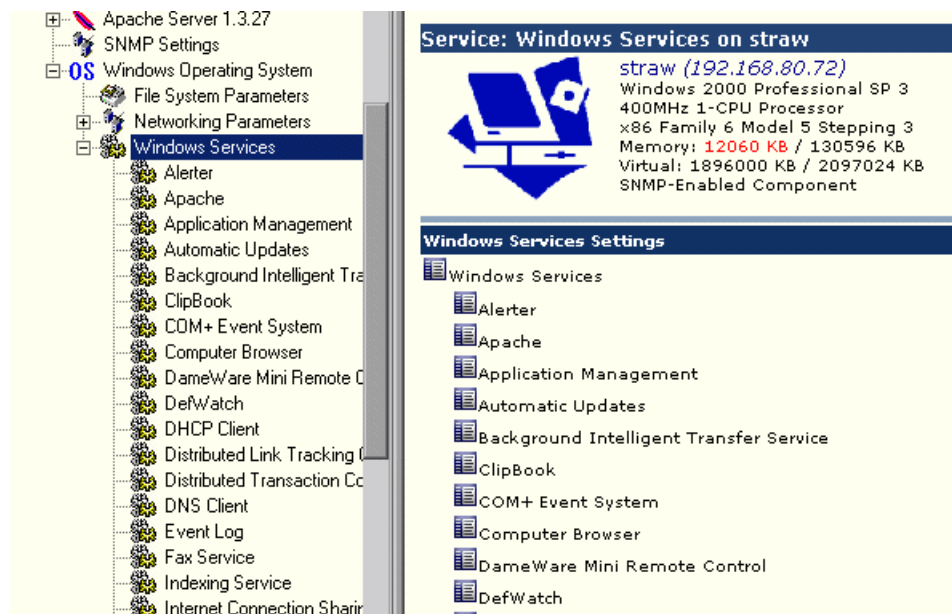
When the host machine is running Windows, a Windows Services element appears as a sub-element of the Windows Operating System element. You should use this element to stop and start the Services settings on hosts with Windows, instead of using the File System Parameters element.

---


**Note:** For details of how to start and stop a Windows service, see “Tune Tab Functions,” on page 607.

---

Expanding the Windows Services element displays the host machine's services.



When you click a service in the Windows Services tree, ProTune displays the service's status in the Information tab.



The following service properties are displayed:

- Status—**Running**, **Stopped** or **Paused**
- Startup Type—**Automatic** (starts automatically when Windows is started), **Manual** (must be started manually by the user) or **Disabled**.

---

**Note:** You can change the startup type via the Tuning tab (see “Changing Tuning Parameter Values,” on page 572).

---

- Path—The path to the service.

## Resynchronizing the Information Tab

If the **Information** tab stops responding, you need to resynchronize.

**To resynchronize the Information tab:**

- Right-click **Server Configuration** and choose **Resync Views**.

## Using Expert Mode

Some tuning parameters are displayed only if you are in Expert mode.

### To enable Expert mode:

- Click the icon of the host or service for which you want to enable it, and then choose **Expert** from the box in the **Tune** tab's toolbar. Alternatively, you can right-click the host or service and choose Expert Mode. Expert mode is enabled for the selected node and its sub-nodes, and the extra parameters are displayed.

### To disable Expert mode:

- Choose Normal from the box in the toolbar, or right-click the host or service and choose Expert Mode. This disables Expert mode for the selected host or service.

## Changing Tuning Parameter Values

You use the **Tuning** tab to change values of tuning and configuration parameters for the selected host or service. For each parameter, ProTune displays the following:

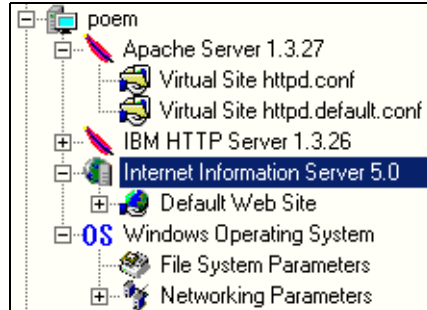
- Current value
- Recommended value—the value recommended by ProTune
- New value—the value that you entered (if you changed the old value)

The different parameters are color-coded as follows:

| Color | Meaning                   |
|-------|---------------------------|
| blue  | important                 |
| black | for advanced users        |
| red   | critical tuning parameter |
| gray  | read-only                 |

### To change the value of a host parameter:

- 1 Select the host or service that you want to tune, by clicking the relevant icon in the Server Configuration tree.



In the **Tuning** tab, ProTune displays the parameters relevant to the selected host or service.

| Service / Attribute Name        | Current Value | Recommended | New Value |
|---------------------------------|---------------|-------------|-----------|
| Internet Information Server 5.0 |               |             |           |
| Listen Backlog                  | 25            | 200         |           |
| Bandwidth Throttle              | 0xFFFFFFFF    | 0xFFFFFFFF  |           |
| Memory Cache Size               | 13372620      |             |           |
| Max Pool Threads                | 10            |             |           |
| WWW Max Connections             | 1000          |             |           |

The values in the Recommended column are based on information in ProTune's knowledge base.

When you click a property, the lower section of the **Tuning** tab displays a description of the property and its values, and may include tuning recommendations.


Description : Bandwidth Throttle

This specifies the amount of bandwidth on the network that the IIS server can use. The server comes with an automatic bandwidth throttler that makes intelligent decisions to meet user requirements. This helps to reduce overloading the network with IIS server activities. For administrators of small corporate servers, where a single server is used for multiple sites, this helps to reduce network usage for IIS servers.

A special value of 0xFFFFFFFF means that you should not do any throttling.

Default Value: 0xFFFFFFFF  
Current Value: 0xFFFFFFFF

- 2 Click the parameter that you want to configure, and click the parameter's New Value column. If the parameter is configurable, a text box opens, or a list box appears.



- 3 Enter the new value. In the case of a list box, choose it from the list.

## Updating the Host or Service with Changes

Changes that you make to parameter values do not take effect until you update the relevant host or service.

When you update a host, all of the services in the host's tree are updated (if any of the parameters have been assigned new values).


When you update a service or service category, only the selected service or category is updated.



To update a host or service:



- 1 Click **Commit Host Changes**, or right-click the host or service icon in the Server Configurations tree and click **Update**. The Update Service Confirmation dialog box is displayed.



- 2 Click **Update** to update the selected host or service. Note that some services need to be restarted for the changes to take effect.  
ProTune displays a message informing you that the changes have taken effect.
- 3 To view the changed values, right-click the host icon and click **Refresh**.

## Configuring Special Tuner Agent Settings

You can configure the following special types of tuner agent settings:

- ▶ F5 BIG-IP iControl Integration Settings
- ▶ SNMP Settings

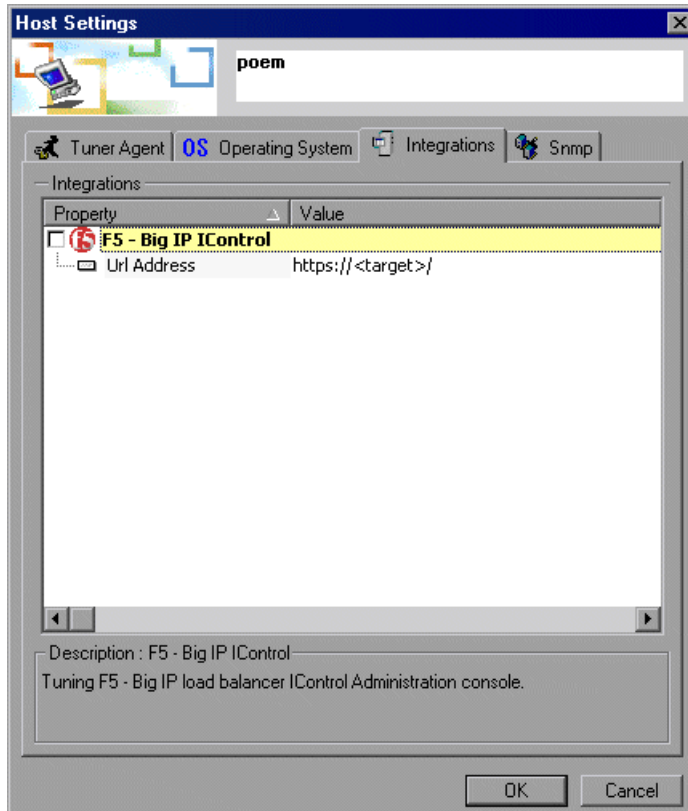
### Enabling F5 BIG-IP iControl Integration

ProTune allows you to integrate with applications that use a Web interface to configure network components. This functionality currently gives you access to the F5 BIG-IP iControl Administration console.

**To access the F5 BIG-IP iControl Administration console:**

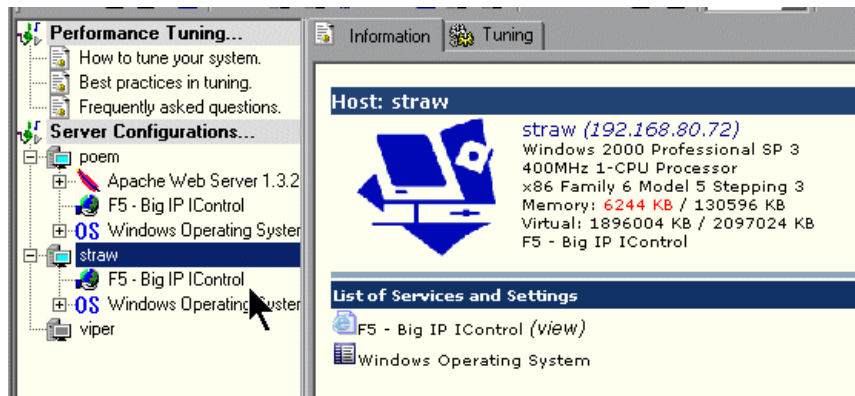


- 1 In the Connect to Server dialog box, click **Options**. Alternatively, for a host that you have already added to the Server Configurations tree, click **View Host Properties**, or right-click the host's icon and choose Properties. The Host Settings dialog box is displayed.
- 2 Click the **Integrations** tab. ProTune displays the F5 BIG-IP iControl properties. In the Property column, check **F5 BIG-IP iControl**.

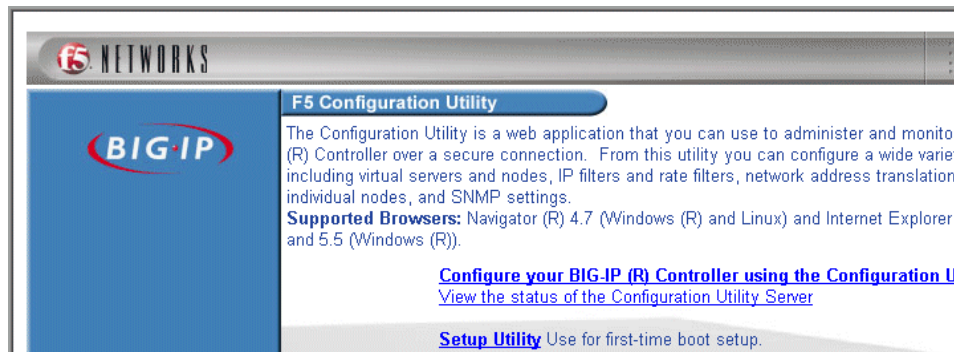


- Click the Url address property's Value section, and replace <target> with the address of the iControl Administration console (for example, username:password@hostname). **Note:** If you use the string "<target>" instead of the hostname, the tuner automatically uses the appropriate hostname. For example, if the hostname is "wizard", specifying the address as https://admin:password@<target> causes the tuner to translate the address to https://admin:password@wizard.

After you apply your changes and exit the Host Settings dialog box, the F5 BIG-IP iControl service is added to your host in the Server Configurations tree.



Clicking the F5 BIG-IP iControl service in the tree causes the iControl Administration console open in the **Information** tab, as in the following example:



You now have access to all the iControl Administration console functions.

## **Enabling SNMP For Tuning**

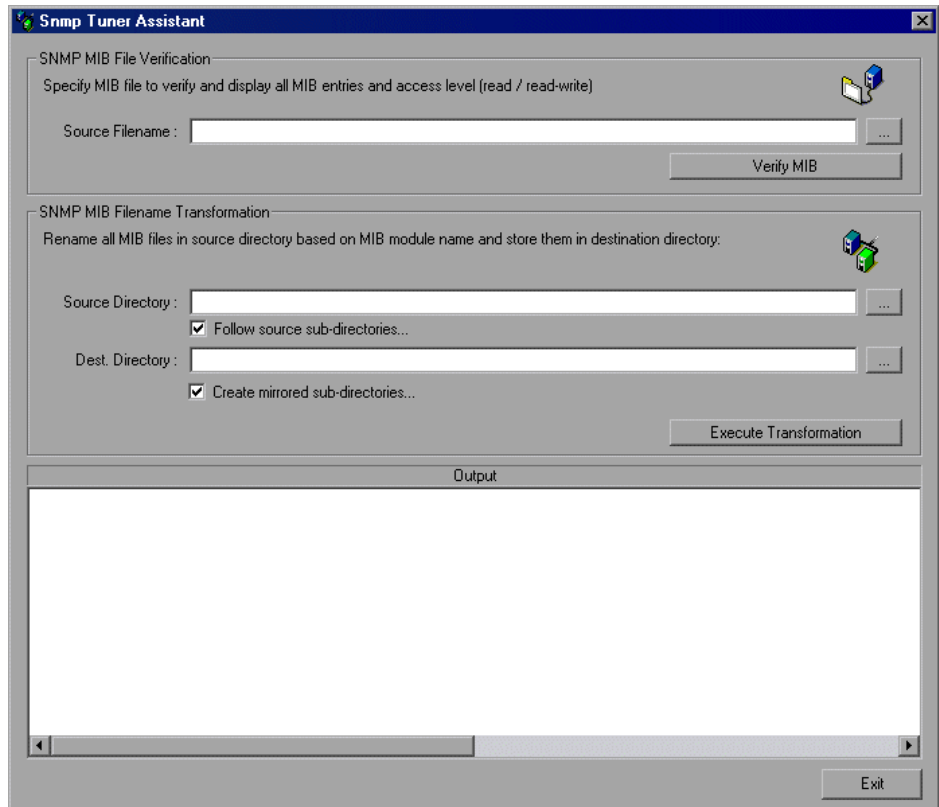
ProTune allows you to use SNMP to tune host machines. Note that the SNMP service must be running on the host machine.

### **To enable SNMP for tuning a host machine:**

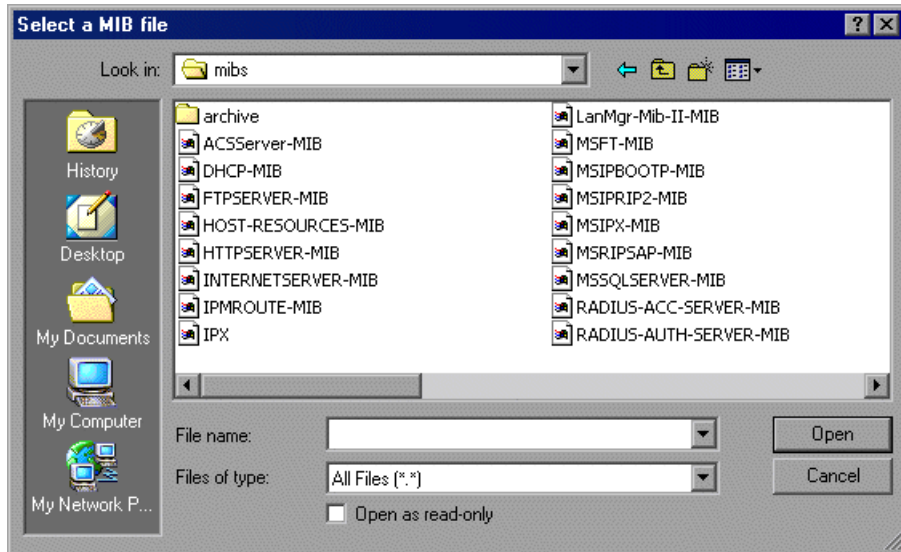
- 1** The ProTune installation directory contains zip files for connecting to different applications via SNMP. These files are located in the `\console\dat\snmp\mibs\archive` subdirectory.

Locate the zip file that applies to the applications to which you want to connect, and extract it to the `\console\dat\snmp\mibs` directory. (For example, if you want to connect to Microsoft applications, extract the file `Microsoft.zip`.) Ensure that you do not create a directory when extracting the file.

- 2 In the console\bin directory, double-click the file snmpasst.exe. The SNMP Tuner Assistant dialog box is displayed.



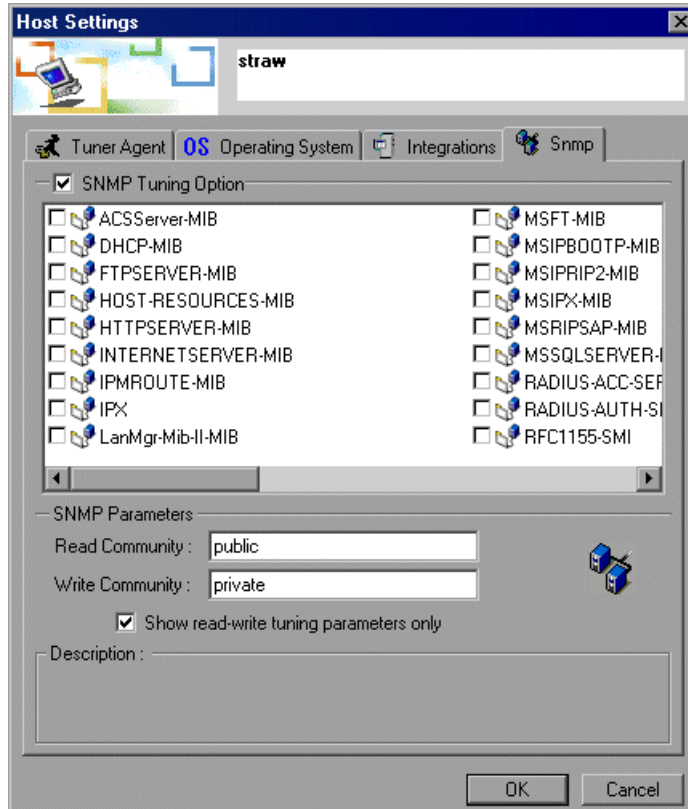
- 3 Using the browse button at the right of the Source Filename field, browse to the console\dat\snmp\mibs directory. The list of .MIB files in the directory is displayed.



- 4 Select the file for the protocol that you want to use, and click **Open**. The filename now appears in the Source Filename field.
- 5 Click **Verify** to check that the MIB file has not been corrupted.
- 6 If error messages are displayed in the Output window, edit the .MIB file and correct the errors. Then click **Verify** to check the file again.
- 7 When the file has been successfully verified, click **Exit** to close the SNMP Tuner Assistant.

**To specify SNMP settings:**

- 1 In the Connect to Server dialog box, click **Options**. Alternatively, for a host that you have already added to the Server Configurations tree, click **View Host Properties**, or right-click the host's icon and choose Properties. The Host Settings dialog box is displayed.
- 2 Click the **SNMP** tab, and check **SNMP Tuning Option**.



- 3 Select the MIB files that you need for tuning the host machine via SNMP. For example, if you want to tune the host machine's MS-SQL Server, select MSSQLSERVER-MIB.
- 4 In the Read Community field, specify the community to use for retrieving values from the server. Default: public.

- 5 In the Write Community field, specify the community to use for writing new values to the server. Default: private.

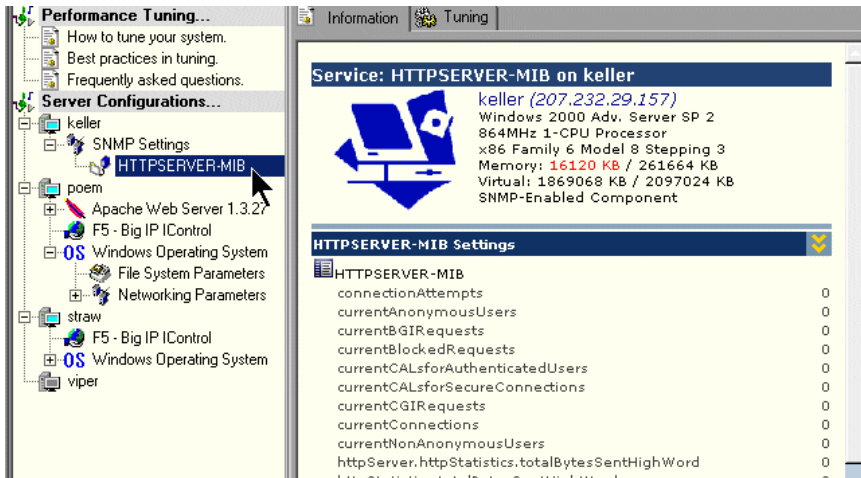
---

**Note:** If a service is not available through the default public and private communities, you need to specify the appropriate community. Please check the server's SNMP configuration, or check with your administrator, for the correct read and write communities.

---

- 6 To cause the **Tuning** tab to display only changeable tuning parameters, check **Show read-write tuning parameters only**.

After you apply your changes and exit the Host Settings dialog box, the SNMP Settings service is added to your host in the Server Configurations tree, with the .MIB file that you chose.





# 33

---

## Exporting and Importing Configuration Settings

This chapter describes how to export and import configuration settings.

It includes the following topics:

- Exporting a Host or Service's Configuration Settings
- Importing Configuration Settings for a Host or Service
- Saving and Loading Profiles
- Creating a New Profile

### About Exporting and Importing Configuration Settings

Once you've viewed (and possibly changed) configuration settings on your hosts and services, you can export the settings to save them for use in the future. You can also import previously saved settings to hosts and services.

Exporting settings allows you to track changes on the hosts and services, and to import the settings into other devices.

Importing previously saved settings for a host or service allows you to replicate desired settings across multiple hosts with similar configurations, avoiding the need to manually update the settings on each host or service individually.

You can export and import settings for a service, a host, or a group of hosts. When you export the settings of a group of hosts, you create a profile. You can subsequently import the profile.

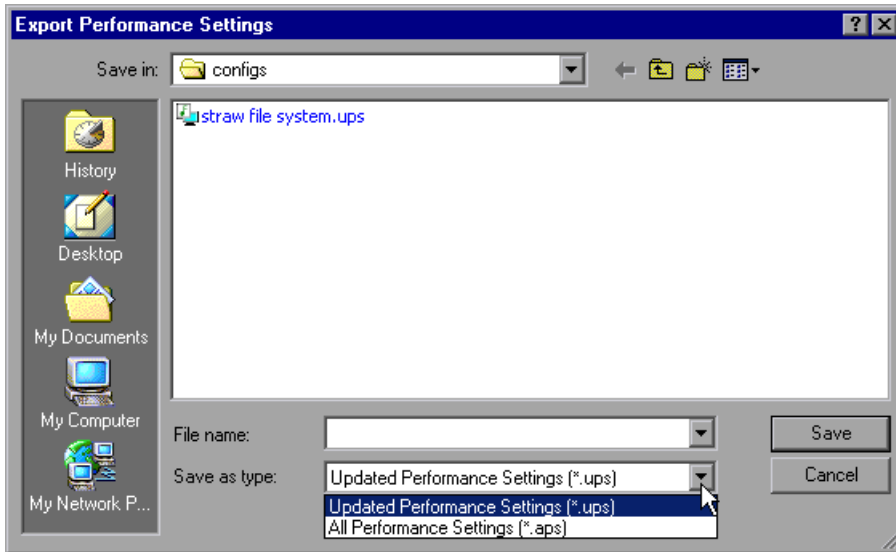
## Exporting a Host or Service's Configuration Settings

The Export function allows you to store the configuration settings of the entire host, or only a subset of these settings (for example, the settings for a specific service). You can export the following sets of values:

- ▶ Updated settings. This means exporting only the settings to which you've assigned new values but have not committed. See "Updating the Host or Service with Changes," on page 574 for details on how to commit changes.
- ▶ All the settings, including those that have not been changed.

### To export configuration settings:

- 1 Click the host or service whose values you want to export and then click **Save/Export configuration settings**. Alternatively, you can right-click the host or service and click Export Settings. The Export Performance Settings dialog box is displayed.



- 2 From the Save as Type box, choose whether to save only the updated (and uncommitted) settings or all the settings for the selected host or service. If you save only the updated settings, ProTune saves them in a file with a *.ups* extension; if you choose to save all the settings, they are saved in a file with an *.aps* extension.

- 3 Enter a meaningful name for your settings file and click **Save**.

ProTune saves your settings in a file with the name you specified.

## Importing Configuration Settings for a Host or Service

The Import function allows you to import previously stored settings and apply them to hosts and services. You can import and apply settings to:

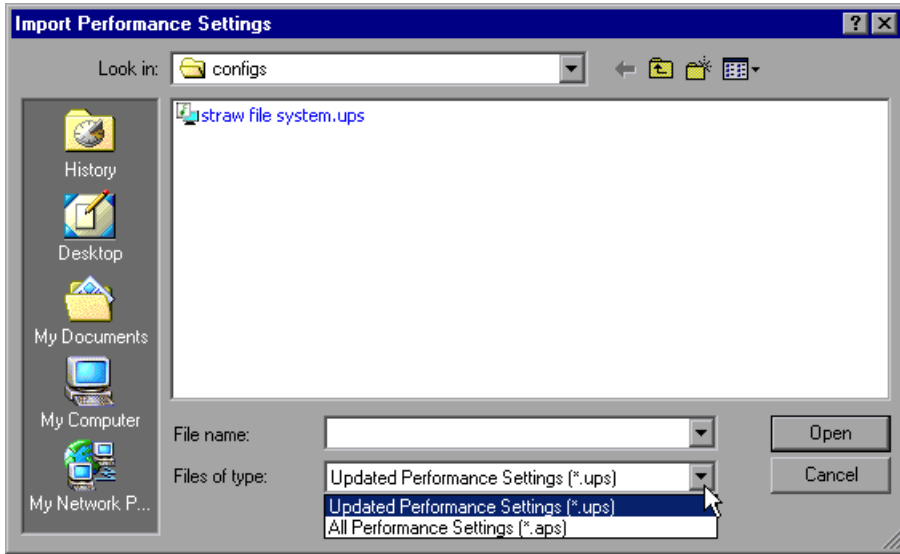
- ▶ Hosts
- ▶ Services
- ▶ Sub-categories of services

Only the entries that are relevant to the host, service or sub-category into which you are importing will be imported from the settings file. For example, you may have a file that contains settings for both Internet Information Server (IIS) and Windows Operating System. If you click the host's Windows Operating System tree element and then import the settings from the file, ProTune loads only the network settings for Windows Operating System, not the IIS settings. Another example: Importing Apache settings into a host that includes IIS does not overwrite the host's IIS settings.

**To import configuration settings:**



- 1 Click the host or service whose values you want to import and then click **Load/Import configuration settings**. Alternatively, you can right-click the host or service and click Import Settings. The Import Performance Settings dialog box is displayed.



- 2 From the Files of type box, choose whether to view the files containing only updated performance settings (files with *.ups* extensions), or files containing all the performance settings for the selected host or service (files with *.aps* extensions).
- 3 Select the settings file from which you want to import, and click **Open**.  
ProTune imports the settings into the specified host, service or category.
- 4 To view the new settings, right-click the host icon (not the service) and click **Refresh**.

## Saving and Loading Profiles

A profile contains all the configuration settings for a group of hosts.

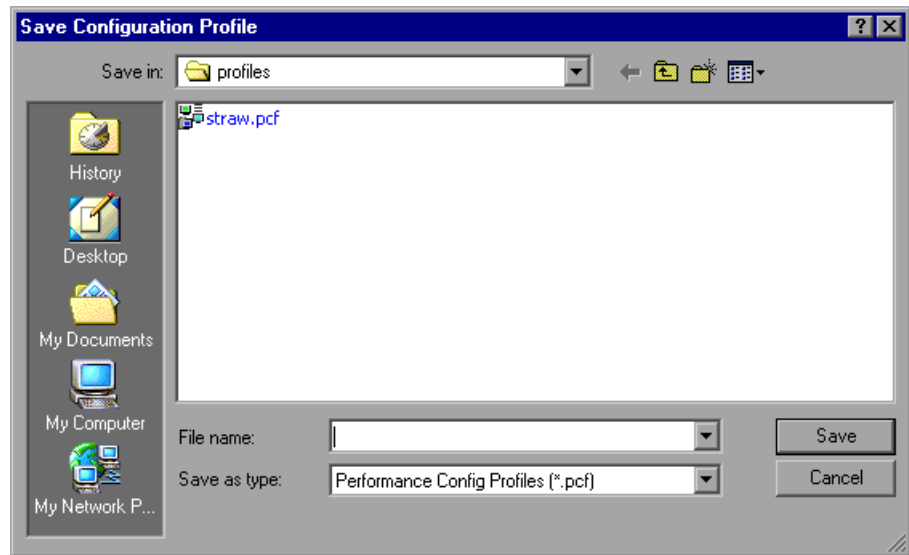
When dealing with large clusters of hosts that may have different operational functions, creating profiles lets you save and load configuration settings for all the hosts in the group. For example, you may find it useful to group a cluster of Web servers into a profile, or group a set of servers relating to an e-commerce or intranet application.

When you save a profile, ProTune saves all the configuration settings for all the hosts that appear in the Tune window.

### To save a profile:



- 1 Click **Save Tuning Profile**, or right-click the Server Configuration tree element and click Save Profile. The Save Configuration Profile dialog box is displayed.

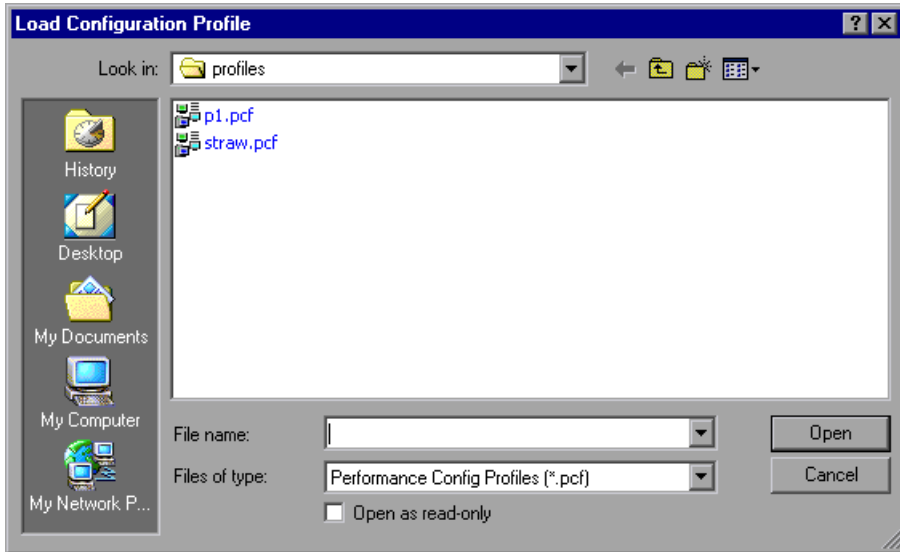


- 2 Enter a name for your profile and click **Save**. ProTune saves the profile with a *.pcf* extension.

**To load a profile:**



- 1 Click **Open Tuning Profile**, or right-click the Server Configuration tree element and click Load Profile. The Load Configuration Profile dialog box is displayed, showing the saved profiles as *.pcf* files.



- 2 Choose the profile you want to load, and click **Open**. ProTune loads the profile and connects to each of the profile's hosts to retrieve and load the profile's settings into the current host configuration. ProTune displays the new configuration in the **Tune** tab.

## Creating a New Profile

---

**Note:** When ProTune creates a new profile, it erases the current profile from memory. This means that all the server settings that you defined will be cleared from the window. If you need to save your settings, make sure you save the existing profile before creating a new one.

---

### To create a new profile:



- 1** Click **New Tuning Profile**. ProTune removes the server icons and their accompanying information from the **Tune** tab.
- 2** Add the required servers and services, and connect to them.
- 3** Save the new profile.





# 34

---

## Configuring Tuning Agents

A tuning agent is an application that runs on the host that you want to tune, and allows you to tune the host remotely. This chapter describes the tuning agent's features and explains how to configure its settings.

The chapter includes the following topics:

- ▶ Changing Tuning Agent Passwords
- ▶ Changing the Tuning Agent's Port
- ▶ Using the Performance Tuner Registry
- ▶ Automatically Starting the Tuning Agent when Booting
- ▶ Starting and Stopping the Apache and IBM HTTP Servers

### About Configuring Tuning Agents

The tuning agent interrogates the host and gathers performance-related information and tuning parameters, and passes the information to the Console. It allows the user to remotely configure and administer the target system.

The tuning agent is a passive service and does not consume any CPU resources when not processing requests. The agent requires only between 10 and 15MB of memory on the host.

The Console uses the tuning agent to change the target system's configurable parameters.

You install the tuning agent over the network from the ProTune Console workstation, or locally on the target system (see “Connecting to the Host Computer,” on page 562).

The tuning agent does not require any registry updates. It uses the PE\_HOME, PE\_USE\_SSL, and PE\_USE\_PORT optional environment variables. Alternatively, you can pass these as arguments to the pe\_agent and pe\_registry commands. For information on these commands, see “Configuring Tuning Agents,” on page 591.

## Changing Tuning Agent Passwords

The tuning agent is supplied with a number of predefined users and passwords (see the table in “Specifying Tuner Agent Settings,” on page 558). It is recommended to change these passwords to prevent unauthorized access to the tuners.

The passwords are defined in the security.properties file, located in the Performance Expert\agent\config directory on the host computer. Note that the passwords must be specified separately for each host.

Following is an example of the relevant section in the security.properties file:

```
# .....
# List of users
# .....
# List of active users that are allowed access to the performance expert agent
# service. By default, admin, mercury, and guest users are supplied. The
# guest user is used if client does not supply any credentials. In order to
# enforce strict user authentication and control, you should remove the guest
# user account from the user list (below).
#
# Format : user.<NAME>=<PASSWORD>
#
# Example: user.admin=changeit
#
# It is strongly recommended that the passwords for admin, mercury and guest
# be changed. The passwords for users will automatically be encrypted when
# tuning agent process (pe_agent) is started.
#
user.admin={DES3S}5197355F27E7E24C28F2D2B9B63C8B2E
user.guest={DES3S}FFD6868073ADF03D
user.mercury={DES3S}9504D70960F68B9C
```

Using a text editor, change the password of the relevant user and save the file. (Note that when the tuning agent is next started, the password will be encrypted.)

## Changing the Tuning Agent's Port

By default, all communication between the Console and the tuning agent is handled by a proprietary messaging protocol encoded and secured over Secure Sockets Layer (SSL). You can change the tuning agent's default port from OTP-SSL 4863 to a user defined port.

When you install the tuning agent, the installation process installs the `pe_agent.bat` (Windows) and `pe_agent` (UNIX) batch files on the host.

---

**Tip:** On a Windows host, the `pe_agent.bat` file is located in the *Program Files\Mercury Interactive\Performance Expert\agent\bin* directory.

---

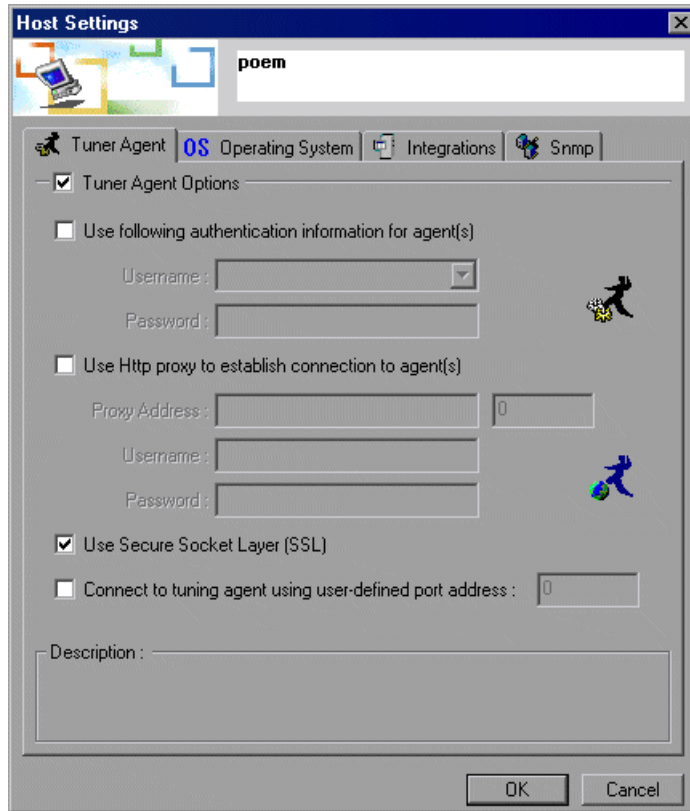
You can modify the agent's listening port using the Host Properties dialog box or the command line.

### Specifying the port using the Host Properties dialog box:



- 1 In the Server Configurations tree, click the host and then click the **View Host Properties** button. Alternatively, right-click the host and choose Properties.

The Host Settings dialog box is displayed:



- 2 Check the **Connect to tuning agent...** box and specify the port in the adjacent field.

**Specifying the port in the command line:**

- Specify the listening port when you invoke the batch file from the command line. Use the following syntax:

```
pe_agent <Port#> <SSL_Flag> <Path_to_PE_Installation>
```

The following examples illustrate how to use the batch file:

| Command                       | Action   |
|-------------------------------|--|
| pe_agent 1234                 | Launches the tuning agent at port 1234 with SSL enabled.   |
| pe_agent 1235 false           | Launches the tuning agent at port 1235 with SSL disabled.  |
| pe_agent 1236 true C:\protune | Launches the tuning agent at port 1236 with SSL enabled, using configuration and tuners from C:\protune. |

- Set the tuning agent's environment variables and then run the batch file. The following example illustrates this method:

```
set PE_USE_PORT=4444
set PE_USE_SSL=true
set PE_HOME=C:\protune
pe_agent
```

---

**Tip:** On a UNIX host, use a similar procedure, depending on your UNIX shell.

---

## Using the Performance Tuner Registry

The Performance Tuner Registry allows you to configure the tuners on a host.

Each environment (for example, IIS, Apache, and Oracle) has a dedicated tuner that is capable of administering the environment. In some cases the tuner needs information on where to find the application that needs tuning, and may also need logon credentials.

The Performance Tuner Registry provides a command-line interface for configuring the individual tuners.

### To invoke the Performance Tuner Registry:

- 1** On the host that is being tuned, set the PE\_HOME environment variable so it points at the Performance Tuner's home directory. **Note:** When you install the tuning agent remotely from the Console machine or from the installation CD, ProTune sets the environment variable to this value.
- 2** From the command line, enter one of the following commands:
  - pe\_registry.bat (Windows)
  - pe\_registry (UNIX)

---

**Note:** Alternatively, if you installed the tuning agent locally from the CD, you can invoke the Performance Tuner Registry by choosing **Start > Programs > Performance Expert > Tuning Agent Configuration**.

---

ProTune invokes the Performance Tuner Registry and displays the Performance Tuner Registry Console:

```

Performance Tuner Registry Console (v. 1.1)
Mercury Interactive Corporation
Agent Directory: C:\PROGRA~1\MERCUR~1\ProTune\agent
Date           : Wed Mar 26 17:19:41 GMT+02:00 2003
-----

List of application performance tuners:

* [1 ] Apache Web Server 1.x/2.x           ( ver. 1.2 )
C [2 ] BEA Weblogic 6.x/7.x               ( ver. 1.3 )
C [3 ] IBM HTTP Server                     ( ver. 1.2 )
C [4 ] IBM Websphere Advanced              ( ver. 1.1 )
C [5 ] IBM Websphere Single Server         ( ver. 1.1 )
C [6 ] iPlanet Enterprise Server           ( ver. 1.1 )
* [7 ] Microsoft IIS/ASP 4/5              ( ver. 1.1 )
* [8 ] Operating System                   ( ver. 1.1 )
  [9 ] Oracle 9iAS                         ( ver. 1.1 )
  [10] Oracle Database                     ( ver. 1.1 )
C [11] PeopleSoft 8.x                     ( ver. 1.1 )
  [12] SAP Enterprise Portals              ( ver. 1.1 )
C [13] Siebel 7.x                         ( ver. 1.1 )
  [14] SQL Server 7.5/2000                ( ver. 1.2 )
-----

[Main Menu] Select an option:
<L>list current tuners
<E>nable a tuner (or E#)
<D>isable a tuner (or D#)
<#> to configure a tuner
<Q>uit
Select [L,E,D,Q,#] ?

```

The window displays a list of all the services for which tuners are available. The services are marked as follows:

| Sign | Indicates   | Comments   |
|------|---|--|
| *    | The tuner is active on the host (the service is installed on the machine).  | If the service is not installed on the host, the tuner is not marked as active.  |
| C    | You may need to configure the tuner before it can become active. For example, you may need to specify where the service is installed. |  |
| X    | The tuner is disabled.  | The user cannot view information about the service for which the tuner is intended, and cannot tune or administer the service. |

The Performance Tuner Registry Console allows you to perform the following actions:

- List the current tuners
- Enable a tuner
- Disable a tuner
- Configure a tuner
- Quit the Performance Tuner Registry Console

**To list the current tuners:**

- Type "L" and press <Enter>.

**To enable a tuner:**

- 1 Type "E" and press <Enter>.

The Performance Tuner Registry Console asks you to enter the ID of the tuner that you want to enable.

- 2 Type the tuner ID (the number in brackets that appears before the tuner's name) and press <Enter>.



---

**Note:** If the tuner requires configuration (indicated by the letter “C” before the tuner’s name), you cannot enable it.

---

The next time you list the current tuners, the enabled tuner appears with an asterisk.

**To disable a tuner:**

- 1 Type “D” and press <Enter>.

The Performance Tuner Registry Console asks you to enter the ID of the tuner that you want to disable.

- 2 Type the tuner ID (the number in brackets that appears before the tuner’s name) and press <Enter>.

The next time you list the current tuners, the enabled tuner appears without an asterisk.

**To configure a tuner:**

- 1 Type the tuner’s number to select it, and press <Enter>.

The configuration menu for the selected tuner is displayed. The following example shows the configuration menu for the WebLogic Application Server tuner:

```
-----
List of current BEA WebLogic Paths:
*** Windows Registry ***
-----
[BEA Weblogic Menu] Select an option:
<L>list BEA Weblogic Home Directories
<A>add a BEA Weblogic Home Directory
<R>remove a BEA Weblogic Home Directory
<Q>quit - to return to main menu
Select [L,A,R,Q] ?
```

- 2 Follow the onscreen instructions for configuring the selected tuner.

**To quit the Performance Tuner Registry Console:**

- Type “Q” and press <Enter>.

## Automatically Starting the Tuning Agent when Booting

You can configure the host to start the tuning agent automatically when the host is started.

**On a Windows system:** Using regedit.exe, create a key in the Windows registry under:

```
[HKLM\Software\Microsoft\Windows\CurrentVersion\Run]
PE_AGENT = "%PE_HOME%\agent\bin\pe_agent.bat" (REG_SZ)
```

**On a UNIX system:** Update your startup file under /etc, and add a command to launch pe\_agent.

When you launch the tuning agent (pe\_agent.bat or pe\_agent), you can include the following optional arguments:

| Argument      | Specifies  | Values        |
|---------------|--|---------------|
| [PE_USE_PORT] | listening port   | Default: 0    |
| [PE_USE_SSL]  | whether SSL is enabled   | True<br>False |
| [PE_HOME]     | location of PE_HOME (if the PE_HOME environment variable is not defined) |               |

Following is the syntax for launching the tuning agent:

```
pe_agent[.bat] [PE_USE_PORT] [PE_USE_SSL] [PE_HOME]
```

Following is an example of how to run the pe\_agent.bat command:

```
pe_agent.bat 0 true C:\Program Files\Mercury Interactive\PerfExpert
```

## Starting and Stopping the Apache and IBM HTTP Servers

To enable the Console to remotely start and stop an Apache Web Server or IBM HTTP Server running on a host machine, you need to perform the following preparatory actions:

- ▶ Create scripts for starting and stopping the Web server, and place them in a directory on the host machine.
- ▶ On the host machine, configure the tuner. This includes specifying the path to the Web server and the scripts for starting and stopping it.

### Creating Scripts for Starting and Stopping the Web Server

To enable the tuning agent to remotely start and stop the Apache Web server or IBM HTTP Server, you need to create scripts and place them in the relevant directory of the Web server installation.

- 1 Use a text editor to create the following batch files:

**On a Windows system:**

| Web Server      | Filename  | Content                            |
|-----------------|-----------|------------------------------------|
| Apache          | start.bat | Apache.exe -w -n "Apache" -k start |
| Apache          | stop.bat  | Apache.exe -w -n "Apache" -k stop  |
| IBM HTTP Server | start.bat | net.exe start "IBM HTTP Server"    |
| IBM HTTP Server | stop.bat  | net.exe stop "IBM HTTP Server"     |

**On a UNIX system:**

| Web Server | Filename | Content                           |
|------------|----------|-----------------------------------|
| Apache     | start    | apachectl -w -n "Apache" -k start |
| Apache     | stop     | apachectl -w -n "Apache" -k stop  |

To enable starting and stopping an IBM HTTP Server running on a UNIX host machine, define shell scripts (start and stop) for starting and stopping the IBM HTTPD service.

**2** Place the batch files in the relevant directory:

- **For Apache Web Server:** In the \bin directory of the Apache installation. For example, on a Windows computer, this might be located at C:\Program Files\Apache Group\Apache\bin. On a Linux system, you need to create the /bin directory (with the start and stop scripts) under /etc/httpd.
- **For IBM HTTP Server:** In the directory containing the Apache.exe file. This is typically the **IBM HTTP Server** directory.

### **Configuring the Tuner**

You use the pe\_registry batch file to configure the tuner, specifying the path to the Web server and the start and stop scripts on the host machine.

**To specify the path to the Web server:**

- 1** On the host machine, invoke the Performance Tuner Registry Console (see "Using the Performance Tuner Registry," on page 596).

The list of tuners is displayed.

```

Performance Tuner Registry Console (v. 1.1)
Mercury Interactive Corporation
Agent Directory: C:\PROGRAM~1\MERCUR~1\ProTune\agent
Date           : Wed Mar 26 17:19:41 GMT+02:00 2003
-----

List of application performance tuners:

* [1 ] Apache Web Server 1.x/2.x           ( ver. 1.2 )
C [2 ] BEA Weblogic 6.x/7.x               ( ver. 1.3 )
C [3 ] IBM HTTP Server                     ( ver. 1.2 )
C [4 ] IBM Websphere Advanced              ( ver. 1.1 )
C [5 ] IBM Websphere Single Server         ( ver. 1.1 )
C [6 ] iPlanet Enterprise Server           ( ver. 1.1 )
* [7 ] Microsoft IIS/ASP 4/5              ( ver. 1.1 )
* [8 ] Operating System                    ( ver. 1.1 )
  [9 ] Oracle 9iAS                          ( ver. 1.1 )
 [10] Oracle Database                      ( ver. 1.1 )
C [11] PeopleSoft 8.x                     ( ver. 1.1 )
 [12] SAP Enterprise Portals               ( ver. 1.1 )
C [13] Siebel 7.x                          ( ver. 1.1 )
 [14] SQL Server 7.5/2000                 ( ver. 1.2 )
-----

[Main Menu] Select an option:
<L>list current tuners
<E>nable a tuner (or E#)
<D>isable a tuner (or D#)
<#> to configure a tuner

```

- 2 Type the number of the Web server tuner and press <Enter>. (In the example above, you would type the number 1 for the Apache Web Server, or 3 for the IBM HTTP Server.) This displays the tuner's configuration menu. The current paths to the Web server are listed above the menu.

```

-----
List of current Apache Paths:
(1) C:\Program Files\Apache Group\Apache
-----

[Apache Menu] Select an option:
<L>list Apache Paths
<A>dd an Apache Path
<R>emove an Apache Path
<Q>uit - to return to main menu
Select [L,A,R,Q] ?

```

The path must specify the directory that contains the Apache.exe file, and must **not** specify a subdirectory of that directory. For example, for an Apache Web server installation on a Windows computer, C:\Program Files\Apache Group\Apache is acceptable; C:\Program Files\Apache Group\Apache\bin is not.

On a Linux computer (for the default Apache installation), the path should be `/etc/httpd`.

- 3 If the correct path is not listed, you need to add it to the path list. To add the path to the list, type **A** and press `<Enter>` to choose the **Add an Apache Path** or **Add an IBM HTTP Server Path** option, specify the path to the directory that contains the `Apache.exe` file, and press `<Enter>`. The path to the Web server is added and displayed by the Performance Tuner Registry.

**To specify the starting and stopping scripts:**

- 1 From the Web server configuration menu, type **U** to update the configuration and press `<Enter>`. The Performance Tuner Registry asks for the ID of the Web server that you are configuring. (This is the number to the left of the relevant path in the list of current Apache paths.)

```
-----
[Apache Menu] Select an option:
  <L>ist Apache Configuration Entries
  <A>dd an Apache Configuration
  <U>pdate Apache Configuration
  <R>emove an Apache Configuration
  <Q>uit - to return to main menu
Select [L,A,U,R,Q] ? u
Please specify the apache configuration id to update [1-1]: 1
```

- 2 Type the relevant ID and press `<Enter>`.

The application displays the path to the Web server's home directory. This is the path that you specified in Step 2 on page 603.

```
Please specify the apache configuration id to update [1-1]: 1
Apache home directory [C:\Program Files\Apache Group\Apache\] :>
```

If the displayed path is incorrect, enter the correct one.

- 3 Press `<Enter>`. The application displays the Web server's version number.

If the version is incorrect, enter the correct one.

- 4 Press `<Enter>`. The application displays the Web server's startup script (if one has been specified).

- 5** Ensure that the displayed script has the name that you specified above in “Creating Scripts for Starting and Stopping the Web Server,” on page 601.

If the name that you defined is not displayed, enter the correct name (start.bat for a Windows system, or start for UNIX).

- 6** Press <Enter>. The application displays the Web server’s shutdown script (if one has been specified).

Ensure that the displayed script has the name that you specified above in “Creating Scripts for Starting and Stopping the Web Server,” on page 601.

- 7** If the name that you defined is not displayed, enter the correct name (stop.bat for a Windows system, or stop for UNIX).

- 8** Press <Enter>.

The Web server’s configuration menu is displayed.

- 9** Type **Q** and press <Enter> to exit the configuration menu. The Performance Tuner Registry main menu is displayed.

- 10** Type **Q** and press <Enter> to exit and save your settings.





# 35

---

## Tune Tab Functions

This chapter describes various functions that are available to you via the **Tune** tab.

The chapter describes the following functions:

- ▶ Start a Service
- ▶ Stop a Service
- ▶ Reboot Host Machines
- ▶ Reconnect the Console to a Server
- ▶ Stop the Tuning Agent
- ▶ Print Host Configurations
- ▶ Reload Host Configuration
- ▶ Remove Host from Server Configurations Tree

### About Tune Tab Functions

Some tuning functions can be performed by all categories of users—whether they have only read-only permissions, or update permissions, or full administrative privileges. Some functions are available only to users who have administrator privileges (that is, users who connect with the *admin* username). If you attempt to perform a function for which you are not authorized, ProTune displays an error message.

For details of the various types of users and their usernames and passwords, see the table on page 560.

## Start a Service

(For users with administrator access only).

### To start a Windows service:

- 1 Expand the host computer's Windows Services element. The host's services are displayed.
- 2 Right-click the service that you want to start, and choose Admin > Start Service.
- 3 To verify that the service has started, right-click the service and choose Refresh. If the service has started, its status (in the right side of the Information tab) changes to **Running**.

### To start other types of services:

- 1 Right-click the service, and choose Admin > Start Service.  
ProTune displays a dialog box requesting confirmation.
- 2 Click **Yes** to start the service.

---

**Note:** After you choose Start Service, you may have to wait for a while until the service has started.

---

## Stop a Service

(For users with administrator access only).

### To stop a Windows service:

- 1 Expand the host computer's Windows Services element. The host's services are displayed.
- 2 Right-click the service that you want to stop, and choose Admin > Stop Service.

- 3** To verify that the service has stopped, right-click the service and choose Refresh. If the service has stopped, its status (in the right side of the Information tab) changes to **Stopped**.

**To stop other types of services:**

- 1** Right-click the service, choose Admin, and then choose Stop Service. ProTune displays a dialog box requesting confirmation.
- 2** Click **Yes** to stop the service.

---

**Note:** After you choose Stop Service, you may have to wait for a while until the service has stopped.

---

## Reboot Host Machines

(For users with administrator access only).

ProTune allows you to reboot host machines from the Console machine.

If the host machine is running Windows, you need to specify its username and password to enable the host to log in automatically after rebooting. You specify the username and password via the Performance Tuner Registry.

**To specify the host machine's username and password:**

- 1** Invoke the Performance Tuner Registry. (See "Using the Performance Tuner Registry," on page 596 for details.)
- 2** Type **8** (for the Operating System tuner) and press <Enter>. The OS Menu is displayed.
- 3** Type **C** and press <Enter>. The Performance Tuner Registry prompts you to enter the username and password.
- 4** Enter the username and password and exit the Performance Tuner Registry.

**To reboot a host machine:**

- 1 Right-click the server's icon, choose Admin, and then choose Reboot Host. ProTune displays a dialog box requesting confirmation.
- 2 Click **Reboot** to reboot the host machine.

**To reboot all the host machines:**

- 1 Right-click the Server Configurations node, choose Admin, and then choose Reboot All. ProTune displays a dialog box requesting confirmation.
- 2 Click **Reboot All** to reboot all the host machines.

## Reconnect the Console to a Server

Once the tuning agent has been installed on a server, you can access the server by clicking the **Connect to Host button** on the toolbar. In the Connect to Server dialog box, choose the server you want to tune and click **Connect**. If the tuning agent is running on the server, the Console connects to the server via the tuning agent, and shows you the server information.



If the tuning agent has been installed on the server but is not currently running, start it by clicking the **Start Tuning Agent** button. Clear the Auto-Install... box, verify that the other fields have the correct values, and click **Start**.

On the Console machine, the server icon in the Server Configurations tree changes to blue, indicating that the connection to the server is alive.

## Stop the Tuning Agent



To stop the tuning agent that is running on the server, click the server's icon and click the **Stop Tuning Agent** button. When ProTune asks you to confirm the action, click **Yes**.

## Print Host Configurations

You may find it useful to keep a hard-copy record of a host's tuning settings. ProTune allows you to print the configuration settings as they appear in the **Information** tab.

To print a host's configuration settings:

- 1 Expand the icons in the **Information** tab so that the data that you want to print is displayed.
- 2 Right-click the server's icon in the Servers Configurations tree, and choose Print.

## Reload Host Configuration



To reload the current settings from a host machine, click the host's icon in the Servers Configurations tree and click the **Reload Host Configuration** button. Alternatively, right-click the host's icon and choose Refresh.

## Remove Host from Server Configurations Tree

To remove a host from the Server Configurations tree, right-click the host's icon and choose Remove.



# 36

---

## Tuning UNIX Hosts

This chapter describes the permissions, access rights and actions that you need to perform before you can tune a host machine running a UNIX operating system.

It includes the following sections:

- ▶ Using Telnet
- ▶ Redirecting Script Output
- ▶ Solaris Requirements
- ▶ IBM AIX Requirements
- ▶ HP-UX Requirements
- ▶ Linux Requirements

### Using Telnet

It is recommended to use Telnet to start or stop a tuning agent on a UNIX host.



To launch Telnet, click **Launch Terminal Client**.

### Redirecting Script Output

It is recommended to redirect the output of the start and stop scripts to a user-defined log file, especially if the scripts print information to the standard output device. This is because some versions of JRE may prevent proper execution of scripts that produce a lot of output.

To redirect the script output, use one of the following methods:

- Specify the script path with redirection to the log file, as in the following example:
 

```
mystartup.sh > startup.log
```
- Update the script code to automatically generate the log file.

## Solaris Requirements

### Access Rights and Permissions

- 1 The user running the tuning agent must have access rights and permissions to execute the following commands:

| Command                  | Gives this information             | Default location |
|--------------------------|------------------------------------|------------------|
| psrinfo                  | CPU speed and number of CPUs       | /usr/sbin        |
| prtconf                  | Total RAM                          | /usr/sbin        |
| vmstat                   | Available RAM                      | /usr/sbin        |
| swap                     | Total and available virtual memory | /usr/sbin        |
| ndd --<br>[program]      | Network tuning parameters          | /usr/sbin        |
| /etc/system --<br>[file] | File system tuning parameters      |                  |

### PATH Environment

Update the PATH environment so it includes the /usr/sbin directory.



## Verification

After updating the path, execute each of the commands listed in the table above to verify that you have the appropriate access rights for executing them. **Note:** Some commands may require root privileges. Also verify the existence of the system file in the /etc directory.

## IBM AIX Requirements

### Access Rights and Permissions

- 1 The user running the tuning agent must have access rights and permissions to execute the following commands:

| Command                                     | Gets this information                    | Default location |
|---|--|------------------|
| uname                                       | Name and version of the operating system | /usr/bin         |
| bootinfo                                    | Total RAM                                | /usr/bin         |
| vmstat                                      | Total and available virtual memory       | /usr/bin         |
| lsdev                                       | Number of processors                     | /usr/bin         |
| /usr/samples/<br>kernel/vmtune -- [program] | File system tuning parameters            |                  |
| no --<br>[program]                          | Network tuning parameters                |                  |

- 2 If the user's Java installation does not support SSL with RSA encryption, the user can launch the tuning agent without using SSL. To do this, launch the tuning agent with the SSL flag set to False, as in the following example:

```
pe_agent 4862 false
```

---

**Note:** The default port for non-SSL connections is 4862.

---

When connecting to a tuning agent whose port and SSL state have non-default settings, the client user should update the host properties for the target host appropriately.

### **PATH Environment**

Update the PATH environment so it includes the `/usr/bin` directory.

### **Verification**

After updating the path, execute each of the commands listed in the table above to verify that you have the appropriate access rights for executing them. **Note:** Some commands may require root privileges.

## **HP-UX Requirements**

### **Access Rights and Permissions**

The user running the tuning agent must have access rights and permissions to execute the following commands:

| <b>Command</b>      | <b>Gets this information</b>             | <b>Default location</b> |
|---------------------|--|-------------------------|
| uname               | Name and version of the operating system | /usr/bin                |
| model               | CPU speed                                | /usr/sbin               |
| swapinfo            | Total and available virtual memory       | /usr/sbin               |
| dmesg               | Total and available RAM                  | /etc                    |
| ioscan              | Number of processors                     | /usr/sbin               |
| ndd --<br>[program] | Network tuning parameters                |                         |

## PATH Environment

Update the PATH environment so it includes the following directories:

- /usr/sbin
- /usr/bin
- /etc

## Verification

After updating the path, execute each of the commands listed in the table above to verify that you have the appropriate access rights for executing them. **Note:** Some commands may require root privileges.

## Linux Requirements

### Access Rights and Permissions

The user running the tuning agent must have access rights and permissions to execute the following commands:

| Command             | Gets this information  | Default location |
|---------------------|--|------------------|
| uname               | Name of the operating system                                 | /usr/bin         |
| /proc/cpuinfo file  | Operating system version, CPU speed and total number of CPUs |                  |
| /proc/meminfo file  | All memory statistics  |                  |
| sysctl -- [program] | File system and network tuning parameters                    |                  |

## **PATH Environment**

Update the PATH environment so it includes the `/usr/bin` directory.

## **Verification**

After updating the path, execute each of the commands listed in the table above to verify that you have the appropriate access rights for executing them. **Note:** Some commands may require root privileges. Also check for the existence of the `meminfo` and `cpuinfo` files in the `/proc` directory.

# **Part VII**

---

## **Appendixes**



# A

---

## Troubleshooting the Console

ProTune enables you to test entire applications. If one of the components of the application is not configured properly, ProTune sessions will not run.

This appendix discusses the most common ProTune problems:

- ▶ ProTune Communications
- ▶ Failure to Communicate with a Load Generator
- ▶ Failure to Connect to the AUT Database
- ▶ Failure to Access Files
- ▶ Failed Vusers or Transactions
- ▶ Increasing the Number of Vusers on a Windows Machine
- ▶ Troubleshooting Firewalls
- ▶ Troubleshooting Remote Tuning

### About Troubleshooting

ProTune relies heavily upon communication between machines on a network. If communication is not established properly, the Console will be unable to send commands to remote load generators and the session will fail. By understanding the reason for the failure and determining when the failure occurred, you can solve most of the communication-related problems.

In order to ensure that the problem lies with your session and not your script, you should verify that your script runs properly on all remote load generators as a stand-alone:

- ▶ Test your GUI scripts on Windows platforms using WinRunner.
- ▶ Test your scripts on UNIX platforms by running them from the command line.
- ▶ Test all other types of scripts on Windows platforms by running them from VuGen, or by running a single user from the Console.

---

**Note:** When a test runs in VuGen, the full browser is used. This differs from a test run in the Console, where only the browser basics are used. There may be occasions when a test passes its run in VuGen, but fails when it is run in the Console. Before running a session in the Console with multiple Vusers, run a single Vuser to ensure the test is bug free.

---

For more information on running scripts in stand-alone mode, refer to the appropriate guide for creating scripts.


## ProTune Communications

Most communication problems can be solved if you understand your ProTune configuration. This knowledge helps you to determine the source of the problem and perform the necessary actions to correct it.

The following diagram illustrates a sample network running ProTune. There are five servers: The ProTune Console, the Web server, the application server, the database server, and the file server which stores the session results (note that result files can also be saved on a non-dedicated server). There are five remote load generators, each one running multiple Vusers.



The arrows indicate the type of communication necessary between the elements of the network. The Users communicate with the Console in both directions (send/receive), but with the file server in one direction (send). The Console must have access to the file server. All Users participating in the session must be able to communicate with the Web server in both directions (send/receive). In order for a client machine to connect to the server machine, it must be able to resolve the server machine name.



If any of the connections are broken, the session will fail.


## Failure to Communicate with a Load Generator

The most common communication error is the failure of the Console machine to connect with a remote load generator. Check the following items:

- TCP/IP setup
- TCP/IP connectivity
- Load generator connections
- UNIX shell

## Checking TCP/IP Setup

The first step in checking your configuration is to verify your machine's TCP/IP setup. ProTune includes a utility called Hostinfo (hostinfo.exe), located under ProTune's bin directory. This utility provides information about the current machine—local name and local address. It also insures that TCP/IP is properly installed on the current machine.



When you invoke Hostinfo, it automatically verifies the TCP stack by:

- retrieving and resolving the local machine name
- retrieving and resolving the IP address

To resolve the IP address, Hostinfo tries to communicate using two UDP sockets on the same machine. It verifies that the IP address obtained while resolving the machine name is the same as the actual IP address of this machine.

To display the results of a test in the Details box, highlight the test name.

Note that the Edit menu in Hostinfo allows you to copy all machine information to the clipboard for sending to support personnel.

## Checking TCP/IP Connectivity

Make sure that TCP/IP connectivity is functional on the Console and Vuser machines. Use a ping utility or type `ping <server_name>` from the DOS command line to verify communication with a remote machine. Make sure that the remote load generator and Console machines can ping each other by IP addresses and machine names.

If the ping does not respond, or fails with a timeout, then the machine name is not recognized. To solve this problem, edit the hosts file, located in the `WINNT\system32\drivers\etc` directory, and add a line with both the IP address, and the name. For example:

```
#      102.54.94.97    rhino.acme.com      # source server
#      38.25.63.10    x.acme.com          # x client host
```

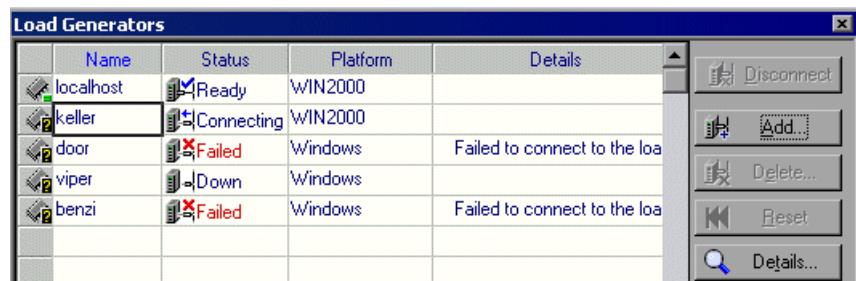


## Load Generator Connections

To verify the load generator connectivity, connect to each one of the remote load generators from the Console's Load Generators dialog box. In the load generator's Platform field, select a Windows or UNIX platform. Select the load generator(s) and click the Connect button. The status changes to *Connecting*.

If the Connection fails, the status changes to *Failed* and details are written to the Details box. Double-click the details box for more information about a failure.

If a connection succeeds, the status changes to *Ready*, and the actual platform name appears in the Platform box (such as WINNT, UNIX, etc.)



If your session uses several domains (for example, Vusers on a different domain than the Console), the Console may have trouble communicating with the load generators. This occurs because the Console uses the short load generator name—not including the domain—by default. To solve this, you must tell the Console to determine the full load generator names, including the domains.

Modify the *miccomm.ini* file in the Console machine's Windows directory as follows:

```
[tcpnet]
LocalHostNameType= 1
```

The possible values for LocalHostNameType are:

- 0 - Attempt to use the full machine name.
- 1 - Use the short machine name. This is the default.

---

**Note:** In certain environments such as WINS, load generators are unable to resolve machine names.

---

### **Connecting to a Console with Multiple IP Addresses**

If the load generator machine does not recognize the Console machine by its short name or full name, and the Console machine has more than one IP address, you can define an alias name for the Console machine in the load generator's *hosts* file, located in the WINNT\system32\drivers\etc directory. The alias name should point to the IP address you want the load generator to recognize. For example: 255.0.0.1 delta.

### **UNIX Shell**

For UNIX Vusers, make sure that the Windows Console can execute a remote shell command. Type the following at the DOS command prompt: `rsh -l <UNIX user login name> <load generator name> <command>`. If you get a message indicating that permission is denied, make sure the *.rhosts* file in your UNIX home directory contains Console machine permission for the user login name. In some cases, a "+" character must be added at the end of the *.rhosts* file. For example, if you log on to the Console as *bill* and connect

to the UNIX load generator as *mike*, you must ensure that *mike* allows *bill* to log on using his name. This can be done by adding the line "+ bill" at the beginning of mike's *.rhosts* file.

For more information on setting user login names, see "Configuring Load Generator Settings" on page 61.

#### To use UNIX without RSH:

- 1 On the UNIX Load Generator machine, run the agent daemon by running the following command from *<ProTune directory>/bin*:

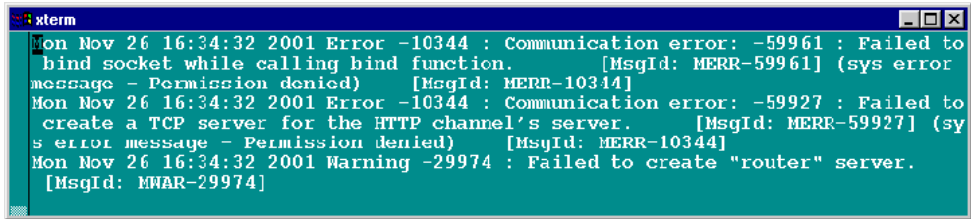
```
m_daemon_setup -install
```

This runs a daemon called *m\_agent\_daemon*, and if successful you will receive a message: *m\_agent\_daemon installed successfully*.

The agent will now keep running, even if the user is logged off. It will only stop running using the command explained in step 3, or by rebooting the machine.

- If you receive the message *ERROR: File m\_agent\_daemon doesn't exist*, this means that you are not in the same directory as the file (meaning not in *<ProTune\_root>/bin* directory, or the file really doesn't exist, which indicates a problem with the installation).
- If a daemon of this name is already being run by the same user you will receive the following warning:  
WARNING: Could not install *m\_agent\_daemon*, reason - user *<user\_name>* is already running *m\_agent\_daemon* on this machine.
- If an error occurred, you will receive the following error message:  
ERROR: Could not install *m\_agent\_daemon*. Check log file *m\_agent\_daemon[xxx].log* in your temp directory.

- If you look at the log file `m_agent_daemon[xxx].log` in the `temp` directory, you will see the following errors, even if the installation succeeded:



```

xterm
Mon Nov 26 16:34:32 2001 Error -10344 : Communication error: -59961 : Failed to
bind socket while calling bind function. [MsgId: MERR-59961] (sys error
message - Permission denied) [MsgId: MERR-10344]
Mon Nov 26 16:34:32 2001 Error -10344 : Communication error: -59927 : Failed to
create a TCP server for the HTTP channel's server. [MsgId: MERR-59927] (sy
s error message - Permission denied) [MsgId: MERR-10344]
Mon Nov 26 16:34:32 2001 Warning -29974 : Failed to create "router" server.
[MsgId: MWAR-29974]

```

These messages appear because the ProTune agent always tries to open port number 443 (because any agent can be a MI Listener, and the MI Listener always listens to this port), and in UNIX machines, this port cannot be opened by any user except for the root user. However, this will not interfere with using this agent for the Load Generator machine.

- 2 In the Console, in the Generators > Load Generator Information > Unix Environment tab, check the **Don't use RSH** option. Then connect as usual.
- 3 To stop the agent daemon, run the following command the `<ProTune_root>/bin` directory: `m_daemon_setup -remove`

This stops the `m_agent_daemon`, and if successful you will receive a message: `m_agent_daemon removed successfully`.

- If no daemon of this name is being run by this user, you will receive the following warning:  
WARNING: Could not remove `m_agent_daemon`, reason - user `<user_name>` is not running `m_agent_daemon` on this machine.
- If an error occurred, you will receive the following error message:  
ERROR: Could not remove `m_agent_daemon`. Check log file `m_agent_daemon[xxx].log` in your `temp` directory.

## Failure to Connect to the AUT Database

If you are running a database application, you must ensure that all remote clients can connect with the database server. If network or configuration errors occur when the client accesses the server, you must correct them before running a session. To ensure that your client application can connect with the database server, perform the following tests.

- Ping
- SQL utilities

**Ping:** Ensure that the client can communicate with the database server using TCP/IP. Use a ping utility or type `ping <server_name>` from the DOS command line.

**SQL Utilities:** Use a simple utility such as ISQL or SQLPLUS to log on to the database server and perform several basic operations.

## Failure to Access Files

A ProTune session will fail if the result path or script is inaccessible to one or more of the participating machines. Check the following items:

- Path Translation
- Script
- Result Path

**Path Translation:** A script's location (path) is always based on the Console machine's mapping of that location. If a Vuser load generator maps to the script's path using a different name, path translation is required. Path translation translates the Console's mapping of a given location to the Vuser load generator's mapping. For example, if one machine maps the script directory as `g:\test`, while another maps it as `h:\test`, the paths should be translated.

Path translation is also effective across platforms—between Windows and UNIX. You use path translation to translate the Windows Console paths into paths recognized by UNIX.

---

**Note:** Path translation is only required if you chose to save all scripts and results to a shared network drive. In the default setup, ProTune saves files locally and collates them to the Console machine; no path translation is required.

---

Suppose that your script is in the `/usr/jon/lr_test1` directory and runs on the UNIX machine, *sunny*. To translate it from the Windows Console machine, *pc1*, where your UNIX directory is mapped as *r*, enter the following line in the path translation table:

|     |     |          |       |
|-----|-----|----------|-------|
| pc1 | r:\ | /usr/jon | sunny |
|-----|-----|----------|-------|

To translate the `f:\qa` Console directory to all load generator machines running `/m/qa/lr_test2/lr_test2.usr` on a UNIX platform, type:

|     |       |       |      |
|-----|-------|-------|------|
| win | f:\qa | /m/qa | UNIX |
|-----|-------|-------|------|

If the paths are not translated properly, the session will fail. For more information about path translation, see Appendix C, "Performing Path Translation."

**Script:** Make sure that the script is accessible to all load generators participating in the session through path translation and permissions. View or run the script as a stand-alone on each of the participating load generators.

**Result Path:** Make sure that the result path is accessible to all load generators participating in the session through path translation and permissions. Check the permissions of the result directory files and modify them if necessary.



## Failed Vusers or Transactions

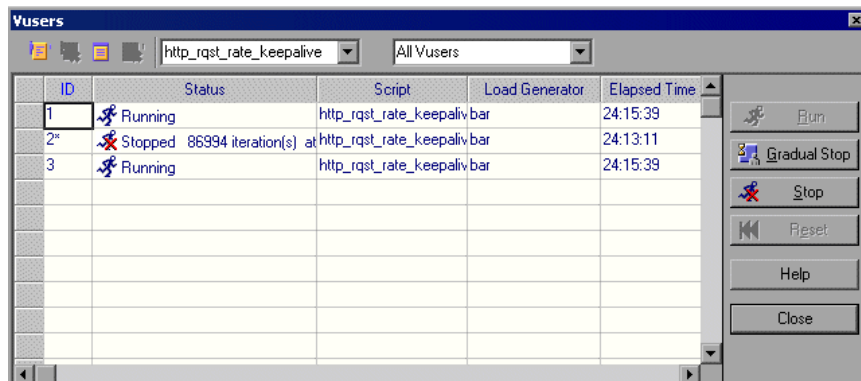
ProTune Vusers or transactions may fail for a variety of reasons relating to the network, database, or actual script. You can find information about session runs from the following sources:

- Run View
- Output Window
- Output File (excluding GUI Vusers)
- Analysis Reports and Graphs

### Run View

The Run view is part of the ProTune Console. The Session Groups window in the top left-hand corner of the view indicates the status of the Vuser groups during and after the session run. During the session run, the columns will show a PENDING, INITIALIZING, READY, RUNNING, or RENDEZVOUS status. You can also view the status of each individual Vuser in the Vusers dialog box. If a Vuser fails and does not complete the script execution, ProTune displays an error status. If a Vuser completes the script execution, ProTune indicates the transaction status of a completed script run using the DONE.FAILED or DONE.PASSED status.

For more information about the Vuser states, see Chapter 9, “Running a Session.”



## Output Window

View the Output window from the Console. The output window contains useful information for debugging a session. The output window lists five types of messages: errors, warnings, notifications, debug, and batch. An error message usually results in a failed script. A warning message indicates that the Vuser encountered a problem, but test execution continued. A notification provides useful information such as recorded think time values and other run-time information. A debug message is sent if you enable the debugging feature in **Tools > Options > Debug Information** (Expert Mode). Batch messages are sent instead of message boxes appearing in the Console, if you are using automation.

| T...      | Message Code (5) | Sample Message Text                            | Total M... | Vusers | Scripts | Generat... |
|-----------|------------------|--|------------|--------|---------|------------|
| [Error]   | -27995           | Action1.c[19]: Error: Requested link ("Te...   | 8          | 8      | 1       | 1          |
| [Warning] | -27798           | -27798 : Action1.c[6]: Warning: could no...    | 70         | 10     | 1       | 1          |
| [Error]   | -27798           | Action1.c[6]: Error: could not resolve ad...   | 10         | 10     | 1       | 1          |
| [Error]   | -19890           | Action1.c[6]: Error -19890 : C-interpreter ... | 10         | 10     | 1       | 1          |
| [Error]   | 0                | Error from ftp_logon_ex at Actions.c (4) : ... | 8921       | 20     | 2       | 1          |

Summary

For more information about the Output window, see Chapter 10, “Viewing Vusers During Execution.”

## Output File

You can view information about script execution in an output file located in the Vuser result directory. The output file, *output.txt*, contains:

- a list of the primary functions called during the session
- error messages from the database server
- transactions and rendezvous information

The extent of the information sent to the output file depends on the output file settings. In the VuGen's run-time settings, you specify a Brief or Extended log. For the Extended log, you can specify a full trace, returned data, or current parameter value. An extended log is helpful for debugging a script, but if you are not debugging, Extended log is not recommended as it introduces extra overhead. For more information about configuring run-time settings, refer to the *ProTune Creating Virtual User Scripts* guide.

## Analysis Reports and Graphs

You can generate graphs and reports to view information about the session run. For example, the Session Summary report displays a table containing the session's run-time data and provides links to the following graphs: Running Vusers, Throughput (Web), Hits Per Second (Web), HTTP Responses per Second, Transaction Summary, and Average Transaction Response Time.

The screenshot displays the 'Analysis Summary' report window. At the top, it shows the scenario name and session details. Below that, a 'Statistics Summary' section lists key performance indicators. The 'Transaction Summary' section includes a table with columns for Transaction Name, Minimum, Average, Maximum, 90 Percent, Pass, Fail, and Abort.

**Analysis Summary** Period: 07/03/2001 10:06:38 - 07/03/2001 10:06:38

Scenario Name: C:\Sanity\_Analysis\Scenario\Scenario1.js  
 Results in session: F:\results\session\W\yes\_amazon\yes\_amazon.lrr.  
 Duration: 20 minutes and 49 seconds.

**Statistics Summary**

- Maximum Running Vusers: 3
- Total Throughput (bytes): 50,222,495
- Throughput (bytes/second): Average: 40,210
- Total Hits: 12,217
- Hits per Second: Average: 10

**Transaction Summary**

Transactions: Total passed: 105 failed: 70 aborted: 3 [Response Time Avg.](#)

| Transaction Name  | Minimum | Average | Maximum | 90 Percent | Pass | Fail | Abort |
|-------------------|---------|---------|---------|------------|------|------|-------|
| tr_amazon_account | 2.145   | 2.56    | 3.245   | 2.68       | 12   | 0    | 0     |
| tr_amazon_book    | 7.871   | 8.825   | 10.425  | 10.31      | 10   | 2    | 0     |
| tr_amazon_books   | 4.346   | 4.588   | 5.177   | 4.68       | 11   | 1    | 0     |
| tr_amazon_list    | 2.454   | 2.615   | 3.575   | 2.66       | 12   | 0    | 0     |

For more information on the available graphs and reports, refer to the *ProTune Analysis User's Guide*.

## Increasing the Number of Vusers on a Windows Machine

Under the normal settings of a Windows machine, you are limited to several hundred Vusers. This limitation is related to the operating system and not to the CPU or memory.

To work around the limitation of the Windows operating system, modify the Windows Kernel as follows:

- 1** Save a copy of the registry file in case you have trouble with these modifications.
- 2** Run Regedit.
- 3** Go to following key in KEY\_LOCAL\_MACHINE:  
System\CurrentControlSet\Control\Session Manager\SubSystems
- 4** Select the Windows key. The default Windows key for NT 4.0 looks like this:  
%SystemRoot%\system32\csrss.exe ObjectDirectory=\Windows  
SharedSection=1024,3072  
Windows=On SubSystemType=Windows ServerDll=basesrv,1  
ServerDll=winsrv:UserServerDllInitialization,3  
ServerDll=winsrv:ConServerDllInitialization,2  
ProfileControl=Off MaxRequestThreads=16

The SharedSection=1024,3072 key has the format xxxx,yyyy where:

xxxx defines the maximum size of the system-wide heap (in kilobytes)

yyyy defines the size of the per desktop heap.

- 5** Increase the SharedSection parameter by changing the yyyy settings from 3072 to 8192 ( which is 8 MB).

This setup successfully allowed 1250 Oracle Vusers to run on a Windows machine using 2 Pentium PRO 200 MHz with 1 GB RAM.

Each Vuser in this setup used approximately 2MB memory. Other Vusers may require more memory.

ProTune was able to load over 2500 Vusers when the Windows terminal server was run as the Operating System and the above registry setting was changed.

The above registry changes enable you to run more threads, allowing you to run more Vusers on the machine. This implies that you are not bound by the Windows operating system; you are only bound by hardware and internal scalability limitations.

## Troubleshooting Firewalls

There are three log files which provide additional information about activity over the firewall.

The **ProTune agent log file** contains information about communication activity between the ProTune agent and the MI Listener.

- ▶ To open the file on Windows machines, right-click the ProTune agent icon in the system tray of the ProTune agent machine, and select **View Log**. Alternatively, open the latest `<temp_directory>\ProTune_agent_startup<unique identifier>.log` file (if the ProTune agent is a process), or `<temp_directory>\ProTune_agent_service<unique identifier>.log` file (if the ProTune agent is a service), in a text editor.
- ▶ To open the file on UNIX machines, open the `<temp_directory>/m_agent_daemon<unique identifier>.log` file in a text editor.
- ▶ To increase the logging level, select Agent Settings from Start->Programs->ProTune->Advanced Settings (or open file `<ProTune_root>\launch_service\dat\br_lrch_server.cfg` in a text editor), and in the Log section, set AgentExtended to 1.

The **MI Listener log file** contains information about MI Listener communication with the ProTune agent and the Console.

To open the file, right-click the MI Listener Agent icon in the system tray of the MI Listener machine, and select **View Log**. Alternatively, open the latest `<temp_directory>\ProTune_agent_startup<unique identifier>.log` file (if the ProTune agent is a process), or `<temp_directory>\ProTune_agent_service<unique identifier>.log` file (if the ProTune agent is a service), in a text editor.

To increase the logging level, select **Start > Programs > ProTune > Advanced Settings > Agent Settings**, or open the `<ProTune_root>\launch_service\dat\br_Inch_server.cfg` file in a text editor. In the Log section, set AgentExtended to 1.

The **Console log file** contains information about communication activity between the Console and the MI Listener.

To open the file on Windows machines, open the `<temp_directory>\drv_log.txt` file in a text editor.

### **Verifying Connection Between ProTune Agent and MI Listener**

If there is a proper connection between the ProTune agent and the MI Listener:

- On Windows platforms, the agent icon's light in the system tray will turn from red to green.
- On UNIX platforms, a file called `<Local_machine_key>_connected_to_MI_Listener` will be created in the temporary directory of the ProTune agent machine. `Local_machine_key` is the value set in the Agent Configuration, as described in Chapter 13, "Working with Firewalls." The file will be removed when the ProTune agent disconnects from the MI Listener.
- On both UNIX and Windows platforms, the following message will appear in the ProTune agent log file: Notify Connected to MI Listener.

---

**Note:** The ProTune agent tries to connect to the MI Listener machine every Timeout seconds (as defined in the Agent Configuration). After a successful connection, if no Console has connected through this MI Listener to the agent after another Timeout, the ProTune will disconnect from the Console. On a Windows machine, the agent icon's light in the system tray will turn from green to red. On UNIX machines, the file `<Local_machine_key>_connected_to_MI_Listener` will be removed from the temporary directory in the ProTune agent machine. In both Windows and UNIX, the message `Disconnected from MI Listener` will appear in the ProTune agent log file.

---

### **UNIX Connection Errors**

After installing the `m_agent_daemon` as described in Chapter 13, "Working with Firewalls," you should receive a message: `m_agent_daemon` installed successfully.

#### **Agent Daemon Errors**

*ERROR: File `m_agent_daemon` doesn't exist.*

This error means that you are not in the same directory as the file (meaning not in `<ProTune_root>/bin` directory, or the file really doesn't exist, which indicates a problem with the installation).

*WARNING: Could not install `m_agent_daemon`, reason - user `<user_name>` is already running `m_agent_daemon` on this machine.*

This warning message occurs when a daemon of this name is already being run by the same user.

*ERROR: Could not install `m_agent_daemon`. Check log file `m_agent_daemon[xxx].log` in your temp directory.*

This error indicates that some error has occurred when loading the daemon. You should check the log file and consult the following troubleshooting tips.

## ProTune Agent Log File Errors

*Error - 10344 : Communication Error: -59961 : Failed to bind a socket while calling bind function.*

*Error -10344 : Communication Error: -59927 : Failed to create a TCP server for the HTTP channel's server.*

*Warning -29974 : Failed to create "router" server.*

These messages appear because the ProTune agent always tries to open port number 443 (because any agent can be a MI Listener, and the MI Listener always listens to this port), and in UNIX machines, this port cannot be opened by any user except for the root user. However, this will not interfere with using this agent for the Load Generator machine.

*Error -10343 : Communication error : -59981 : Failed to connect to remote host - <MI\_Listener\_name> .*

The MI Listener is not being run at the time of the connection attempt on the machine set in MI Listener Name in the Agent Configuration.

*Error -10343 : Communication error: -59928 : Unresolved server name .*

The name passed in MI Listener Name in the Agent Configuration is not a name, full name or IP address of a valid machine, or no value was set.

*Error -10343 : Communication error: -59928 : Unresolved server name .*

The name passed in Proxy Name in the Agent Configuration is not a name, full name or IP address of a valid machine.

*Error -10343 : Communication error: -59945 : Client failed to connect to a PROXY Server with the following settings:*

*(-server\_port=<proxy\_server\_port>)(-server\_fd\_primary=2)(-server\_type=8)(-allowed\_msg\_size=0)(-allowed\_msgs\_num=0)(-proxy\_configuration\_on)(-tcp\_tunnel\_configuration\_on).*

The Proxy Name field is empty.



*Error -10343 : Communication error: -59982 : Failed to connect to remote host - <MI\_Listener\_Name>. The remote address is not a valid address.*

*Error -10343 : Communication error: -59945 : Client failed to connect to a PROXY Server with the following settings:  
(-server\_name=<proxy\_server\_name>)(-server\_port=<proxy\_server\_port>)(-server\_fd\_primary=2)(-server\_type=8)(-allowed\_msg\_size=0)(-allowed\_msgs\_num=0)(-proxy\_configuration\_on)(-tcp\_tunnel\_configuration\_on).*

The Proxy Port set in Agent Configuration, has been set to the wrong port number.

*Error -10343 : Communication error: -59913 : NTLM authentication to proxy server error - connection to proxy refused.*

The proxy server is configured in for NTLM authentication and the Proxy User Name, Proxy Password and/or Proxy Domain are not set correctly in the Agent Configuration.

*Error -10343 : Communication error: - 59880 : Basic authentication to proxy server error - connection to proxy refused.*

The proxy server is configured in for Basic authentication and the Proxy User Name and/or Proxy Password are not set correctly in the Agent Configuration.

*Error -10343 : Communication error: -59907 : SSL connect error : verify host failed : wrong DNS test .*

This error occurs when you have set the Check Server Certificates setting to True, and have not issued a new certificate to the MI Listener machine (see Appendix F, “Working with Digital Certificates” for more details).

*Error -10343 : Communication error: -59907 : SSL connect error : certificate verify failed.*

*Error -10343 : Communication error: -59907 : SSL connect error : sslv3 alert handshake failure.*

*Error -10343 : Communication error: -59907 : SSL connect error : sslv3 alert bad certificate.*

*Error -10343 : Communication error: -59907 : SSL connect error : sslv3 alert certificate expired.*

These errors occur when you set the Check Server Certificates setting to True. See Appendix F, "Working with Digital Certificates" to learn how to issue a valid certificate.

*Error -10343 : Communication error: -59910 : SSL initialization error : Certificate not found .*

*Error -10343 : Communication error: -59910 : SSL initialization error : No such file or directory.*

*Error -10343 : Communication error: -59910 : SSL initialization error : system lib.*

These errors occur when the Client Certificate owner setting in the Agent Configuration is set to True, but no certificate was installed in the ProTune agent machine (see Appendix F, "Working with Digital Certificates" for more details).

### **MI Listener Log File Errors**

*Error - 10344 : Communication Error: -59961 : Failed to bind a socket while calling bind function.*

*Error -10344 : Communication Error: -59927 : Failed to create a TCP server for the HTTP channel's server.*

*Warning -29974 : Failed to create "router" server.*

This error means that another process on the MI Listener machine is occupying port 443 (for instance the IIS service).

*Error -10343 : Communication error: -59904 : SSL accept error : sslv3 alert certificate expired.*

These errors occur when you have set the Check Server Certificates setting to True, and the MI Listener's certificate is expired.

*Error -10343 : Communication error: -59904 : SSL accept error : sslv3 alert bad certificate.*

These errors occur when you have set the Check Server Certificates setting to True, and either:

- The MI Listener's certificate does not have a signature that is included in the ProTune agent's CA List.
- The MI Listener's certificate has a future verification date.

See Appendix F, "Working with Digital Certificates" to learn how to issue a valid certificate and how to add a Certification Authority to a CA list, or how to create a certificate with a new validation date.

*Error -10343 : Communication error: -59904 : SSL accept error : peer did not return a certificate.*

These errors indicate that the Check Client Certificates setting in the MI Listener Configuration is set to True, but the Client Certificate owner setting in the Agent Configuration is set to False.

*Error -10343 : Communication error: -59904 : SSL accept error : no certificate returned.*

These errors indicate that the Check Client Certificates setting in the MI Listener Configuration is set to True, and the Client Certificate owner setting in the Agent Configuration is set to True, but either:

- The ProTune agent's certificate does not have a signature that is included in the MI Listener's CA List.
- The ProTune agent's certificate has a future verification date.

See Appendix F, "Working with Digital Certificates" to learn how to issue a valid certificate and how to add a Certification Authority to a CA list, or how to create a certificate with a new validation date.

*Error -10343 : Communication error: -59904 : SSL accept error : no certificate returned.*

These errors indicate that the Check Client Certificates setting in the MI Listener Configuration is set to True, and the Client Certificate Owner setting in the Agent Configuration is set to True, but the ProTune agent's certificate has expired.

## General Connection Errors

These errors can occur when using all configurations.

If no errors appear both in the ProTune agent log, and the MI Listener log, but the agent does not connect to the MI Listener, make sure that the `FireWallServiceActive` attribute in the Firewall section in the `<ProTune_Installation>\dat\br_Inch_server.cfg` file on the ProTune agent machine, is set to 1.

## Verifying Connection Between the Console and Agent through the MI Listener

When there is a successful connection between the ProTune agent and the MI Listener, and the Console machine fails to connect, you should check the following:

- ▶ The **Name** field in the Load Generators dialog in the Console should match the name set in the **Local Machine Key** in the Agent Configuration.
- ▶ The **MI Listener** field in the **Load Generators > Details > Firewall** tab of the above host matches the name set in the **MI Listener Name** in the Agent Configuration.
- ▶ In the Tools menu of the Console, in the **Options > Timeout** tab, the **Load Generator Connect timeout** might need to be increased, because the Firewalls may slow down the communication.
- ▶ Make sure that the Console machine recognizes the ProTune agent machine (e.g., by using the ping utility). If this fails, there is a configuration problem in the system not related to ProTune, and it must be solved before the connection can be made.
- ▶ Make sure that the Console has successfully connected to the MI Listener by checking port 50500 on the MI Listener machine (you can use the netstat utility, on the MI Listener machine).

## Troubleshooting Remote Tuning

This section covers some issues you need to handle when using ProTune's remote tuning functions.

### **Not Viewing Information about a Host Running Windows NT**

If an NT host machine runs an old version of atl.dll, ProTune displays only the host's name without any information about its services. In addition, the following error message is displayed on the host machine:

"The ordinal 57 could not be located in the dynamic link library ATL.DLL."

**Solution:** If the host machine runs Windows NT, ensure that its version of atl.dll is 3.00.8449 or higher. (The atl.dll file is located in the \\<WINNT installation path>\system32\ directory.)

If the file on the host machine is an old one, replace it with a newer version and then register the new file. (For example, if the atl.dll file is located in C:\WINNT\system32\, register it by running regsvr32 C:\WINNT\system32\atl.dll from the command line).

### **Not Viewing Websphere Server Information**

To enable you to view (and tune) a WebSphere host from the Console machine, ensure that the WebSphere Administration Server is running on the host machine.



# B

---

## Working in Expert Mode

Advanced users can fine-tune the ProTune configuration settings while working in *Expert Mode*. In Expert mode, additional options are displayed in the Options dialog box and in the Load Generator Information dialog box. This appendix describes the additional settings that are available in the Expert mode:

- ▶ Entering Expert Mode
- ▶ Options - Agent Settings
- ▶ Options - General Settings
- ▶ Options - Debug Information Settings
- ▶ Options - Output Settings
- ▶ Options - Monitor Settings
- ▶ Load Generator Information - UNIX Environment Settings
- ▶ Load Generator Information - Connection Log Settings

### Entering Expert Mode

The ProTune Console Expert mode is intended for support personnel to provide access to system information. When you work in the Expert mode, the Console dialog boxes contain additional options for fine tuning the Console operation.

To activate the Expert mode, choose **Tools > Expert Mode**. An active Expert mode is indicated by a check mark.

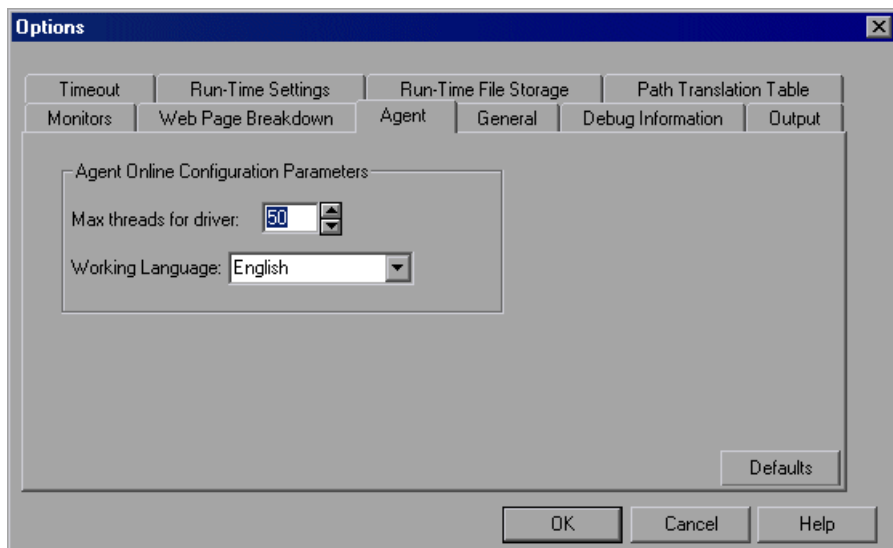
To exit the Expert mode, repeat the above process.

## Options - Agent Settings

The Agent settings allow you to customize the behavior of the ProTune agent on a remote load generator machine. Using the Options dialog box, you set the online configuration parameters for the agent.

**To set the Agent settings:**

- 1 Enter Expert mode (see above).
- 2 Choose **Tools > Options**. The Options dialog box appears. Select the **Agent** tab.



- 3 Select the maximum number of threads to be executed for the current User's driver.
- 4 Select the agent's working language (English or Japanese).
- 5 Click **OK** to accept the settings and close the dialog box.



## Options - General Settings

The General tab in the Options dialog box allows you to specify global settings for data table storage and multiple IP address allocation, and instruct ProTune not to collate log files.

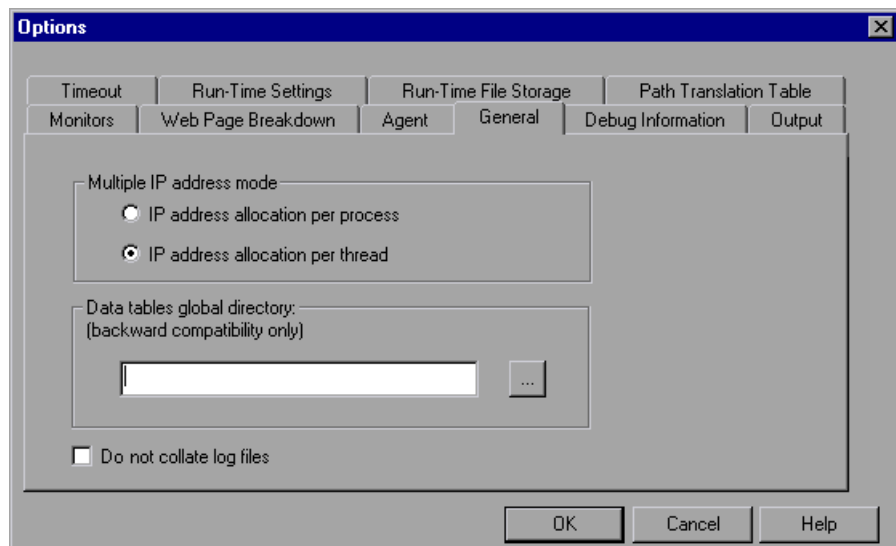
**Multiple IP address mode:** The mode used to allocate IP addresses when the multiple IP address option is enabled (**Session > Enable IP Spoofer**). The Console can allocate an IP address per process or per thread. Web Vusers require IP address allocation per process. WinSock Vuser IP addresses can be allocated per thread or per process. Allocation per thread results in a more varied range of IP addresses in a session.

**Data tables global directory:** The network location for data tables used as a source for parameter values. This setting is only required for scripts created with earlier versions of ProTune.

**Do not collate log files:** Instructs ProTune to collate only result files, and not log files.

**To set the General Expert mode settings:**

- 1 Choose **Tools > Options**. The Options dialog box appears. Select the **General** tab.



- 2 Select the Multiple IP address mode.
- 3 Enter the global directory for data tables.
- 4 If you want ProTune to collate only result files and not log files, check **Do not collate log files**.
- 5 Click **OK** to accept the settings and close the dialog box.

## Options - Debug Information Settings

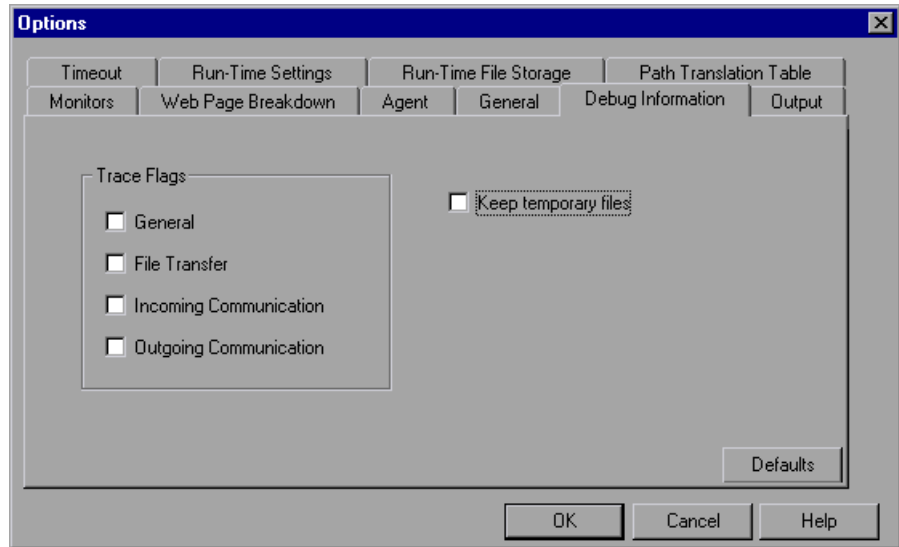
The Debug settings in the Options dialog box allow you to determine the extent of the trace to be performed during session execution. The debug information is written to the Output window.

The following trace flags are available: General, File Transfer, Incoming Communication, and Outgoing Communication. You only need to select the flags relating to your problem. For example, if you encounter specific problems with the transfer of files, select the File Transfer flag.

The ProTune agent and Console create some temporary files, which collect information such as the parameter file sent to the Vuser, the output compilation file, and the configuration file. The ProTune agent files are saved in *brr* folders in the TMP or TEMP directory of the agent machine. The Console files are saved in *lrr* folders in the TMP or TEMP directory of the Console machine. At the end of the session, all these files are automatically deleted. However, using the Debug Information Expert mode settings, you can instruct ProTune to keep these temporary files.

**To set the Debug Information settings:**

- 1 Choose **Tools > Options**. The Options dialog box appears. Select the **Debug Information** tab.



- 2 Select the check boxes for the desired trace flags.
- 3 To save the temporary run-time files, select the **Keep temporary files** check box.
- 4 Click **OK** to accept the settings and close the dialog box.

## Options - Output Settings

When Expert mode is enabled, the Options dialog box includes the **Output** tab. This tab contains the following settings:

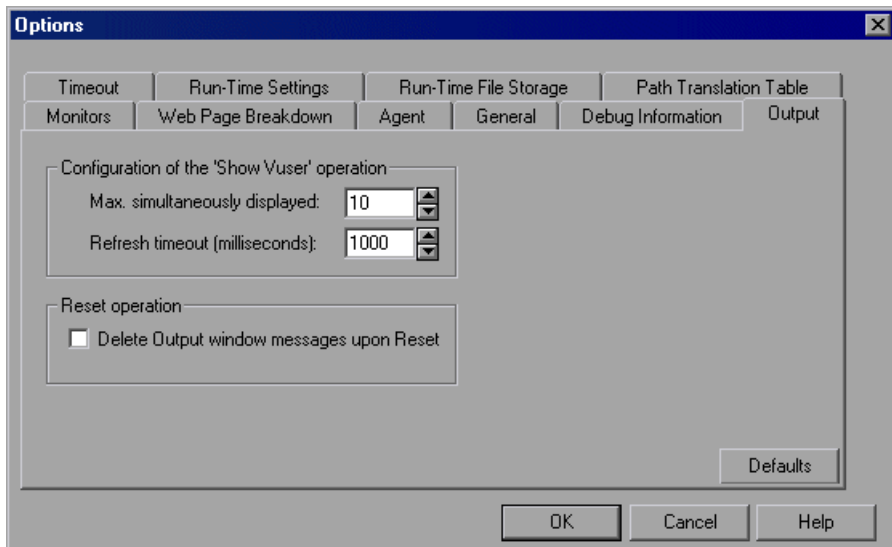
**Max simultaneously displayed:** Specifies the maximum number of Vuser logs that may be displayed simultaneously, as well as the maximum number of active UNIX, GUI, RTE, or Web Vusers that the Console should display by opening up Run-Time Viewers on your machine. The default number is 10.

**Refresh timeout:** Defines how often to refresh the Vuser log. The default is every 1000 milliseconds.

**Delete Output window messages upon Reset:** Instructs ProTune to clear all messages in the Output window when you reset a session step.

To set the Output settings:

- 1 Choose **Tools > Options**. The Options dialog box appears. Select the **Output** tab.



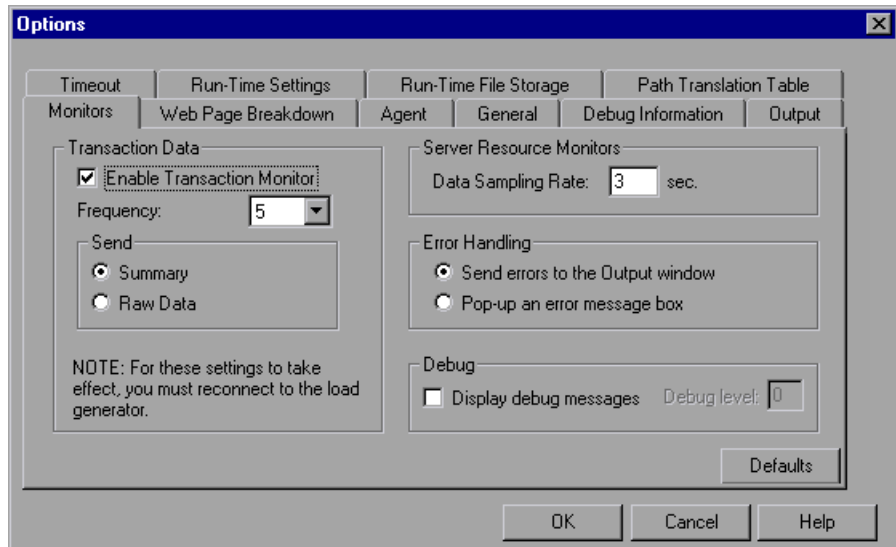
- 2 Specify the maximum number of Vuser logs to be displayed simultaneously, in the **Max. simultaneously displayed** box.

- 3 Specify the frequency at which ProTune refreshes the Vuser log, in the **Refresh timeout** box.
- 4 To clear the messages in the Output window when you reset a session step, select the **Delete Output window messages upon Reset** check box.
- 5 Click **OK** to accept the settings and close the dialog box.

## Options - Monitor Settings

Expert mode provides the following additional monitor setting:

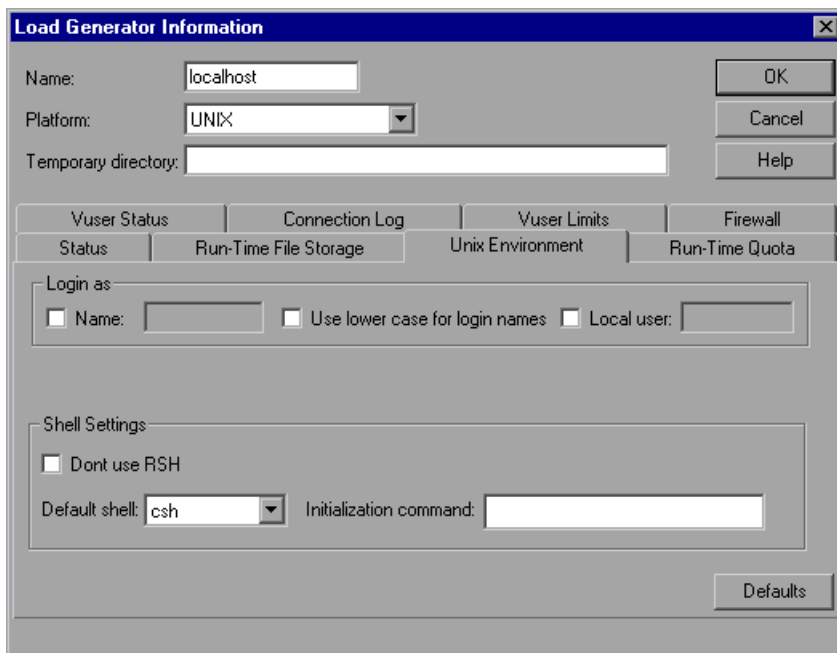
**Send Summary or Raw Data:** sends a summary of the data collected back to the Console, or sends all of the data in raw form. Sending the data in raw form saves time because the data does not need to be processed. However, since all of the data is being transferred to the Console, it may cause more network traffic. If the transfer speed is significant to you, it is recommended that you choose **Summary**.



## Load Generator Information - UNIX Environment Settings

Expert mode provides the following additional UNIX Environment setting:

**Local User:** UNIX load generators that use the *rsh* shell establish a connection as the current NT user (due to security considerations). To “mislead” *rsh* and log in as a user other than the current NT login, select the **Local user** check box and specify the desired UNIX login name. Since modifying the local user name is a security breach for *rsh*, this option should only be used when you encounter a problem connecting to the remote machine.




## Load Generator Information - Connection Log Settings

The Connection Log tab in the Load Generator dialog box allows you to view the standard output and standard errors generated as the Console connects to the selected UNIX load generator. You can also change the command that the Console sends to the remote bridge in order to connect to the load generator.

**To set the Connection Log settings:**



- 1** Click the **Generators** button, or select **Session > Load Generators**. The Load Generators dialog box opens.
- 2** Click **Connect** to change the Status of a load generator from Down to Ready.
- 3** Click the **Details** button. The Load Generator Information dialog box opens. Select the **Connection Log** tab.



You can view the rsh standard output and rsh standard errors generated as the Console sends the connection command to the selected UNIX load generator.

In the Bridge cmd box, enter a new command if you want to change the default bridge command being sent by the Console to the remote bridge in order to connect the UNIX load generator.



# C

---

## Performing Path Translation

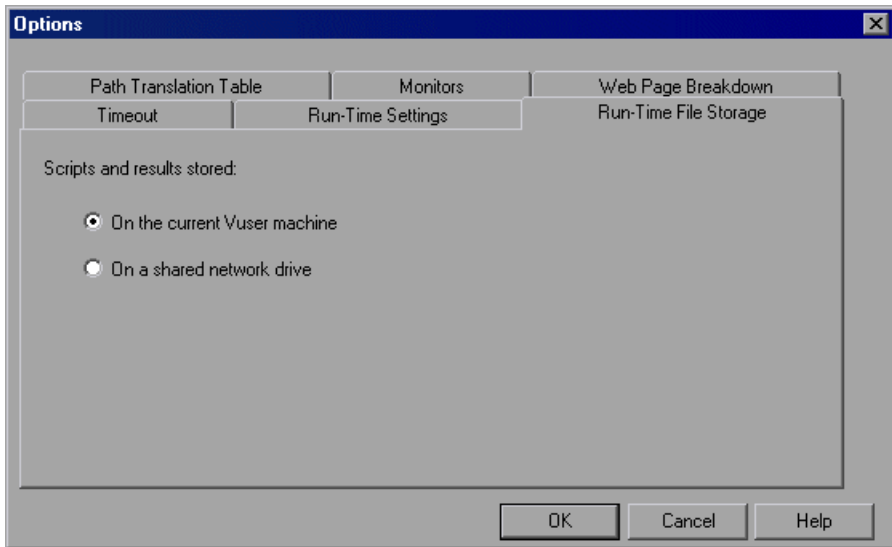
When you run a session step, ProTune gathers run-time data from the participating Vusers. By default, ProTune stores the data in temporary files on each Vuser machine. After the session step, the data is collated in the general results directory.

Alternatively, you can instruct ProTune to write the run-time data directly to a shared network drive. (See Chapter 5, “Configuring Session Steps.”) This method is not recommended, since it increases network traffic and necessitates path translation.

### Understanding Path Translation

Path Translation is a mechanism used by ProTune to convert a remote path name for the Console. A typical session step might have the ProTune Console running on a Windows-based machine and include multiple Vusers running on both Windows-based and UNIX load generators. One remote load generator may map the network drive as *F*, while another load generator maps the same drive as *H*. In a complex session step such as this, you need to ensure that all participating machines recognize the same network drive.

You instruct ProTune to store scripts and run-time data results on a shared network drive from the Run-time File Storage tab of the Options dialog box.



Result and script files stored on a shared network drive require you to perform path translation.

The Script view contains a list of all the Vuser scripts associated with a session step—and their locations. A script's location (path) is always based on the Console machine's mapping of that location. If a Vuser load generator maps to the script's path using a different name, path translation is required.

For example, assume that the Console is running on a Windows-based machine named *pc2*, and that a Vuser script is located on a network drive. The Console machine maps the network drive as *m:\lr\_tests*. If the remote Vuser machine (load generator) hosting the Vusers also maps the path as *m:\lr\_tests*, no translation is necessary. However, if the remote machine maps the path as another drive or path, for example *r:\lr\_tests*, you must translate the path to enable the load generator to recognize the script location.

Similarly, when saving run-time result files to a shared drive that is mapped differently by the Console and remote load generator, you must perform path translation.

Path translation is also effective across platforms—between Windows and UNIX. You use path translation to translate Windows-based paths (as seen by the Console) into paths recognized by the UNIX Vuser load generator.

## Adding Entries to the Path Translation Table

To translate a path from one Windows-based computer to another, or between Windows-based and UNIX machines, you create an entry in the Path Translation table. This table contains a list of paths translated into formats that can be recognized by different machines.

Each line of the Path Translation table has the following format:

```
<console_host><console_path><remote_path>[<remote_host>]
```

*console\_host*

The name or type of the machine that is running the Console. For example, if the Console is running on a Windows-based computer, you could type `win` in the host field. Alternatively, you could enter the name of the machine running the Console (for example, `LOADPC1`).

The value of *console\_host* can be:

|                 |   |
|-----------------|---|
| <b>hostname</b> | the name of the machine running the Console                 |
| <b>win</b>      | the Console is running on a Windows-based computer          |
| <b>unix</b>     | the Console is running on a UNIX machine                    |
| <b>all</b>      | the Console is running on a Windows-based or a UNIX machine |

*console\_path*

The path of a specific directory—as recognized by the Console. For example, if the directory *scripts* is located on the network drive *r*—as mapped by the Console—type the path *r:\scripts* in the *console\_path* field.

*remote\_path*

The path of a specific directory—as recognized by the remote machine. For example, if the directory *scripts* is located on the network drive *n*—as mapped by the remote load generator—type the path *n:\scripts* in the *remote\_path* field.

If a Vuser on the remote UNIX load generator recognizes the above path as */m/tests*, you would type this path in the *remote\_path* field.

*remote\_host*

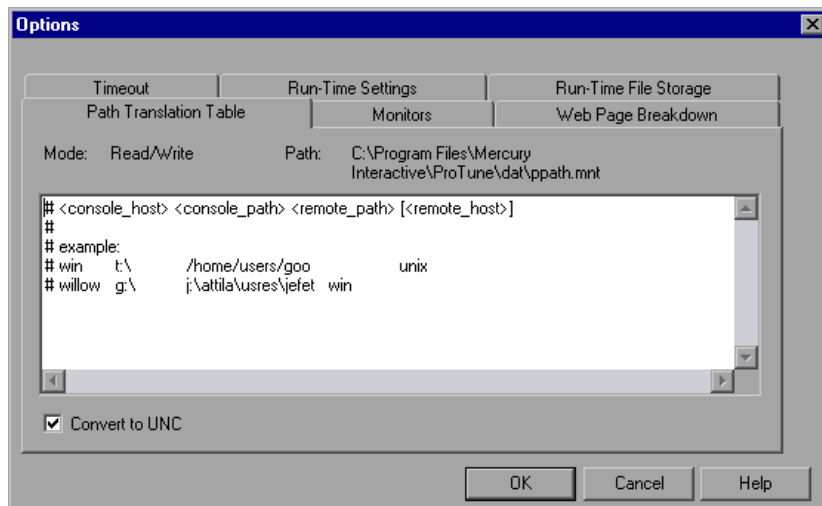
The name or type of the remote load generator. For example, if all the remote machines are UNIX workstations, you could type *unix* in the *remote\_host* field. The options for the *remote\_host* field are the same as the options for the *console\_host* field, listed above. The *remote\_host* parameter is optional.

## Editing the Path Translation Table

You maintain the Path Translation table using the ProTune Console. ProTune saves the Path Translation table as an ASCII file, *ppath.mnt*. This file, stored in ProTune\_directory/dat, has a one-line entry for each network path to translate.

**To edit the Path Translation table:**

- 1 Start the ProTune Console.
- 2 Choose **Tools > Options** and select the **Path Translation Table** tab. The Path Translation Table view opens.



- 3 Before you enter path translation information, consider using the Universal Naming Convention method. If your machines are Windows machines, you can tell the Console to convert all paths to UNC, and all machines will be able to recognize the path without requiring path translation. An example of UNC format is `\\machine_a\results`.

Select the Convert to UNC check box to tell ProTune to ignore the path translation table and to convert all paths to the Universal Naming Convention.

- 4 If your machines are not Windows machines and you require path translation, type the path information into the table. You can insert comments by typing the “#” symbol at the start of a line in the table.
- 5 Click **OK** to close the table and save the information.

## Path Translation Examples

The following section illustrates sample Path Translation Table entries.

Note that when you translate a Windows-based path to a UNIX path, you must enter the appropriate slashes—forward slashes for UNIX and back slashes for Windows-based paths.

The examples below show the use of the Path Translation table for a Windows-based Console called Merlin.

In the first example, Vusers are running on a Windows 2000 machine, Oasis. Merlin maps the network drive as f:, while Oasis maps it as g:\loadtest.

|        |     |              |       |
|--------|-----|--------------|-------|
| merlin | f:\ | g:\loadtest\ | Oasis |
|--------|-----|--------------|-------|

In the second example, Vusers are running on a UNIX machine, Ultra. Ultra maps the networks drive as /u/tests/load/.

|        |     |                |       |
|--------|-----|----------------|-------|
| merlin | f:\ | /u/tests/load/ | Ultra |
|--------|-----|----------------|-------|

In the third example, the mapping of the network drive by the remote load generator Jaguar, is identical to the Console's mapping, so no translation is required. This line can be excluded from the Path Translation table.

|        |     |     |        |
|--------|-----|-----|--------|
| merlin | n:\ | n:\ | Jaguar |
|--------|-----|-----|--------|

In the fourth example, all Windows-based Vuser load generators map the network drive as m:\loadtest.

|        |         |              |     |
|--------|---------|--------------|-----|
| merlin | l:\mnt\ | m:\loadtest\ | win |
|--------|---------|--------------|-----|

# D

---

## Working with Server Monitor Counters

When you configure the System Resource, Microsoft IIS, Microsoft ASP, ColdFusion, and SQL Server monitors, you are presented with a list of default counters that you can measure on the server you are monitoring. Using the procedure described below, you can create a new list of default counters by including additional counters, or deleting existing counters.

In addition, there are specific counters that are especially useful for determining server performance and isolating the cause of a bottleneck during an initial stress test on a server.

The following sections describe:

- Changing a Monitor's Default Counters
- Useful Counters for Stress Testing

### Changing a Monitor's Default Counters

You can change the default counters for the System Resource, Microsoft IIS, Microsoft ASP, ColdFusion, or SQL Server monitors by editing the *res\_mon.dft* file found in the ProTune/dat directory.

**To change the default counters:**

- 1** Open a new session and click the **Run** tab.
- 2** For each of the monitors, select the counters you want to measure.
- 3** Save the session and open the session *.irs* file with an editor.

- 4 Copy the MonItemPlus section of the each counter you selected into the *res\_mon.dft* file.
- 5 Count the number of new counters in the file and update the **ListCount** parameter with this number.

## Useful Counters for Stress Testing

Certain counters are especially useful for determining server performance and isolating the cause of a bottleneck during an initial stress test on a server.

The following is a list of counters that are useful for monitoring Web server performance:

| Object              | Counter                       |
|---------------------|-------------------------------|
| Web Service         | Maximum Connections           |
| Web Service         | Bytes Total/sec               |
| Web Service         | Current NonAnonymous Users    |
| Web Service         | Current Connections           |
| Web Service         | Not Found Errors              |
| Active Server Pages | Requests/sec                  |
| Active Server Pages | Errors/sec                    |
| Active Server Pages | Requests Rejected             |
| Active Server Pages | Request Not Found             |
| Active Server Pages | Memory Allocated              |
| Active Server Pages | Requests Queued               |
| Active Server Pages | Errors During Script Run Time |
| Memory              | Page Faults/sec               |
| Server              | Total Bytes/sec               |
| Process             | Private Bytes/Inetinfo        |



The following is a list of counters that are useful for monitoring SQL Server performance:

| Object          | Counter               |
|-----------------|-----------------------|
| SQLServer       | User Connections      |
| SQLServer       | Cache Hit Ratio       |
| SQLServer       | Net-Network Reads/sec |
| SQLServer       | I/O-Lazy Writes/sec   |
| SQLServer-Locks | Total Blocking Locks  |
| PhysicalDisk    | Disk Queue Length     |

The following is a list of counters that are useful for monitoring both Web and SQL server performance:

| Object       | Counter                |
|--------------|------------------------|
| Processor    | % Total Processor Time |
| PhysicalDisk | % Disk Time            |
| Memory       | Available Bytes        |
| Memory       | Pool Nonpaged Bytes    |
| Memory       | Pages/sec              |
| Memory       | Committed Bytes        |
| System       | Total Interrupts/sec   |
| Object       | Threads                |
| Process      | Private Bytes:_Total   |

---

**Note:** The % Disk Time counter requires that you run the `diskperf -y` utility at the command prompt and reboot your machine.

---



# E

---

## Configuring Multiple IP Addresses

When you run a session step, the Vusers on each load generator machine use the machine's IP address. You can define multiple IP addresses on a load generator machine to emulate a real-life situation in which users sit on different machines.

This appendix describes:

- Adding IP Addresses to a Load Generator
- Using the IP Wizard
- Configuring Multiple IP Addresses on UNIX
- Updating the Routing Table
- Enabling Multiple IP Addressing from the Console

## About Multiple IP Addresses

Application servers and network devices use IP addresses to identify clients. The application server often caches information about clients coming from the same machine. Network routers try to cache source and destination information to optimize throughput. If many users have the same IP address, both the server and the routers try to optimize. Since Vusers on the same load generator machine have the same IP address, server and router optimizations do not reflect real-life situations.

ProTune's multiple IP address feature enables Vusers running on a single machine to be identified by many IP addresses. The server and router recognize the Vusers as coming from different machines and as a result, the testing environment is more realistic.

---

**Note:** The maximum number of IP addresses that can be spoofed per network card for Windows NT SP3 is 35 IPs; Solaris (version 2.5.1) up to 255 IPs; Solaris (version 2.6 and higher) up to 8192 IPs.

---

### Applicable Protocols

The multiple IP address feature is applicable to the following protocols:

- ▶ **Client/Server:** DNS, Windows Sockets
- ▶ **Custom:** Java Vuser, Javascript Vuser, VB Vuser, VB Script Vuser
- ▶ **E-business:** FTP, Palm, SOAP, Web (HTTP/HTML) protocols, WinSock\Web Dual Protocol
- ▶ **ERP:** Oracle NCA, PeopleSoft 8 multi-lingual, Siebel-Web
- ▶ **Mailing Services:** Internet Messaging (IMAP), MS Exchange (MAPI), POP3, SMTP
- ▶ **Streaming Data:** Real
- ▶ **Wireless:** i-Mode, VoiceXML, WAP

This feature can be implemented on Windows and UNIX platforms.

## Adding IP Addresses to a Load Generator

ProTune includes an IP Wizard program that you run on each Windows NT or Windows 2000 load generator machine to create multiple IP addresses. You add new IP addresses to a machine once and use the addresses for all session steps. For information about adding IP addresses on UNIX machines, see “Configuring Multiple IP Addresses on UNIX” on page 672.

**The following procedure summarizes how to add new IP addresses to a load generator:**

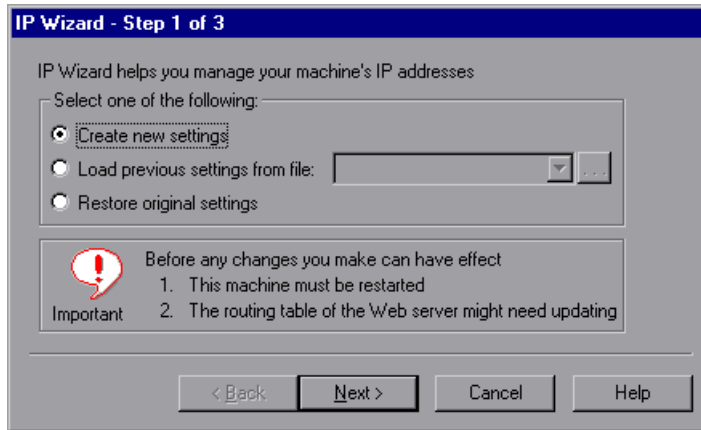
- 1** Run the IP Wizard on the load generator machine to add a specified number of IP addresses. Manually configure the new IP addresses for UNIX load generator machines.
- 2** Restart the machine.
- 3** Update the server’s routing table with the new addresses, if necessary.
- 4** Enable this feature from the Console. Refer to “Enabling Multiple IP Addressing from the Console” on page 674.

## Using the IP Wizard

The IP Wizard resides on each load generator machine. You run this process once to create and save new IP addresses on Windows machines. The new addresses can be a range of addresses defined by the Internet Assignment Numbers Authority. They are for internal use only, and cannot connect to the Internet. This range of addresses is the default used by the IP Wizard.

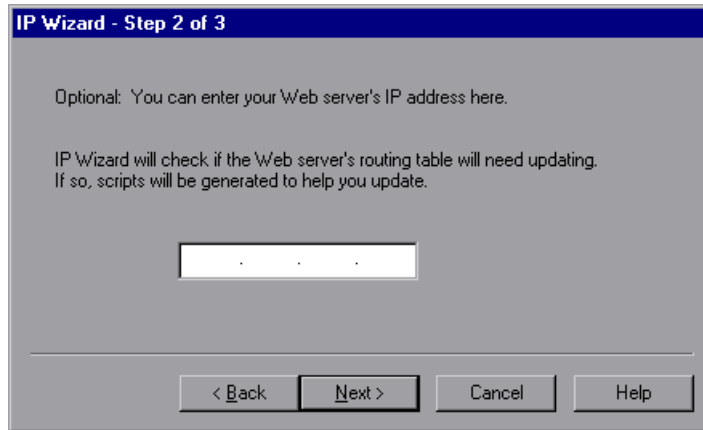
**To add new IP addresses to a load generator machine:**

- 1** Invoke the IP Wizard by clicking **Start > Programs > ProTune > Tools > IP Wizard**.




- 2** If you have an existing file with IP address settings, select **Load previous settings from file** and choose the file.
- 3** If you are defining new settings, select **Create new settings**.
- 4** Click **Next** to proceed to the next step. If you have more than one network card, choose the card to use for IP addresses and click **Next**.

The optional Web server IP address step enables the IP Wizard to check the server's routing table to see if it requires updating after the new IP addresses are added to the load generator.




- 5 To check the server's routing table directly after adding the addresses, enter the server IP address. Refer to "Updating the Routing Table" on page 673 for more information.
- 6 Click **Next** to see a list of the machine's IP address(es). Click **Add** to define the range of addresses.



IP addresses include two components, a *netid* and *hostid*. The submask determines where the netid portion of the address stops and where the hostid begins.

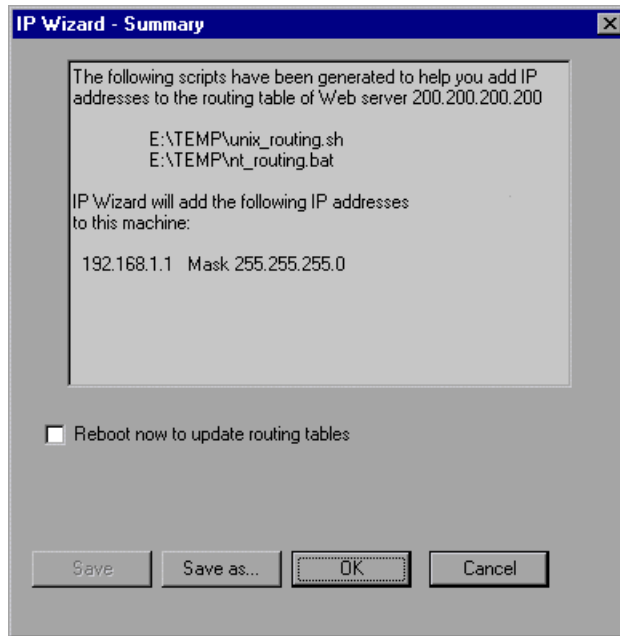
- 7 Select a class that represents the correct submask for the machine's IP addresses.
- 8 Specify the number of addresses to create. Select **Verify that new IP addresses are not already in use** to instruct the IP Wizard to check the new addresses. The IP Wizard will only add the addresses not in use.
- 9 Click **OK** to proceed.

After the IP Wizard creates the new addresses, the summary dialog box lists all of the IP addresses.





- 10 Click **Finish** to exit the IP Wizard. The IP Wizard Summary dialog box is displayed.



- 11 Note the address of the *.bat* file, and see “Updating the Routing Table” on page 673 for information about using the batch file to update the routing table, if necessary.
- 12 After you update the routing table, check **Reboot now to update routing tables** to initialize the NT device drivers with the new addresses.
- 13 Click **OK**.

## Configuring Multiple IP Addresses on UNIX

To configure multiple IP addresses on UNIX, manually configure the addresses on the load generator machine.

### **Solaris 2.5, 2.6, 7.0, 8.0**

To configure the `hme0` device to support more than one IP address:

- 1 Create entries in `/etc/hosts` for each hostname on your physical machine:

```
128.195.10.31 myhost
128.195.10.46 myhost2
128.195.10.78 myhost3
```

- 2 Create `/etc/hostname.hme0:n` files that contain the hostname for the virtual host `n`. Note that `hostname.hme0:0` is the same as `hostname.hme0`.

```
/etc/hostname.hme0 (Contains name myhost)
/etc/hostname.hme0:1 (Contains name myhost2)
/etc/hostname.hme0:2 (Contains name myhost3)
```

The above changes will cause the virtual hosts to be configured at boot time.

- 3 You can also directly enable/modify a logical hosts configuration by running `ifconfig` directly on one of the logical hosts, using the `hme0:n` naming scheme:

```
% ifconfig hme0:1 up
% ifconfig hme0:1 129.153.76.72
% ifconfig hme0:1 down
```

To verify the current configuration, use `ifconfig -a`.

## Linux

To define multiple IP addresses for a single Ethernet card, you need IP Aliasing compiled into the kernel. To do this, use the *ifconfig* command:

```
/sbin/ifconfig eth0:0 x.x.x.x netmask 255.255.x.x up
```

Substitute the new IP address for x.x.x.x, and insert the correct information for subnet mask. Place this command in the *rc.local* file so that it executes upon boot.

## HP 11.0 or higher

To define multiple IP addresses for a single Ethernet card, you need IP Aliasing compiled into the kernel. To do this, use the *ifconfig* command:

```
/sbin/ifconfig lan1:0 x.x.x.x netmask 255.255.x.x up
```

Substitute the new IP address for x.x.x.x, and insert the correct information for subnet mask. Place this command in the *rc.local* file so that it executes upon boot.

## Updating the Routing Table

Once the client machine has new IP addresses, the server needs the addresses in its routing table, so that it can recognize the route back to the client. If the server and client share the same netmask, IP class, and network, the server's routing table does not require modification.

---

**Note:** If there is a router between the client and server machines, the server needs to recognize the path via the router. Make sure to add the following to the server routing table: route from the Web server to the router, and routes from the router to all of the IP addresses on the load generator machine.

---

### To update the Web server routing table:

- 1 Edit the batch file that appears in the IP Wizard Summary screen. An example *.bat* file is shown below.

```
REM This is a bat file to add IP addresses to the routing table of a
server
REM Replace [CLIENT_IP] with the IP of this machine that the server
already recognizes
REM This script should be executed on the server machine

route ADD 192.168.1.50 MASK 255.255.255.255 [CLIENT_IP] METRIC 1
route ADD 192.168.1.51 MASK 255.255.255.255 [CLIENT_IP] METRIC 1
route ADD 192.168.1.52 MASK 255.255.255.255 [CLIENT_IP] METRIC 1
route ADD 192.168.1.53 MASK 255.255.255.255 [CLIENT_IP] METRIC 1
route ADD 192.168.1.54 MASK 255.255.255.255 [CLIENT_IP] METRIC 1
```

- 2 For each occurrence of [CLIENT\_IP], insert your IP address instead.
- 3 Run the batch file on the server machine.

## Enabling Multiple IP Addressing from the Console

Once you define multiple IP addresses, you set an option to tell the Console to use this feature.

### To enable multiple IP addressing from the Console:

- 1 In the Console Design view, select **Session > Enable IP Spoofer**.

---

**Note:** You must select this option before connecting to a load generator.

---

- 2 Use the **General Options** of the Console Expert Mode to specify how the Console should implement this feature.

For more information, refer to Appendix B, “Working in Expert Mode.”

# F

---

## Working with Digital Certificates

A Digital Certificate is an electronic “credit card” that establishes your credentials when doing business or other transactions on the Web. It is issued by a Certification Authority (CA). It contains the IP address of the machine for which it was issued, a validation date, and the digital signature of the certificate-issuing authority.

This appendix describes:

- Using Digital Certificates with Firewalls
- Creating and Using Digital Certificates

### Using Digital Certificates with Firewalls

When the MI Listener sends its Public Key to the ProTune agent, it always sends its certificate as well (this is the server-side certificate). The ProTune agent can be configured to authenticate the certificate which it received, as described in Chapter 13, “Working with Firewalls.” If the agent is configured to authenticate the certificate, it can verify whether the sender is really the machine that it claims to be by:

- Comparing the certificate's IP address with the sender's IP address.
- Checking the validation date.
- Looking for the digital signature in its Certification Authorities list.

The MI Listener may also require the ProTune agent to send a certificate at any point in the session. This is called the client-side certificate, as described in the MI Listener Configuration Settings in Chapter 13, “Working with Firewalls.” If the ProTune agent owns a certificate, it sends it to the MI Listener for the same authentication process. If the ProTune agent does not own a certificate, the communication might not be continued.

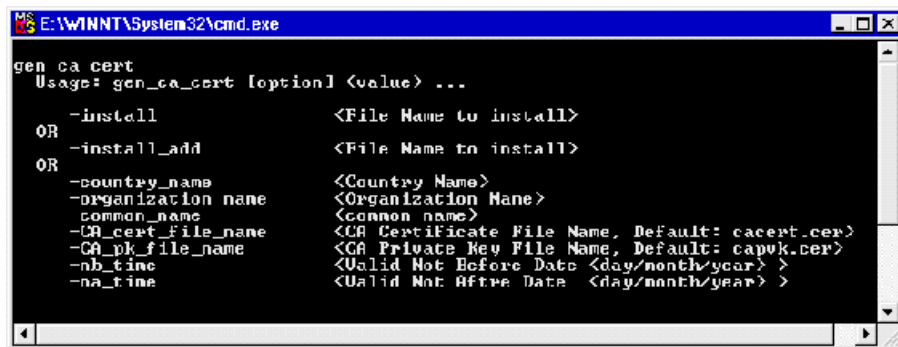
An SSL CA list and an SSL Certificate are included in each ProTune installation. This certificate is the same for all ProTune installations, which means that it can be obtained by third parties. Therefore, if you are interested in a more secure process, you should create your own Certificate Authority and include it in the list, and issue matching certificates for your machines.

## Creating and Using Digital Certificates

You create a Certification Authority using the `gen_ca_cert.exe` (on UNIX platforms `gen_ca_cert`) utility, and a Digital Certificate using the `gen_cert.exe` (on UNIX platforms `gen_cert`) utility. Both utilities can be used on UNIX and Windows platforms, using a command-line interface.

**To creating a Certificate Authority using `gen_ca_cert`:**

- 1 To view the format and usage, run the `gen_ca_cert` utility from the <ProTune root folder>\launch\_service\bin directory.



```

E:\WINNT\System32\cmd.exe
gen_ca_cert
Usage: gen_ca_cert [option] <value> ...

  -install          <File Name to install>
OR
  -install_add     <File Name to install>
OR
  -country_name    <Country Name>
  -organization_name <Organization Name>
  -common_name     <Common Name>
  -ca_cert_file_name <CA Certificate File Name, Default: cacert.cer>
  -ca_pk_file_name <CA Private Key File Name, Default: capvk.cer>
  -nb_time         <Valid Not Before Date <day/month/year> >
  -na_time         <Valid Not After Date <day/month/year> >

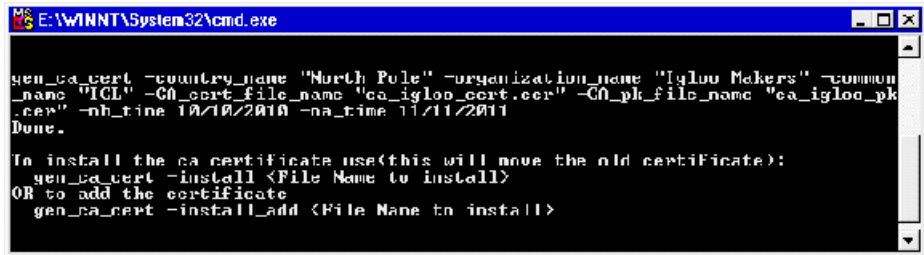
```

- 2 Create a new Certificate Authority by running the `gen_ca_cert` command with at least one of the options: `-country_name <country name>`, `-organization_name <organization name>` and `-common_name <the name of the CA>`.

This process creates two files in the directory from which the utility was run: the CA Certificate (`ca_cert.cer`), and the CA Private Key (`ca_pk.cer`). To provide different file names, use the `-CA_cert_file_name` and the `-CA_pk_file_name` options respectively.

By default, the CA is valid for three years, from the time that the CA is generated. To change the validation dates, use the options `-nb_time <beginning of validity in dd/mm/yyyy format>` and/or `-na_time <ending of validity in dd/mm/yyyy format>`.

The following example creates two files: `ca_igloo_cert.cer` and `ca_igloo_pk.cer` in the current directory.:



```

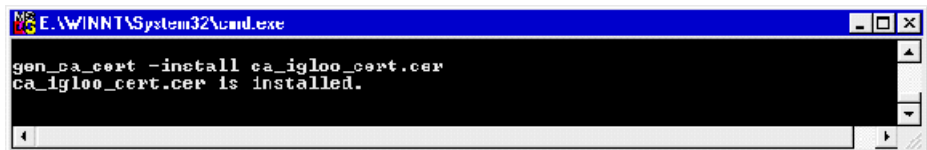
E:\WINNT\System32\cmd.exe
gen_ca_cert -country_name "North Pole" -organization_name "Igloo Makers" -common_name "IGL" -CA_cert_file_name "ca_igloo_cert.cer" -CA_pk_file_name "ca_igloo_pk.cer" -nb_time 10/10/2010 -na_time 11/11/2011
Done.

To install the ca certificate use(this will move the old certificate):
gen_ca_cert -install <File Name to install>
OR to add the certificate
gen_ca_cert -install_add <File Name to install>

```

- 3 To install this CA, use the `-install <name of certificate file>` option. This option replaces any previous CA list and creates a new one that includes only this CA.

To add the new CA to the existing CA list, use the `-install_add <name of certificate file>`.



```

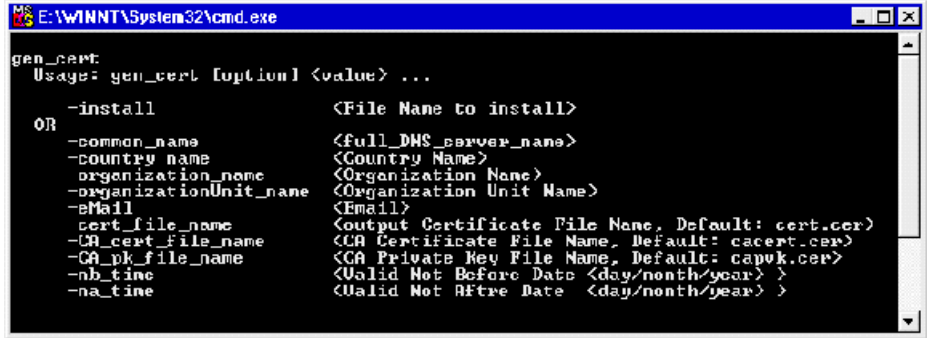
E:\WINNT\System32\cmd.exe
gen_ca_cert -install ca_igloo_cert.cer
ca_igloo_cert.cer is installed.

```

- 4 The `-install` and `-install_add` options install the certificate file only. Keep the private key file in a safe place and use it only for issuing certificates.

### To create a Digital Certificate using gen\_cert:

- 1 To view the format and usage, run the *gen\_cert* utility from the <ProTune root folder>\launch\_service\bin directory.



```

E:\WINNT\System32\cmd.exe
gen_cert
Usage: gen_cert [option] <value> ...

-install          <File Name to install>
OR
-common_name     <full_DNS_server_name>
-country_name    <Country Name>
-organization_name <Organization Name>
-organizationUnit_name <Organization Unit Name>
-eMail           <Email>
-cert_file_name  <Output Certificate File Name, Default: cert.cer>
-CA_cert_file_name <CA Certificate File Name, Default: cacert.cer>
-CA_pk_file_name <CA Private Key File Name, Default: capvk.cer>
-nb_time         <Valid Not Before Date <day/month/year> >
-na_time         <Valid Not After Date <day/month/year> >
  
```

- 2 Create a new Digital Certificate by running the *gen\_cert* command with at least one of the options: *-country\_name* <country name>, *-organization\_name* <organization name>, *-organization\_unit\_name* <organization unit name>, *-eMail* <email address> and *-common\_name* <the name, full name or IP address of the machine>.

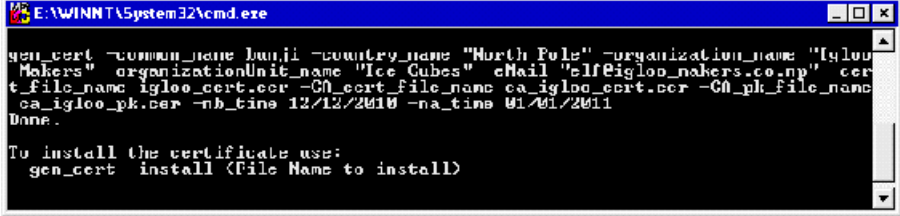
The CA Certificate and the CA Private Key files are necessary for the creation of the certificate. By default, it is assumed that they are in the current directory, and are named *cacert.cer* and *capvk.cer* respectively. In any other case, use the *-CA\_cert\_file\_name* and *-CA\_pk\_file\_name* options to give the correct files and locations.

In this process, the certificate file is created in the directory from which the utility was run. By default, the file name is *cert.cer*. To provide a different name, use the *-cert\_file\_name* option.

By default, the CA is valid for three years, from the time that the CA is generated. To change the validation dates, use the *-nb\_time* <beginning of validity in dd/mm/yyyy format> and/or *-na\_time* <ending of validity in dd/mm/yyyy format> options .



The following example creates the *igloo\_cert.cer* file in the current directory:



```
E:\WINNT\System32\cmd.exe
gen_cert -common_name buuji -country_name "North Pole" -organization_name "Igloo
Makers" -organizationUnit_name "Ice Cubes" -email "elf@igloo_makers.co.np" -cer
t_file_name igloo_cert.cer -CA_cert_file_name ca_igloo_cert.cer -CA_pk_file_name
ca_igloo_pk.cer -nb_time 12/12/2010 -na_time 01/01/2011
Done.

To install the certificate use:
gen_cert install <File Name to install>
```

- 3 If you wish to install this certificate, use the `-install <name of certificate file>` option. This option replaces any previous certificate, as it is possible to own only one certificate per machine.



# Host Resolution Functions Copyright Agreement

Copyright (c) 1980, 1983, 1985, 1987, 1988, 1989, 1990, 1993

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Portions Copyright (c) 1993 by Digital Equipment Corporation.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Digital Equipment Corporation not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

THE SOFTWARE IS PROVIDED "AS IS" AND DIGITAL EQUIPMENT CORP. DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL DIGITAL EQUIPMENT CORPORATION BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Copyright (c) 1996 by Internet Software Consortium.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND INTERNET SOFTWARE CONSORTIUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.



---

# Index

## A

- Acrobat Reader xiii
- activating rstatd 236
- Add Destination Machines for Network Delay Monitoring dialog box 269
- Add Load Generator dialog box 58
- Add Machine dialog box
  - DB2 monitor 389
- Add Oracle Measurements dialog box 407
- Add SAP Monitor Measurements dialog box 433
- Add Step dialog box 41
- Add Sybase Measurements dialog box 416
- Add TUXEDO Measurements dialog box 242, 531
- adding load generators 48
- adding measurements 208
- adding scripts 47
- adding session steps 29, 40
- adding steps 44
- administrator services 607, 608, 609
- ADO-DB Connection Rate 35
- ADO-DB SQL Query Rate 35
- agent 160
  - daemon 628
  - settings 646
- agent summary window 145
- AIX requirements for tuning agent 615
- alert actions 83
  - specifying 89
- alert categories 86
- alert conditions
  - specifying 88
- Alert Description pane 85, 87
- alert details
  - viewing 92
- alert notification types 84
- alert schemes 80
- alert types 78
- alerts
  - automatic assignment 23
  - configuring 85
  - creating 85
  - defining 77
  - deleting 85
  - recurring 83
  - session actions 84
  - specifying actions 83
  - specifying conditions 80
  - viewing descriptions 85
- Alerts in Output window 92
- Alerts window
  - Actions tab 83
  - Condition tab 82
- Antara FlameThrower monitor 247–258
- Apache
  - Add Measurements dialog box 280
  - monitor 280
- Application Deployment Resource monitors
  - Citrix MetaFrame XP monitor 510
- Application Deployment Solutions monitors 509–517
- Application Traffic Management monitors 535–540
- Application Traffic monitors
  - F5 BIG-IP monitor 535, 540
- Ariba
  - monitor 300
  - Monitor Configuration dialog box 300
- ASP
  - monitor 342

ATG Dynamo  
  monitor 304  
  Resources dialog box 304  
Auto Assign feature 23  
Auto Assign Monitors dialog box 24

## **B**

BEA WebLogic  
  monitor 357  
  Monitor Configuration dialog box  
    359  
benchmark profiles 96  
benchmark scheduling 106  
Bonk DoS Attack 38  
Books Online xiii  
BroadVision  
  monitor 310  
  Monitor Configuration dialog box  
    311

## **C**

CA 675  
canned scripts 29, 42  
certification authority 675  
Check Point FireWall-1  
  monitor 275  
  SNMP Resources dialog box 276  
choosing a profile 99  
choosing scripts 41  
Citrix MetaFrame XP  
  dialog box 511  
  monitor 510  
Citrix Monitor Configuration dialog box  
  510, 513  
ColdFusion  
  dialog box 320  
  monitor 320  
collating session results 117  
command line  
  Vuser script options 52  
component position  
  changing 15  
configuration settings  
  export, import 583–589  
  exporting 584

configuration settings (*cont'd*)  
  file types 584  
  importing 585  
  replicating across multiple hosts 583  
configuring  
  alerts 85  
  load generator 57–60  
  load generator settings 61  
  load generators 49  
  session steps 69–76  
configuring connection parameters 557  
configuring tuners 596, 599  
configuring tuning agents 591–600  
Connect to Host button 558  
Connect to Server dialog box 558  
connecting to database 629  
Connection Capacity (SSL) 32  
Connection Log tab 653  
connection parameters  
  configuring 557  
Connection Rate, HTTP 32  
Connection Settings dialog box 24  
Connections graph 225  
Connections per Second graph 225  
console\_host 657  
console\_path 658  
Context Sensitive Help xiv  
counters, for stress testing 662  
creating  
  alerts 85  
  goal-oriented sessions 29–53  
  manual profiles 101  
  new session steps 41  
  profiles 99  
creating goal oriented profiles 107–111  
creating new profiles 589  
custom queries  
  Oracle monitor 410  
Custom Scripts pane 43

## **D**

Data Points graph (online) 213  
Database Server Resource monitors 387–422  
  DB2 monitor 388  
  Oracle monitor 405

- Database Server Resource monitors (*cont'd*)
  - SQL Server monitor 412
  - Sybase monitor 415
- Database Server session steps 35
- database, connecting to 629
- DB2
  - monitor 388
  - Monitor Configuration dialog box 389
- DDoS attacks
  - SYN Flooding 39
  - UDP Echo 40
- debug
  - information settings 648
  - level 185
- default counter, changing 661
- Define Benchmark pane 106
- defining load behavior 111
- defining step goals 109
- defining step settings 110
- delaying step execution 97
- deleting alerts 85
- deleting profiles 100
- Denial of Service attacks 38
- Details button 142
- digital certificate
  - agent configuration settings 167
  - MI Listener configuration settings 170
  - overview 675
- disabling scripts 47
- disabling steps 45
- disabling tuners 599
- displayed rows, limiting 143
- distributing scripts by percentage 48
- DNS Request Rate 32
- documentation set xiv
- Done.Failed - Vuser state
  - Session Groups window 136
- Done.Passed - Vuser state
  - Session Groups window 136
- DoS attacks 38
  - Bonk 38
  - Ping of Death 38
  - Teardrop 39
- DoS tests
  - Targa3 39

- Down - Vuser state
  - Session Groups window 136
- duplicating steps 44
- Duration tab 103

## E

- EJB
  - monitor 452–467
  - Monitor Configuration dialog box 466
- Element Properties tab 16
- enabling scripts 47
- enabling steps 45
- enabling tuners 598
- ERP Server Resource monitors 431–437
- ERP/CRM Resource Server monitors
  - SAP Portal monitor 438
  - Siebel monitor 441
  - Siebel Server Manager monitor 445
- Error - Vuser state
  - Running Vusers graph 212
  - Session Groups window 136
- error handling 185
- Error log 139
- Error Statistics graph 212
- Ethernet-bus based network 547
- executing without delay 97
- Exiting - Vuser state
  - Session Groups window 136
- Expert mode 645–654
  - agent settings 646
  - connecting to UNIX load generator 653
  - debug settings 648
  - enabling and disabling 572
  - general settings 647
  - monitor settings 651
  - output settings 650
  - tuning 572
- Export Performance Settings dialog box 584
- exporting all settings 584
- exporting configuration settings 583–589
- exporting updated settings 584

## F

- F5 BIG-IP monitor 536
- filtering Vusers 132
- Finished - Vuser state
  - Running Vusers graph 212
- firewalls
  - configuring agent to operate over firewall 166
  - configuring the Console 170
  - Firewall Server monitors 275–278
  - installation configurations 158
  - installing MI\_Listener 169
  - monitors over 199–209
  - network monitoring 271
  - preparing for data collection 205
  - running Vusers 157–172
  - troubleshooting 635
- FTP
  - Connection Capacity 33
  - Get File Rate 33
  - Put File Rate 33
  - Server session steps 33
- Fujitsu INTERSTAGE
  - monitor 323
  - SNMP Resources dialog box 323
- Function Reference xiii

## G

- generating reports 153–155
- goal oriented profile scheduling 107–111
- goal oriented profiles 95, 96
  - creating 107–111
  - defining load behavior 111
  - defining step goal 109
  - defining step settings 110
- goal-oriented session 29–53
- Gradual Exiting - Vuser state
  - Session Groups window 136
- Graph Configuration dialog box 190
- graph time 189
- graphs, *See* online graphs

## H

- hardware
  - checking communications 622

- Hide button 132
- Hits per Second graph 220
- hme0 device 672
- host configuration
  - printing 611
  - reloading 611
- host parameters
  - committing 575
  - updating 575
- Host Properties dialog box 559, 561, 576, 581, 594
- hostid, IP address component 670
- Hostinfo utility 624
- hosts
  - removing 611
- hosts file 625
- HP, configuring IP addresses 673
- HP-UX requirements for tuning agent 616
- HTTP
  - Response per Second graph 221
  - HTTP Connection Rate 32
  - HTTP Downstream Bandwidth 33
  - HTTP Request Rate 32

## I

- IBM WebSphere MQ monitor 520–529
- IIS monitor 283
- IMAP Connection Capacity 34
- IMAP Search Mail 34
- IMAP Store Mail 34
- Import Performance Settings dialog box 586
- importing all settings 585
- importing configuration settings 583–589
- importing updated settings 585
- increasing number of Vusers 634
- Information tab 552, 567
- infrastructure session steps 32
- initial load 50
- initialization quota 65
- Initializing - Vuser state
  - Session Groups window 136
- initializing Vusers 131
- installing tuning agents
  - remotely 563
- installing tuning agents from CD
  - Windows 565



- installing tuning agents locally
    - UNIX 566
    - Windows 565
  - IP addresses
    - adding to a load generator 667
    - class 670
    - configuring multiple 665–674
    - configuring on HP 673
    - configuring on Linux 673
    - configuring on Solaris 672
    - enabling from the Console 674
    - hostid 670
    - IP Wizard 667
    - load generator machine 665
    - netid 670
    - per process 647
    - per thread 647
    - submask 670
  - iPlanet (NAS)
    - dialog box 329, 331
    - monitor 325
  - iPlanet (SNMP)
    - dialog box 289
    - monitor 289
  - iPlanet/Netscape
    - Add Measurements dialog box 286
    - monitor 285
- J**
- J2EE
    - monitor 492–507
  - Java
    - Java Performance monitors 451–489, 491–507
  - JProbe
    - dialog box (monitor) 469
    - monitor 468–470
- L**
- limiting displayed rows 143
  - Linux
    - configuring IP addresses 673
  - Linux requirements for tuning agent 617
  - listing available tuners 598
  - lists
    - load generator list 57–60
    - script list (goal-oriented session) 51–53
  - load behavior
    - defining 111
  - Load Behavior tab 111
  - load generator configuration
    - checking Console communication 623
    - connecting load generators 58
    - disabling load generators 58
    - disconnecting load generators 58
    - enabling load generators 58
    - Expert mode 652
    - firewall 67
    - initializing quota 65
    - limiting of Vusers 66
    - run-time files 62
    - UNIX shell 63
  - Load Generator Information dialog box 61
    - Firewall tab 67
    - Run-Time File Storage tab 62
    - Run-Time Quota tab 65
    - Status tab 61
    - Unix Environment tab 63
    - Vuser Limits tab 66
    - Vuser Status tab 68
  - load generators 57
    - adding 48, 59
    - adding an IP address 667
    - configuration 57–60
    - configuring 49
    - managing 57
    - modifying 59
    - multiple IP addresses 647
    - setting attributes 61–68
    - viewing load generator details 59
  - Load Generators window 58
  - Load Topology Template dialog box 14
  - loading profiles 587
  - lr\_user\_data\_point 213

## M

- Mail Server session steps 33
- manual benchmark scheduling 106
- manual profile scheduling 101
- manual profiles 95, 96
  - creating 101
  - script scheduling 105
  - session step scheduling 101
- MAPI Connection Capacity
  - session step
    - MAPI Connection Capacity 34
- MAPI Send Mail 34
- Measurement Configuration dialog box
  - Configuration tab 197
  - Description tab 197
- measurement frequency, setting 209
- measurements
  - selecting 18
- Media Player Client
  - monitor 429
- ML\_Listener 169
  - installing 205
- Microsoft
  - ASP monitor 342
  - IIS monitor 283
- Middleware Performance monitors 519–534
  - IBM WebSphere MQ monitor 520–529
  - TUXEDO monitor 530–534
- MMS Play Media 35
- modifying profile properties 99
- monitors
  - Application Deployment Solution 509–517
  - Application Traffic Management 535–540
  - automatic assignment 23
  - database server resources 387–422
  - ERP Server Resources 431–437
  - Firewall Server 275–278
  - Java Performance 451–489, 491–507
  - Middleware Performance 519–534
  - network 263–273
  - online 175–178
  - run-time 212
  - selecting 18

- monitors (*cont'd*)
  - streaming media 423–429
  - system resources 227–262
  - transaction 214
  - Web application server resources 299–386
  - Web resources 219–226
  - Web server resources 279–288
- monitors over firewall 199–209
  - adding and removing measurements 208
  - configuring measurement frequency 209
  - configuring properties 205
  - installing ML\_Listener 205
  - overview 200
  - preparing for data collection 205
- MQ monitor 520–529
- MS Active Server Pages dialog box 342
- MS IIS dialog box 283
- MS SQL Server dialog box 412
- multiple IP addresses 647
  - allocating 647
  - connecting to Console 626
  - enabling 647
- multiple script steps 41

## N

- netid, IP address component 670
- Network
  - Breakdown dialog box 273
  - Delay Time dialog box 269
  - Delay Time graph 272
  - Monitor Settings for Defined Path dialog box 269
  - Network Delay options 189
- Network monitor 263–273
  - configuring 268
  - determining bottlenecks 264
  - monitoring over a firewall 271
  - on UNIX 265
  - overview 263
- network segment delay, viewing 273

New Monitored Server Properties dialog box  
206

New Schedule dialog box 99

## O

ODBC SQL Connection Rate 36

ODBC SQL Query Rate 36

ODBC SQL Query Rate (reuse cursor) 36

online graphs 187

  configuring 187

  customizing display view 184

  data point 213

  exporting 198

  merging two graphs 192

  modifying a measurement scale 194

  opening graphs 182

  sampling rate 185

  viewing data offline 198

  x-axis style 188

  y-axis style 189

online monitors 175–178

  changing default counters 661

  configuring graphs 187

  configuring measurements 194

  debugging 185

  display type 189

  error handling 185

  graph time 189

  graphs 193

  line color 194

  pausing 194

  show/hide lines 195

  starting 180

  viewing data offline 198

online transaction monitoring

  adding transactions 216

  graphs 214

  setup 215

online Web server resource monitoring

  using a proxy server 297

Open a New Graph dialog box 182

Options dialog box

  Agent tab 646

  Debug Information tab 649

  General tab 647

Options dialog box (*cont'd*)

  Path Translation Table tab 659

  Run-Time File Storage tab 75

  Run-Time Settings tab 71

  Timeout tab 73

  Web Page Breakdown tab 217

Oracle

  custom queries 410

  Logon dialog box 407

  monitor 405

Oracle9iAS HTTP

  monitor 344

  Server Monitor Configuration dialog  
  box 344

output file 632

Output window 140–143

  clearing 142

  debugging information 632

  drilling down on log information 142

  Error log 139

  filtering messages 141

  refreshing 143

  saving messages to a file 142

  sorting messages 141

  Summary tab 141

  viewing message detail 142

Overlay Graphs dialog box 192

## P

packets 264

Pages Downloaded per Second graph 223

path translation

  debugging file locations 629

  defined 655

  editing the Path Translation Table 659

  examples 660

  script path 54

  session configuration 76

  using the Path Translation Table 657

pausing

  monitors 194

pausing Vusers 131

Pending - Vuser state

  Session Groups window 136

PeopleSoft session steps 36

- Performance Tuner Registry Console 597
- Ping of Death DoS Attack 38
- POP3 Connection Capacity 34
- POP3 Retrieve Mail 34
- prepared scripts 29
- print host configuration 611
- profiles 583
  - choosing 99
  - creating 99
  - creating new 589
  - deleting 100
  - file type 587
  - goal oriented 96
  - loading 587
  - manual, goal oriented 95
  - modifying properties 99
  - renaming 100
  - saving 587
- properties
  - specifying 16
- ProTune
  - agent 160
- Proxy Server 297

## R

- ramp down 101
- Ramp Down tab 104
- ramp up 101
- Ramp Up tab 102
- Ready - Vuser state
  - Running Vusers graph 212
  - Session Groups window 136
- Real Connection Capacity 35
- Real Server dialog box 426
- RealPlayer
  - Client monitor 428
  - Server monitor 426
- RealPlayer Play Media 35
- reboot host machines 609
- reconnecting the Console 610
- recurring alerts 83
- refreshing run-time settings 52
- registry, modifying 634
- relative script paths 54
- reload host configuration 611

- remote\_host 658
- remote\_path 658
- removing hosts 611
- removing measurements 208
- removing steps 44
- renaming profiles 100
- renaming scripts 47
- renaming steps 44
- rendezvous
  - Vuser state 136
- rendezvous points 52
- rendezvous tab 52
- reports
  - generating 153–155
- Reset button 131
- resetting Vusers 131
- restarting hosts and services 575
- results
  - collating 117
  - directory file structure 116
  - files for debugging 629
  - naming 114
  - specifying location for 114
- Retries per Second graph 224
- routing table 673
- rsh
  - checking Console connection 626
  - connection for UNIX network
    - monitor 267
  - running UNIX without 627
- rstatd process
  - activating 236
  - resource monitors 236
- Running - Vuser state
  - Running Vusers graph 212
  - Session Groups window 136
- running over firewall 157–172
- Run-Time graphs 211–216
- Run-Time settings
  - in the Console 51
- Run-Time Settings dialog box 51
- Run-Time Viewer 129, 132

**S**

- sampling rate 185
- SAP
  - monitor 433
  - Monitor Logon dialog box 433
- SAP Portal
  - monitor 438
- Save Configuration Profile dialog box 587
- saving profiles 587
- Schedule Builder
  - choosing a profile 99
  - deleting a profile 100
  - invoking 98
  - modifying a profile 99
  - renaming a schedule 100
  - selecting a profile 98–100
- Schedule Builder dialog box
  - Duration tab 103
  - Ramp Down tab 104
  - Ramp Up tab 102
- Schedule Builder window 98
- scheduling by benchmark 106
- scheduling by script 105
- scheduling by session step 101
- scheduling session steps 95
- script details, viewing 51
- Script Information dialog box 51
- Script List 47
- script log 144
  - default refresh rate 144
  - disable refreshing 144
  - run-time viewing 143
  - searching 145
  - viewing in text format 144
  - viewing Vuser 143
- Script Parameters pane 43
- script paths, relative 54
- script run-time settings 51
- script scheduling 105
- scripts
  - adding to script list 47
  - canned 29, 31
  - choosing 41
  - disabling 47
  - enabling 47
  - list 31
  - scripts (*cont'd*)
    - managing 46–49
    - renaming 47
    - sorting 47
- Security session steps 38
- Select measurements to monitor dialog box 19
- Select Online Graphs dialog box 183
- Server Configuration tree 552
- server monitors
  - adding and removing measurements 208
  - configuring properties 205
  - setting a measurement frequency 209
- Server Monitors dialog box 206
- server routing table 673
- service parameters
  - committing 575
  - updating 575
- services
  - starting 608
  - stopping 608
- session
  - collating results 117
  - creating a goal-oriented session 29–53
  - result directory 116
  - viewing output messages 140
- session configuration
  - path translation 76
  - run-time file location 74
  - run-time settings 70
  - specifying results location 114
  - timeout intervals 71
- session execution
  - activating additional Vusers 126
  - controlling individual Vusers 129
  - initializing Vusers 131
  - messages 140
  - monitoring active Vusers 136
  - running session unattended 123
- session step
  - ADO-DB Connection Rate 35
  - ADO-DB SQL Query Rate 35
  - Bonk DoS Attack 38
  - creating new 41
  - delaying execution 97

session step (*cont'd*)

- DNS Request Rate 32
  - duration 101
  - FTP Connection Capacity 33
  - FTP Get File Rate 33
  - FTP Put File Rate 33
  - HTTP Connection Rate 32
  - HTTP Downstream Bandwidth 33
  - HTTP Request Rate 32
  - IMAP Connection Capacity 34
  - IMAP Search Mail 34
  - IMAP Store Mail 34
  - MAPI Send Mail 34
  - MMS Play Media 35
  - ODBC SQL Connection Rate 36
  - ODBC SQL Query Rate 36
  - ODBC SQL Query Rate (reuse cursor)  
36
  - Ping of Death DoS Attack 38
  - POP3 Connection Capacity 34
  - POP3 Retrieve Mail 34
  - Real Connection Capacity 35
  - RealPlayer Play Media 35
  - running 121–132
  - SMTP Connection Capacity 33
  - SMTP Send Mail 33
  - specifying duration 103
  - specifying start time 97
  - SSL Connection Capacity 32
  - start without delay 97
  - SYN Flooding DDoS Attack 39
  - Targa3 DoS Test 39
  - TCP Connection Capacity 32
  - Teardrop DoS Attack 39
  - UDP Echo DDoS Attack 40
- session step execution 121–132
- specifying step start time 96
- session step scheduling 101
- Session Step Start dialog box 97
- session steps 31
- adding 29, 40
  - configuring 69–76
  - Database Server 35
  - FTP Server 33
  - infrastructure 32
  - Mail Server 33

session steps (*cont'd*)

- PeopleSoft 36
  - scheduling 95
  - Security 38
  - Siebel 38
  - Streaming Server 35
  - Web Server 32
  - Winsock 32
- session summary
- viewing 147–151
- Set Alert Name dialog box 86
- Set Results Directory dialog box
- local or remote location 115
- settings
- agent 646
  - debug 648
  - general 647
  - load generator 61–68
  - measurement frequency 209
  - monitors 651
  - output 650
  - timeout 71
- shared network drives 75
- show/hide a measurement
- online monitors 194
  - Transaction monitor 195
- Siebel
- monitor 441
- Siebel Server Manager
- monitor 445
- Siebel session steps 38
- SilverStream
- Add Measurements dialog box 350
  - monitor 349
- SiteScope Administration Console 26
- SiteScope monitor 259
- SiteScope Monitor Configuration dialog box  
260
- SiteScope resources 176
- Sitraka JMonitor 470
- Monitor Configuration dialog box  
485
- SMTP Connection Capacity 33
- SMTP Send Mail 33
- SNMP
- Resources monitor 238–240

- SNMP Resources
    - dialog box 238
  - Solaris
    - configuring IP addresses 672
  - Solaris requirements for tuning agent 614
  - sorting scripts 47
  - sorting Vusers 132
  - specifying alert actions 89
  - specifying alert conditions 88
  - specifying step start time 97
  - SQL Server
    - dialog box 412
    - monitor 412
  - SSL
    - agent configuration settings 167
    - MI Listener configuration settings 170
    - overview 675
  - SSL Connection Capacity 32
  - SSL Connections per Second graph 225
  - Start Agent Service dialog box 563
  - Start ProTune Session dialog box 12
  - Start Time tab 105
  - starting services 608
  - step goal, defining 109
  - step settings, defining 110
  - Steps
    - adding 44
    - disabling 45
    - duplicating 44
    - enabling 45
    - moving down 45
    - moving up 45
    - removing 44
    - renaming 44
  - steps
    - specifying order 45
  - Stopped - Vuser state
    - Session Groups window 136
  - stopping services 608
  - stopping tuning agent 610
  - stopping Vusers 131
  - stopping Vusers gradually 131
  - Streaming Media monitors 423–429
    - Media Player Client monitor 429
    - RealPlayer Client monitor 428
  - Streaming Media monitors (*cont'd*)
    - RealPlayer Server monitor 426
    - Windows Media Server monitor 424
  - Streaming Server session steps 35
  - Summary tab 141
  - Support Information xiv
  - Support Online xiv
  - supported operating systems 552
  - Sybase
    - Logon dialog box 416
    - monitor 415
  - SYN Flooding DDoS Attack 39
  - System Resource monitors 227–262
    - Antara FlameThrower monitor 247–258
    - SiteScope monitor 259
    - SNMP Resources monitor 238–240
    - TUXEDO monitor 241–246
    - UNIX Resources monitor 234–236
  - System Topology window 13
- T**
- Targa3 DoS Test 39
  - TCP Connection Capacity 32
  - TCP/IP setup 624
  - Teardrop DoS Attack 39
  - Throughput graph 220
  - timeout settings 71
  - topology 9
    - adding component to 15
    - creating 9, 15
    - modifying 15
  - topology components
    - clearing 15
    - connecting 15
    - deleting all components 15
    - deleting single component 15
    - zooming in on 15
    - zooming out 15
  - topology design palette 11
  - topology diagram
    - building 11
    - creating from a template 14
    - exporting 15

- topology diagram (*cont'd*)
  - importing 13
  - saving 15
- topology templates 10
- Total Transactions per Second (Passed) graph 214
- TowerJ
  - monitor 489
- Transaction monitor 211–216
- transactions
  - failed 631
  - Transaction Response Time graph 214
  - Transactions dialog box 138
  - Transactions per Second (Failed, Stopped) graph 214
  - Transactions per Second (Passed) graph 214
- Transactions dialog box 138
- troubleshooting
  - Console 621
  - firewalls 635
  - monitors 541–548
  - network considerations 546
- tunable applications 553
- Tune tab 552
- Tune tab functions 607–611
- tuners
  - configuring 596, 599
  - disabling 599
  - enabling 598
  - listing available 598
  - Performance Expert Registry 596
- tuning
  - changing host and service parameters 572
  - Connect to Host button 558
  - connecting to host 557
  - connection parameter configuration 557
  - Expert mode 572
  - from the Console 551–618
  - Information tab 552, 567
  - parameter value colors 572
  - predefined usernames and passwords 560
  - procedure 554
  - tuning (*cont'd*)
    - recommended values 572
    - secure shell (UNIX) 562
    - Server Configuration tree 552
    - specifying host 557
    - specifying port for agent 560
    - SSL use 560
    - summary host information 567
    - Tune tab 552
    - tuner username and password 560
    - Tuning tab 552, 572
    - updating host parameters 574
    - viewing host and service information 568
    - viewing host details 567
    - viewing host information 567
    - viewing host properties 559, 561, 576, 581
  - tuning agent
    - stopping 610
  - tuning agent port
    - changing 593
  - tuning agents 591
    - AIX requirements 615
    - batch files (Windows and UNIX) 593
    - configuring 591–600
    - host activity example 564
    - HP-UX requirements 616
    - installing and starting 562
    - installing from CD (Windows) 565
    - installing from the Console 563
    - installing locally (UNIX) 566
    - installing locally (Windows) 565
    - Linux requirements 617
    - Solaris requirements 614
    - supported operating systems 552
    - tunable applications 553
  - tuning hosts and services 551
  - tuning procedure flow 554
  - tuning recommendations 573
  - Tuning tab 552, 572
- TUXEDO
  - monitor 241–246, 530–534
  - monitor measurements 244, 533



**U**

- UDP Echo DDoS Attack 40
- UNIX
  - activating rstatd 236
  - connection to load generator 653
  - Resources monitor 234–236
  - rsh 626
  - shell 626
  - without rsh 627
- Unix Kernel Statistics dialog box 234
- UNIX ProTune agent 162
- Update Service Confirmation dialog box 575
- updating host parameters 574
- User-Defined Data Points graph 213

**V**

- viewing alert details 92
- viewing host information 567
- viewing script details 51
- viewing session summary 147–151
  - overview 147
- viewing Vuser execution 132
- viewing Vuser script log 132
- viewing Vusers 135–145
  - agent summary 145
  - Output window 140
  - overview 135
- viewing Web page with error 144
- Vuser information 144
- Vuser percentage
  - modifying 48
- Vuser script log
  - viewing 132
- Vuser scripts
  - command line options 52
- Vuser states
  - Run-Time graphs 212
- Vusers
  - activating additional during session execution 126
  - error, warning, and notification messages 140
  - filtering 132
  - initializing 131
  - monitoring 136

**Vusers (*cont'd*)**

- pausing 131
- resetting 131
- running 131
- sorting 132
- status in Session Groups window 136
- stopping 131
- stopping gradually 131
- viewing 135–145
- viewing execution 132
- Vuser Log 143
- Vuser script log 132
- Vusers with Error Statistics graph 212
- Vusers window 18

**W**

- Web Application Server Resource monitors 299–386
  - Ariba monitor 300
  - ATG Dynamo monitor 304
  - BroadVision monitor 310
  - ColdFusion monitor 320
  - Fujitsu INTERSTAGE monitor 323
  - iPlanet (NAS) 325
  - Microsoft ASP 342
  - Oracle9iAS HTTP monitor 344
  - SilverStream monitor 349
  - WebLogic (JMX) monitor 357
  - WebLogic monitor 353
  - WebSphere (EPM) monitor 377
  - WebSphere monitor 363
- Web page breakdown, enabling 217
- Web Resource monitors 219–226
- Web Server Resource monitors 279–288
  - Apache monitor 280
  - iPlanet (SNMP) monitor 289
  - iPlanet/Netscape monitor 285
  - Microsoft IIS monitor 283
- Web Server session steps 32
- Web Vusers, multiple IP addresses 647
- WebLogic
  - (JMX) monitor 357
  - (SNMP) Resources dialog box 353
  - monitor 353

WebSphere

(EPM) monitor 377

(EPM) Monitor Configuration dialog  
box 385

monitor 363

Monitor Configuration dialog box  
366

Windows

Media Server - Add Measurements  
dialog box 424

Media Server monitor 424

Resources dialog box 229

Resources monitor 229–233

Windows Resource monitors

Windows Resources monitor 229–233

Winsock session steps 32





MERCURY INTERACTIVE

Mercury Interactive Corporation  
1325 Borregas Avenue  
Sunnyvale, CA 94089  
Tel. (408)822-5200 (800) TEST-911  
Fax. (408)822-5300

**Main Telephone:** (408) 822-5200  
**Sales & Information:** (800) TEST-911, (866) TOPAZ-4U  
**Customer Support:** (877) TEST-HLP  
**Fax:** (408) 822-5300

**Home Page:** [www.mercuryinteractive.com](http://www.mercuryinteractive.com)  
**Customer Support:** [support.mercuryinteractive.com](http://support.mercuryinteractive.com)



\* PTC ONUG1. 5/ 01 \*