

LoadRunner®
Controller User's Guide
Version 7.6



MERCURY INTERACTIVE

LoadRunner Controller User's Guide, Version 7.6

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 pending. All rights reserved.

Mercury Interactive, the Mercury Interactive logo, WinRunner, XRunner, FastTrack, LoadRunner, LoadRunner TestCenter, TestDirector, TestSuite, WebTest, Astra, Astra SiteManager, Astra SiteTest, Astra QuickTest, QuickTest, Open Test Architecture, POPs on Demand, TestRunner, Topaz, Topaz ActiveAgent, Topaz Delta, Topaz Observer, Topaz Prism, Twinlook, ActiveTest, ActiveWatch, SiteRunner, Freshwater Software, SiteScope, SiteSeer and Global SiteReliance are registered trademarks of Mercury Interactive Corporation or its wholly-owned subsidiaries, Freshwater Software, Inc. and Mercury Interactive (Israel) Ltd. in the United States and/or other countries.

ActionTracker, ActiveScreen, ActiveTune, ActiveTest SecureCheck, Astra FastTrack, Astra LoadTest, Change Viewer, Conduct, ContentCheck, Dynamic Scan, FastScan, LinkDoctor, ProTune, RapidTest, SiteReliance, TestCenter, Topaz AIMS, Topaz Console, Topaz Diagnostics, Topaz Open DataSource, Topaz Rent-a-POP, Topaz WeatherMap, TurboLoad, Visual Testing, Visual Web Display and WebTrace are trademarks of Mercury Interactive Corporation or its wholly-owned subsidiaries, Freshwater Software, Inc. and Mercury Interactive (Israel) Ltd. in the United States and/or other countries. The absence of a trademark from this list does not constitute a waiver of Mercury Interactive's intellectual property rights concerning that trademark.

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 USA
Tel: (408) 822-5200
Toll Free: (800) TEST-911, (866) TOPAZ-4U
Fax: (408) 822-5300

© 1999 - 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.

Table of Contents

Welcome to LoadRunner.....	xi
Online Resources.....	xi
LoadRunner Documentation Set.....	xii
Using the LoadRunner Documentation Set.....	xiii
Typographical Conventions.....	xiv

PART I: UNDERSTANDING LOADRUNNER

Chapter 1: Introduction	3
Application Load Testing	3
The LoadRunner Solution	4
Using LoadRunner.....	5
Working with LoadRunner	6
LoadRunner Vuser Technology.....	7
LoadRunner Vuser Types.....	8
Chapter 2: The LoadRunner Testing Process	15
Step I: Planning the Test	16
Step II: Creating the Vuser Scripts	16
Step III: Creating the Scenario	16
Step IV: Running the Scenario	17
Step V: Monitoring a Scenario	18
Step VI: Analyzing Test Results	18
Chapter 3: Load Test Planning	19
About Load Test Planning.....	19
Analyzing the Application	20
Defining Testing Objectives	23
Planning LoadRunner Implementation.....	25
Examining Load Testing Objectives.....	31

Chapter 4: The LoadRunner Controller at a Glance	35
Opening the Controller	35
Introducing the LoadRunner Controller	38
Managing Scenario Files	42
Running a Scenario	44

PART II: DESIGNING A SCENARIO

Chapter 5: Creating a Manual Scenario	49
About Creating a Scenario	50
Creating Vuser Groups	50
Configuring Vusers in a Vuser Group	56
Configuring Vuser Run-Time Settings	59
Configuring Load Generators	61
Configuring Load Generator Settings	64
Configuring WAN Emulation Settings.....	72
Configuring Scripts	80
Using Relative Paths for Scripts.....	83
Chapter 6: Creating a Manual Scenario Using the Percentage Mode	87
About Creating a Manual Scenario Using the Percentage Mode.....	87
Defining the Total Number of Vusers	88
Assigning Properties to Scripts	89
Configuring Scripts	92
Converting a Scenario to the Vuser Group Mode.....	95
Chapter 7: Creating a Goal-Oriented Scenario	97
About Planning a Goal-Oriented Scenario.....	97
Defining Scenario Goals	100
Assigning Properties to Scripts	105
Configuring Scripts	108
Chapter 8: Scheduling a Scenario	113
About Scenario Scheduling	114
Delaying the Start of a Scenario	115
Selecting a Schedule	116
Scheduling a Scenario	117
Scheduling Vuser Groups	121
Adding Vusers to a Scheduled Scenario	125

Chapter 9: Using Rendezvous Points	127
About Using Rendezvous Points	127
Setting the Rendezvous Attributes	129
Setting the Rendezvous Policy	131
Disabling and Enabling Rendezvous Points.....	133
Disabling and Enabling Vusers at Rendezvous Points.....	133
Viewing Rendezvous Information	135
Chapter 10: Configuring a Scenario	137
About Configuring a Scenario.....	137
Configuring Scenario Run-Time Settings.....	138
Setting Timeout Intervals	139
Setting the Run-Time File Location	142
Specifying Path Translation	144
Chapter 11: Preparing to Run a Scenario	145
About Preparing to Run a Scenario	145
Specifying a Results Location	146
Results Directory File Structure	148
Collating Results	149
Setting Scenario Summary Information	151
Chapter 12: Managing Scenarios Using TestDirector	153
About Managing Scenarios Using TestDirector	153
Connecting to and Disconnecting from TestDirector	154
Opening Scenarios from a TestDirector Project	157
Saving Scenarios to a TestDirector Project	158
Saving Results to a TestDirector Project	160
Adding Vuser Scripts from a TestDirector Project	161

PART III: EXECUTING A SCENARIO

Chapter 13: Running a Scenario	167
About Running a Scenario	167
Running an Entire Scenario	169
Controlling Vuser Groups	170
Controlling Individual Vusers.....	173
Manually Releasing Vusers from a Rendezvous.....	175
Manually Adding Vusers to a Running Scenario	176

Chapter 14: Viewing Vusers During Execution	181
About Viewing Vusers During Execution	181
Monitoring Vuser Status	182
Viewing the Output Window	187
Viewing the Vuser Script Log	190
Logging Execution Notes	192
Viewing the Agent Summary	193
Chapter 15: Working with Firewalls	195
About Using Firewalls in LoadRunner	195
Overview of Running or Monitoring over the Firewall	197
Configuring the LoadRunner Agents in LAN1	199
Configuring the Firewall to Allow Agent Access.....	207
Installing and Configuring the MI_Listener in LAN2.....	208
Configuring the Controller to Run or Monitor Vusers over the Firewall.....	209

PART IV: MONITORING A SCENARIO

Chapter 16: Online Monitoring	215
About Online Monitoring	216
Starting the Monitors	218
Opening Online Monitor Graphs	220
Customizing the Graph Display View	222
Configuring Online Monitors	222
Setting Monitor Options	223
Configuring Online Graphs	225
Merging Graphs	229
Understanding Online Monitor Graphs	230
Configuring Online Measurements	231
Exporting Online Monitor Graphs	236
Viewing Data Offline	236
Chapter 17: Monitoring over a Firewall	237
About Monitoring over the Firewall	238
Installing Monitors over Firewall	238
Installing MI_Listener	243
Preparing for Data Collection	243
Configuring Server Monitor Properties	244
Adding and Removing Measurements	246
Configuring Measurement Frequency	247
Configuring the Network Delay Monitor over a Firewall.....	247

Chapter 18: Remote Performance Monitoring	249
About Remote Performance Monitoring	249
Installing the Remote Performance Monitor Server	250
Configuring the Remote Performance Monitor User Settings.....	251
Connecting to the LoadRunner Remote Performance Monitor	254
Monitoring Load Test Data	256
Viewing Online Graphs.....	256
Customizing Online Graph Settings.....	259
Chapter 19: Run-Time and Transaction Monitoring	265
About Run-Time and Transaction Graphs	265
Run-Time Graphs	266
User-Defined Data Points Graph	267
Transaction Monitor Graphs	268
Enabling the Transaction Monitor	269
Adding Transactions to a Script	270
Enabling Web Page Breakdown	271
Chapter 20: Web Resource Monitoring	273
About Web Resource Monitoring.....	273
Hits per Second Graph	274
Throughput Graph	274
HTTP Responses per Second Graph	275
Pages Downloaded per Second Graph	277
Retries per Second Graph	278
Chapter 21: System Resource Monitoring	279
About System Resource Monitoring.....	279
Configuring the Windows Resources Monitor	281
Configuring the UNIX Resources Monitor	286
Configuring an rstatd Daemon on UNIX	289
Configuring the SNMP Resources Monitor	290
Configuring the Antara FlameThrower Monitor	292
Configuring the SiteScope Monitor	304
Chapter 22: Network Monitoring	307
About Network Monitoring.....	307
Network Monitoring from a UNIX Source Machine	309
Configuring the Network Monitor	312
Viewing the Network Delay Time Graph	316
Chapter 23: Firewall Server Performance Monitoring	319
About the Firewall Server Monitor.....	319
Configuring the Check Point FireWall-1 Server Monitor	319

Chapter 24: Web Server Resource Monitoring	323
About Web Server Resource Monitors.....	323
Configuring the Apache Monitor	324
Configuring the Microsoft IIS Monitor	327
Configuring the iPlanet/Netscape Monitor	329
Configuring the iPlanet (SNMP) Monitor	334
Monitoring Using a Proxy Server	342
Chapter 25: Web Application Server Resource Monitoring	343
About Web Application Server Resource Monitors	344
Configuring the Ariba Monitor	344
Configuring the ATG Dynamo Monitor	348
Configuring the BroadVision Monitor	355
Configuring the ColdFusion Monitor	365
Configuring the Fujitsu INTERSTAGE Monitor	368
Configuring the iPlanet (NAS) Monitor	370
Configuring the Microsoft Active Server Pages Monitor	387
Configuring the Oracle9iAS HTTP Monitor	389
Configuring the SilverStream Monitor	394
Configuring the WebLogic (SNMP) Monitor	397
Configuring the WebLogic (JMX) Monitor	401
Configuring the WebSphere Monitor	407
Configuring the WebSphere (EPM) Monitor	421
Chapter 26: Database Resource Monitoring	431
About Database Resource Monitoring.....	431
Configuring the DB2 Monitor	432
Configuring the Oracle Monitor	449
Configuring the SQL Server Monitor	456
Configuring the Sybase Monitor	459
Chapter 27: Streaming Media Monitoring	467
About Streaming Media Monitoring.....	467
Configuring the Windows Media Server Monitor	468
Configuring the RealPlayer Server Monitor	470
Viewing the RealPlayer Client Online Graph.....	472
Viewing the Media Player Client Online Graph.....	473
Chapter 28: ERP Server Resource Monitoring	475
About ERP Server Resource Monitoring	475
Setting Up the SAP Monitor.....	476
Configuring the SAP Monitor	476

Chapter 29: Java Performance Monitoring	481
About Java Performance Monitoring	481
EJB Performance Monitoring	482
JProbe Performance Monitoring	498
Sitraka JMonitor Performance Monitoring	501
Chapter 30: J2EE Performance Monitoring	521
About J2EE Performance Monitoring.....	522
Installing the J2EE Monitor on the Application Server	523
Initial J2EE Monitor Configuration Settings.....	525
Activating the J2EE Monitor on the Client Machine	526
Examples of Modifying Application Server Configurations	530
Troubleshooting the J2EE Monitor	538
Chapter 31: Application Deployment Solutions	539
About Application Deployment Solutions Monitoring	539
Configuring the Citrix MetaFrame Server Monitor	540
Chapter 32: Middleware Performance Monitoring	551
About Middleware Performance Monitoring.....	551
Configuring the IBM WebSphere MQ Monitor	552
Configuring the TUXEDO Monitor	564
Chapter 33: Troubleshooting Online Monitors	569
Troubleshooting Server Resource Monitors	569
Troubleshooting the Network Delay Monitor	572
Network Considerations.....	574

PART V: APPENDIXES

Appendix A: Interpreting LoadRunner Online Graphs	579
Online Monitoring Graphs	579
Appendix B: Performing Path Translation	583
Understanding Path Translation	583
Adding Entries to the Path Translation Table.....	585
Editing the Path Translation Table	586
Path Translation Examples.....	588

Appendix C: Working in Expert Mode	589
Entering Expert Mode	589
Options - Agent Settings	590
Options - General Settings	591
Options - Debug Information Settings	592
Options - Output Settings	594
Options - Monitor Settings	595
Load Generator Information - UNIX Environment Settings	596
Load Generator Information - Connection Log Settings	597
Appendix D: Troubleshooting the Controller	599
About Troubleshooting	599
LoadRunner Communications	600
Failure to Communicate with a Load Generator	601
Failure to Connect to the AUT Database	607
Failure to Access Files	607
Failed Vusers or Transactions.....	609
Increasing the Number of Vusers on a Windows Machine	613
Troubleshooting Firewalls	614
Appendix E: Working with Server Monitor Counters	623
Changing a Monitor's Default Counters	623
Useful Counters for Stress Testing	624
Appendix F: Configuring Multiple IP Addresses	627
About Multiple IP Addresses	628
Adding IP Addresses to a Load Generator	629
Using the IP Wizard	629
Configuring Multiple IP Addresses on UNIX.....	634
Updating the Routing Table.....	635
Enabling Multiple IP Addressing from the Controller.....	636
Appendix G: Controller Command Line Arguments	637
About Controller Command Line Arguments	637
Invoking the Controller from the Command Line	637
TestDirector Arguments	638
Run-Time Arguments	639
Appendix H: Working with Digital Certificates	641
Using Digital Certificates with Firewalls.....	641
Creating and Using Digital Certificates	642
Index	647

Welcome to LoadRunner

Welcome to LoadRunner, Mercury Interactive's tool for testing the performance of applications. LoadRunner stresses your entire application to isolate and identify potential client, network, and server bottlenecks.

LoadRunner enables you to test your system under controlled and peak load conditions. To generate load, LoadRunner runs thousands of Virtual Users that are distributed over a network. Using a minimum of hardware resources, these Virtual Users provide consistent, repeatable, and measurable load to exercise your application just as real users would. LoadRunner's in-depth reports and graphs provide the information that you need to evaluate the performance of your application.

Online Resources



LoadRunner includes the following online tools:

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

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 LoadRunner online books.

LoadRunner Function Reference gives you online access to all of LoadRunner'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 *LoadRunner Function Reference*.

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

LoadRunner Documentation Set

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

- install LoadRunner
- create Vuser scripts
- use the LoadRunner Controller
- use the LoadRunner Analysis

Using the LoadRunner Documentation Set

The LoadRunner documentation set consists of one installation guide, a Controller user's guide, an Analysis user's guide, and two guides for creating Virtual User scripts.

Installation Guide

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

- ▶ the LoadRunner Controller—on a Windows-based machine
- ▶ Virtual User components—for both Windows and UNIX platforms

Controller User's Guide

The LoadRunner documentation pack includes one Controller user's guide:

The *LoadRunner Controller User's Guide (Windows)* describes how to create and run LoadRunner scenarios using the LoadRunner Controller in a Windows environment. The Vusers can run on UNIX and Windows-based platforms. The Controller user's guide presents an overview of the LoadRunner testing process.

Analysis User's Guide

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

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

Guides for Creating Vuser Scripts

The LoadRunner documentation pack has two guides that describe how to create Vuser scripts:

- ▶ The *LoadRunner Creating Vuser Scripts Guide* describes how to create all types of Vuser scripts. When necessary, supplement this document with the online *LoadRunner Function Reference* and the following guide.

- The *WinRunner User's Guide* describes in detail how to use WinRunner to create GUI Vuser scripts. The resulting Vuser scripts run on Windows platforms. The *TSL Online Reference* should be used in conjunction with this document.

For information on	Look here...
Installing LoadRunner	<i>LoadRunner Installation Guide</i>
The LoadRunner testing process	<i>LoadRunner Controller User's Guide (Windows)</i>
Creating Vuser scripts	<i>LoadRunner Creating Vuser Scripts Guide</i>
Creating and running scenarios	<i>LoadRunner Controller User's Guide (Windows)</i>
Analyzing test results	<i>LoadRunner Analysis User's Guide</i>

Typographical Conventions

This book uses the following typographical conventions:

- 1, 2, 3** Bold numbers indicate steps in a procedure.
- Bullets indicate options and features.
- > The greater than sign separates menu levels (for example, **File > Open**).
- Stone Sans** The **Stone Sans** font indicates names of interface elements on which you perform actions (for example, "Click the **Run** button.>").
- Bold** **Bold** text indicates method or function names
- Italics* *Italic* 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, *<Product installation folder>\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 LoadRunner

1

Introduction

To load test your application, LoadRunner emulates an environment where multiple users work concurrently. While the application is under load, LoadRunner accurately measures, monitors, and analyzes a system's performance and functionality.

Application Load Testing

Modern system architectures are complex. While they provide an unprecedented degree of power and flexibility, these systems are difficult to test. Whereas single-user testing focuses primarily on functionality and the user interface of a system component, application testing focuses on performance and reliability of an entire system.

For example, a typical application testing scenario might depict 1000 users that log in simultaneously to a system on Monday morning: What is the response time of the system? Does the system crash? To be able to answer these questions—and more—a complete application performance testing solution must:

- ▶ test a system that combines a variety of software applications and hardware platforms
- ▶ determine the suitability of a server for any given application
- ▶ test the server before the necessary client software has been developed
- ▶ emulate an environment where multiple clients interact with a single server application
- ▶ test an application under the load of tens, hundreds, or even thousands of potential users

Manual Testing Limitations

Traditional or manual testing methods offer only a partial solution to load testing. For example, you can test an entire system manually by constructing an environment where many users work simultaneously on the system. Each user works at a single machine and submits input to the system. However, this manual testing method has the following drawbacks:

- it is expensive, requiring large amounts of both personnel and machinery
- it is complicated, especially coordinating and synchronizing multiple testers
- it involves a high degree of organization, especially to record and analyze results meaningfully
- the repeatability of the manual tests is limited

The LoadRunner Solution

The LoadRunner automated solution addresses the drawbacks of manual performance testing:

- LoadRunner reduces the personnel requirements by replacing human users with virtual users or *Vusers*. These *Vusers* emulate the behavior of real users—operating real applications.
- Because numerous *Vusers* can run on a single computer, LoadRunner reduces the hardware requirements.
- The LoadRunner Controller allows you to easily and effectively control all the *Vusers*—from a single point of control.
- LoadRunner monitors the application performance online, enabling you to fine-tune your system during test execution.
- LoadRunner automatically records the performance of the application during a test. You can choose from a wide variety of graphs and reports to view the performance data.
- LoadRunner checks where performance delays occur: network or client delays, CPU performance, I/O delays, database locking, or other issues at the database server. LoadRunner monitors the network and server resources to help you improve performance.

- Because LoadRunner tests are fully automated, you can easily repeat them as often as you need.

Using LoadRunner

<i>Scenarios</i>	Using LoadRunner, you divide your application performance testing requirements into <i>scenarios</i> . A scenario defines the events that occur during each testing session. Thus, for example, a scenario defines and controls the number of users to emulate, the actions that they perform, and the machines on which they run their emulations.
<i>Vusers</i>	In the scenario, LoadRunner replaces human users with <i>virtual users</i> or <i>Vusers</i> . When you run a scenario, Vusers emulate the actions of human users working with your application. While a workstation accommodates only a single human user, many Vusers can run concurrently on a single workstation. In fact, a scenario can contain tens, hundreds, or even thousands of Vusers.
<i>Vuser Scripts</i>	The actions that a Vuser performs during the scenario are described in a Vuser script. When you run a scenario, each Vuser executes a <i>Vuser script</i> . The Vuser scripts include functions that measure and record the performance of your application's components.
<i>Transactions</i>	To measure the performance of the server, you define <i>transactions</i> . A transaction represents an action or a set of actions that you are interested in measuring. You define transactions within your Vuser script by enclosing the appropriate sections of the script with <i>start</i> and <i>end</i> transaction statements. For example, you can define a transaction that measures the time it takes for the server to process a request to view the balance of an account and for the information to be displayed at the ATM.
<i>Rendezvous points</i>	You insert <i>rendezvous points</i> into Vuser scripts to emulate heavy user load on the server. <i>Rendezvous points</i> instruct Vusers to wait during test execution for multiple Vusers to arrive at a certain point, in order that they may simultaneously perform a task. For example, to emulate peak load on the bank server, you can insert a rendezvous point instructing 100 Vusers to deposit cash into their accounts at the same time.

Controller You use the *LoadRunner Controller* to manage and maintain your scenarios. Using the Controller, you control all the Vusers in a scenario from a single workstation.

Load generator When you execute a scenario, the LoadRunner Controller distributes each Vuser in the scenario to a *load generator*. The load generator is the machine that executes the Vuser script, enabling the Vuser to emulate the actions of a human user.

Performance analysis Vuser scripts include functions that measure and record system performance during load-testing sessions. During a scenario run, you can monitor the network and server resources. Following a scenario run, you can view performance analysis data in reports and graphs.

Working with LoadRunner

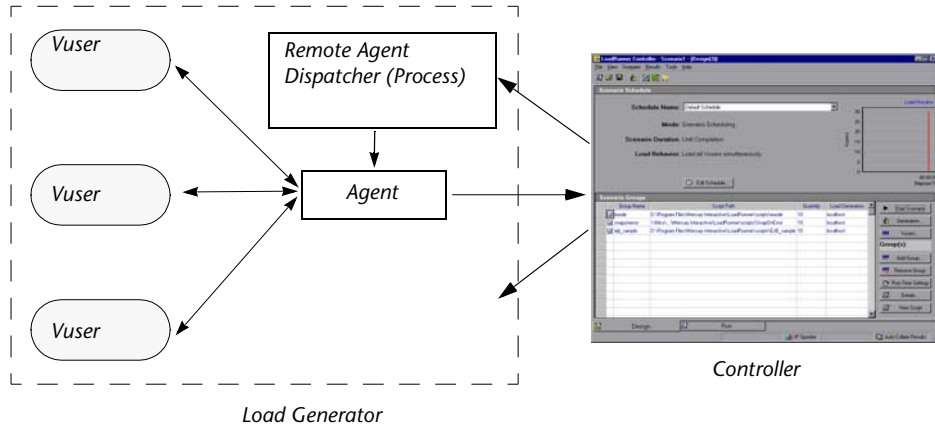
Suppose you want to test an online banking Web server that is accessed by many Internet users. The Web site provides a full range of banking services to the customers—such as the ability to transfer funds and check account balances. To test this server, you create a scenario. The scenario defines the actions that are performed on the server during the load test.

During the scenario that loads and monitors the bank server, you want to:

- ▶ emulate conditions of controlled load on the server
- ▶ emulate conditions of maximum load on the server
- ▶ measure server performance under load
- ▶ check where performance delays occur: network or client delays, CPU performance, I/O delays, database locking, or other issues at the server
- ▶ monitor the network and server resources under load

LoadRunner Vuser Technology

On each Windows load generator, you install the *Remote Agent Dispatcher (Process)* and a *LoadRunner Agent*.



Remote Agent Dispatcher (Process)

The Remote Agent Dispatcher (Process) enables the Controller to start applications on the load generator machine.

Agent

The LoadRunner Agent enables the Controller and the load generator to communicate with each other. When you run a scenario, the Controller instructs the Remote Agent Dispatcher (Process) to launch the LoadRunner agent. The agent receives instructions from the Controller to initialize, run, pause, and stop Vusers. At the same time, the agent also relays data on the status of the Vusers back to the Controller.

LoadRunner Vuser Types

LoadRunner has various types of Vusers. Each type is designed to handle different aspects of today's system architectures. You can use the Vuser types in any combination in a scenario in order to create a comprehensive application test. The following Vuser types are available:

► **Client/Server**

For MSSQLServer, ODBC, Oracle (2-tier), DB2 CLI, Sybase Ctlib, Sybase Dblib, Windows Sockets and DNS protocols.

► **Custom**

For C templates, Visual Basic templates, Java templates, Javascript and VBScript type scripts.

► **Distributed Components**

For COM/DCOM, Corba-Java, and Rmi-Java protocols.

► **E-business**

For FTP, LDAP, Media Player, Multi Protocol Web/WS, Web (HTTP, HTML), Palm, and RealPlayer protocols.

► **Enterprise Java Beans**

For EJB Testing and Rmi-Java protocols.

► **ERP**

For Oracle NCA, Peoplesoft (Tuxedo), SAP, and Siebel protocols.

► **Legacy**

For Terminal Emulation (RTE).

► **Mailing Services**

Internet Messaging (IMAP), MS Exchange (MAPI), POP3, and SMTP.

► **Middleware**

For the Tuxedo (6, 7) protocol.

► **Wireless**


For i-Mode, VoiceXML, and WAP protocols.

GUI Vusers

GUI Vusers

GUI Vusers operate graphical user interface (GUI) applications. These applications can run in a Microsoft Windows environment. Each GUI Vuser that you develop emulates a real user by submitting input to, and receiving output from, GUI applications. For example, a GUI Vuser could operate Microsoft Paint as follows:

1. Select Open from the File menu.
2. Select a graphic file called test.bmp.
3. Click the Open button.
4. Select Flip/Rotate from the Image menu.
5. Click the Flip Horizontal radio button.
6. Click the OK button.
7. Select Save from the File menu.




The operations that a GUI Vuser performs on an application are defined in a GUI Vuser script. You create GUI Vuser scripts using Mercury Interactive's GUI testing tools: WinRunner (for Microsoft Windows applications), and Astra QuickTest (for Web applications).

You can run only a single GUI Vuser on a Windows-based load generator. Use Citrix to run multiple GUI Vusers. Refer to the Readme file for additional information about configuring your load generators using Citrix. For additional information on Windows-based GUI Vusers, refer to the *LoadRunner Creating Vuser Scripts User's Guide*.

Note: You can run GUI and SAP Vusers on remote load generators only if you install the Remote Agent Dispatcher as a process. If you install the Remote Agent Dispatcher as a service, you cannot run GUI Vusers on remote load generators.

Vuser Technology

Vusers (except for GUI and RTE Vusers) generate load on a server by submitting input directly to the server. Vusers do not operate client applications—they access the server using LoadRunner API functions. These API functions emulate the input from an actual application.



Because Vusers are not reliant on client software, you can use Vusers to test server performance even before the client software has been developed. Further, since Vusers do not have a user interface, the amount of system resources required is minimal. This allows you to run large numbers of Vusers on a single workstation.

The following example illustrates the use of Vusers: Suppose that you have a Web-based database server that maintains your customer information. The information is accessed by numerous customer service personnel who are located throughout the country. The server receives the queries, processes the requests, and returns responses, via the Web, to field personnel.

You want to test the response times of the entire system when numerous service personnel simultaneously access the server. Using LoadRunner, you could create several hundred Vusers, each Vuser accessing the server database. The Vusers enable you to emulate and measure the performance of your database and Web servers under the load of many users.

You develop a Vuser script to define the actions of a Vuser. A Vuser script includes functions that control the script execution, specify the input that the Vuser submits to the server, and measure the server performance.

You develop Vuser scripts either by recording with LoadRunner's Vuser Script Generator (VuGen) or by using LoadRunner's Vuser script templates.

For the database server example above, you could create a Vuser script that performs the following actions:

- logs in to the Web application
- connects to the database server
- submits an SQL query
- retrieves and processes the server response
- disconnects from the server and the Web

You can create Vuser scripts on a Windows-based platform, or program them on a UNIX platform. For a list of the supported UNIX platforms, see the LoadRunner Readme file. For more information about Vusers, refer to the *LoadRunner Creating Vuser Scripts User's Guide*.

RTE Vusers

RTE Vusers

RTE Vusers operate character-based applications. Each RTE Vuser that you develop emulates a real user by submitting input to, and receiving output from, character-based applications.

```

View - Vuser RTE/0
=====
operations:
  1) Withdraw Cash.
  2) Deposit Cash.
  3) Balance Report.
  4) Exit ATM.

Please select (1-4): 2
Enter amount of money to deposit: 168
Depositing $168 in process, Please wait...
Operation has been successfully completed.

ATM Services
=====
operations:
  1) Withdraw Cash.
  2) Deposit Cash.
  3) Balance Report.
  4) Exit ATM.


Please select (1-4): 

```

The following example illustrates the use of RTE Vusers: Suppose that you have a database server that maintains customer information. The information is accessed by numerous field service representatives who are located throughout the country. Every time a field service representative makes a repair, he accesses the server database by modem. Using a character-based application, the service representative records the customer complaint and accesses additional information about the customer.

You want to test the response times of the server when many service personnel simultaneously access the server. Using LoadRunner, you could create several hundred RTE Vusers, each Vuser accessing the server database using a character-based application. The RTE Vusers enable you to emulate and measure the performance of your server under the load of many users.

The operations that an RTE Vuser performs on an application are defined in an RTE Vuser script. You create RTE Vuser scripts by using the Vuser Script Generator (VuGen). The generator enables you to record the actions that you perform on a character-based application.

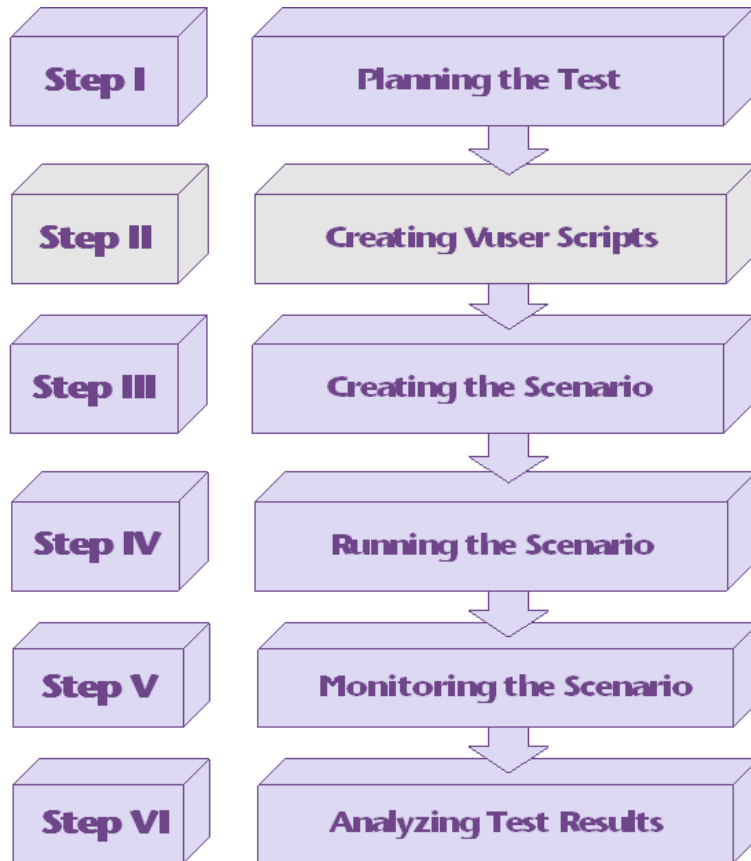


For further information on RTE Vusers, refer to the *LoadRunner Creating Vuser Scripts User's Guide*.

2

The LoadRunner Testing Process

You can easily create and run load test scenarios by following the LoadRunner testing process below. The following illustration outlines the testing process:



This chapter gives you an overview of LoadRunner's six-step process for testing your Web-based application under load.

Step I: Planning the Test

Successful load testing requires that you develop a thorough test plan. A clearly defined test plan will ensure that the LoadRunner scenarios that you develop will accomplish your load testing objectives. For more information, see Chapter 3, "Load Test Planning."

Step II: Creating the Vuser Scripts

Vusers emulate human users interacting with your Web-based application. A Vuser script contains the actions that each virtual user performs during scenario execution.

In each Vuser script you determine the tasks that will be:

- performed by each Vuser
- performed simultaneously by multiple Vusers
- measured as transactions

For more information on creating Vuser scripts, refer to the *LoadRunner Creating Vuser Scripts User's Guide*.

Step III: Creating the Scenario

A scenario describes the events that occur during a testing session. A scenario includes a list of machines on which Vusers run, a list of scripts that the Vusers run, and a specified number of Vusers or Vuser groups that run during the scenario. You create scenarios using the LoadRunner Controller. For an introduction to the Controller, see Chapter 4, "The LoadRunner Controller at a Glance."

Creating a Manual Scenario

You create a scenario by defining Vuser groups to which you assign a quantity of individual Vusers, Vuser scripts, and load generators to run the scripts. For instructions on creating a manual scenario, see Chapter 5, “Creating a Manual Scenario.”

You can also create a scenario using the Percentage Mode, in which you define the total number of Vusers to be used in the scenario, and the load generator machines and percentage of the total number of Vusers to be assigned to each Vuser script. For instructions on creating a manual scenario in Percentage Mode, see Chapter 6, “Creating a Manual Scenario Using the Percentage Mode.”

Creating a Goal-Oriented Scenario

For Web tests, you can create a goal-oriented scenario, in which you define the goals you want your test to achieve. LoadRunner automatically builds a scenario for you, based on these goals. For instructions on creating a goal-oriented scenario, see Chapter 7, “Creating a Goal-Oriented Scenario.”

Step IV: Running the Scenario

You emulate user load on the server by instructing multiple Vusers to perform tasks simultaneously. You can set the level of load by increasing and decreasing the number of Vusers that perform tasks at the same time. For more information, see Chapter 9, “Using Rendezvous Points.”

Before you run a scenario, you set the scenario configuration and scheduling. This determines how all the load generators and Vusers behave when you run the scenario. For more information, see Chapter 10, “Configuring a Scenario” and Chapter 8, “Scheduling a Scenario.”

You can run the entire scenario, groups of Vusers (Vuser groups), or individual Vusers. While a scenario runs, LoadRunner measures and records the transactions that you defined in each Vuser script. You can also monitor your system’s performance online. For more information, see Part III, “Executing a Scenario.”

Step V: Monitoring a Scenario

You can monitor scenario execution using the LoadRunner online run-time, transaction, system resource, Web resource, Web server resource, Web application server resource, database server resource, network delay, streaming media resource, firewall server resource, ERP server resource, and Java performance monitors. For more information, see Part IV, “Monitoring a Scenario.”

Step VI: Analyzing Test Results

During scenario execution, LoadRunner records the performance of the application under different loads. You use LoadRunner's graphs and reports to analyze the application's performance. For more information about LoadRunner's reports and graphs, see the *LoadRunner Analysis User's Guide*.

3

Load Test Planning

Developing a comprehensive test plan is a key to successful load testing. A clearly defined test plan ensures that the LoadRunner scenarios you develop will accomplish your load testing objectives.

This chapter introduces the load test planning process:

- ▶ Analyzing the Application
- ▶ Defining Testing Objectives
- ▶ Planning LoadRunner Implementation
- ▶ Examining Load Testing Objectives

About Load Test Planning

As in any type of system testing, a well-defined test plan is the first essential step to successful testing. Planning your load testing helps you to:

- ▶ Build test scenarios that accurately emulate your working environment.

Load testing means testing your application under typical working conditions, and checking for system performance, reliability, capacity, etc.

- ▶ Understand which resources are required for testing.

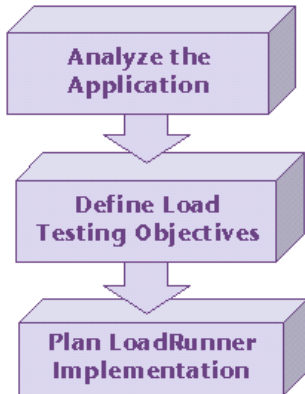
Application testing requires hardware, software, and human resources. Before you begin testing, you should know which resources are available and decide how to use them effectively.

- ▶ Define success criteria in measurable terms.

Focused testing goals and test criteria ensure successful testing. For example, it's not enough to define vague objectives like "Check server response time

under heavy load.” A more focused success criteria would be “Check that 50 customers can check their account balance simultaneously, and that the server response time will not exceed one minute.”

Load test planning is a three-step process:



Analyzing the Application

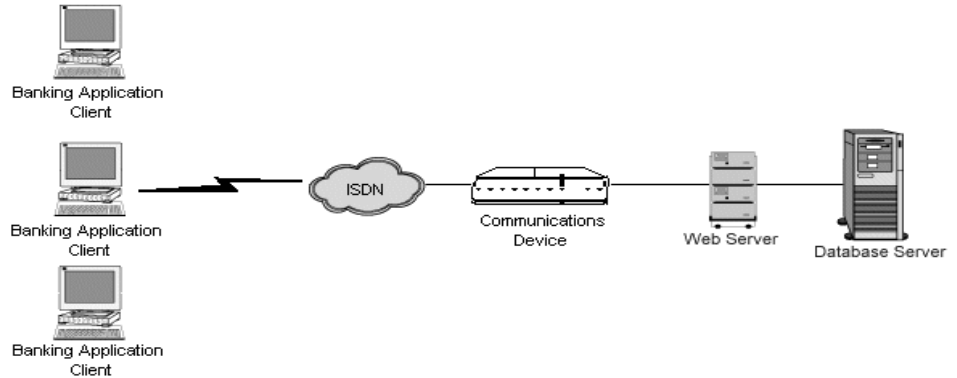
The first step to load test planning is analyzing your application. You should become thoroughly familiar with the hardware and software components, the system configuration, and the typical usage model. This analysis ensures that the testing environment you create using LoadRunner will accurately reflect the environment and configuration of the application under test.

Identifying System Components

Draw a schematic diagram to illustrate the structure of the application. If possible, extract a schematic diagram from existing documentation. If the application under test is part of a larger network system, you should identify the component of the system to be tested. Make sure the diagram includes all system components, such as client machines, network, middleware, and servers.

The following diagram illustrates an online banking system which is accessed by many Web users. The Web users each connect to the same

database to transfer funds and check balances. The customers connect to the database server through the Web, using multiple browsers.



Describing the System Configuration

Enhance the schematic diagram with more specific details. Describe the configuration of each system component. You should be able to answer the following questions:

- How many users are anticipated to connect to the system?
- What is the application client's machine configuration? (hardware, memory, operating system, software, development tool, etc.)
- What types of database and Web servers are used? (hardware, database type, operating system, file server, etc.)
- How does the server communicate with the application client?
- What is the middleware configuration and application server between the front-end client and back-end server?
- What other network components may affect response time? (modems etc.)
- What is the throughput of the communications devices? How many concurrent users can each device handle?

For example, the schematic diagram above specified that there are multiple application clients accessing the system.

Front-End Client Configuration	
Anticipated number of application clients	50 concurrent application clients
Hardware / Memory	586 / 32MB
Operating system & version	Windows NT 4.0
Client browser	Internet Explorer 4.0

Analyzing the Usage Model

Define how the system is typically used, and decide which functions are important to test. Consider who uses the system, the number of each type of user, and each user's common tasks. In addition, consider any background load that might affect the system response time.

For example, suppose 200 employees log on to the accounting system every morning, and the same office network has a constant background load of 50 users performing various word processing and printing tasks. You could create a LoadRunner scenario with 200 virtual users signing in to the accounting database, and check the server response time.

To check how background load affects the response time, you could run your scenario on a network where you also simulate the load of employees performing word processing and printing activities.

Task Distribution

In addition to defining the common user tasks, examine the distribution of these tasks. For example, suppose the bank uses a central database to serve clients across many states and time zones. The 250 application clients are located in two different time zones, all connecting to the same Web server. There are 150 in Chicago and 100 in Detroit. Each begins their business day at 9:00 AM, but since they are in different time zones, there should never be more than 150 users signing in at any given time.

You can analyze task distribution to determine when there is peak database activity, and which activities typically occur during *peak load* time.

Defining Testing Objectives

Before you begin testing, you should define exactly what you want to accomplish.

Following are common application testing objectives that LoadRunner helps you test, as described in Robert W. Buchanan, Jr's *The Art of Testing Network Systems* (John Wiley & Sons, Inc., 1996).

Objective	Answers the Question
Measuring end-user response time	How long does it take to complete a business process?
Defining optimal hardware configuration	Which hardware configuration provides the best performance?
Checking reliability	How hard or long can the system work without errors or failures?
Checking hardware or software upgrades	How does the upgrade affect performance or reliability?
Evaluating new products	Which server hardware or software should you choose?
Measuring system capacity	How much load can the system handle without significant performance degradation?
Identifying bottlenecks	Which element is slowing down response time?

A more detailed description of each objective appears at the end of this chapter.

Stating Objectives in Measurable Terms

Once you decide on your general load testing objectives, you should provide more focused goals by stating your objectives in measurable terms. To provide a baseline for evaluation, determine exactly what constitutes acceptable and unacceptable test results.

For example:

General Objective - Product Evaluation: choose hardware for the Web server.

Focused Objective - Product Evaluation: run the same group of 300 virtual users on two different servers, HP and NEC. When all 300 users simultaneously browse the pages of your Web application, determine which hardware gives a better response time.

Deciding When to Test

Load testing is necessary throughout the product life cycle. The following table illustrates what types of tests are relevant for each phase of the product life cycle:

Planning and Design	Development	Deployment	Production	Evolution
Evaluate new products	Measure response time	Check reliability	Measure response time	Check HW or SW upgrades
Measure response time	Check optimal hardware configuration	Measure response time	Identify bottlenecks	Measure system capacity
	Check HW or SW upgrades	Measure system capacity		
	Check reliability			

Planning LoadRunner Implementation

The next step is to decide how to use LoadRunner to achieve your testing goals.


Defining the Scope of Performance Measurements

You can use LoadRunner to measure response time at different points in the application. Determine where to run the Vusers and which Vusers to run according to the test objectives:

► Measuring end-to-end response time:


You can measure the response time that a typical user experiences by running a GUI Vuser or RTE Vuser at the front end. GUI Vusers emulate real users by submitting input to and receiving output from the client application; RTE Vusers emulate real users submitting input to and receiving output from a character-based application.

You can run GUI or RTE Vusers at the front end to measure the response time across the entire network, including a terminal emulator or GUI front end, network, and server.



► Measuring network and server response times:


You can measure network and server response time, excluding response time of the GUI front end, by running Vusers (not GUI or RTE) on the client machine. Vusers emulate client calls to the server without the user interface. When you run many Vusers from the client machine, you can measure how the load affects network and server response time.



► Measuring GUI response time:


You can determine how the client application interface affects response time by subtracting the previous two measurements:

GUI response time = end-to-end - network and server




► Measuring server response time:

You can measure the time it takes for the server to respond to a request without going across the network. When you run Vusers on a machine directly connected to the server, you can measure server performance.



► Measuring middleware-to-server response time:

You can measure response time from the server to middleware if you have access to the middleware and its API. You can create Vusers with the middleware API and measure the middleware-server performance.



Defining Vuser Activities

Create Vuser scripts based on your analysis of Vuser types, their typical tasks and your test objectives. Since Vusers emulate the actions of a typical end-user, the Vuser scripts should include the typical end-user tasks. For example, to emulate an online banking client, you should create a Vuser script that performs typical banking tasks. You would browse the pages that you normally visit to transfer funds or check balances.

You decide which tasks to measure based on your test objectives and define *transactions* for these tasks. Transactions measure the time that it takes for the server to respond to tasks submitted by Vusers (end-to-end time). For example, to check the response time of a bank Web server supplying an account balance, define a transaction for this task in the Vuser script.

In addition, you can emulate peak activity by using *rendezvous points* in your script. Rendezvous points instruct multiple Vusers to perform tasks at exactly the same time. For example, you can define a rendezvous to emulate 70 users simultaneously updating account information.

Selecting Vusers

Before you decide on the hardware configuration to use for testing, determine the number and type of Vusers required. To decide how many Vusers and which types to run, look at the typical usage model, combined with the testing objectives. Some general guidelines are:

- ▶ Use one or a few GUI users to emulate each type of typical user connection.
- ▶ Use RTE Vusers to emulate terminal users.
- ▶ Run multiple non-GUI or non-RTE Vusers to generate the rest of the load for each user type.

For example, suppose that you have five kinds of users, each performing a different business process:

Usage Model	GUI	RTE	Other
100 customer service users in New York (LAN connection)	2	–	98
30 customers in Europe (dial-in ISDN connection)	2	–	28
5 background batch processes	–	–	5
150 customers (terminal connection)	–	150	–
6 managers (2 users with 486 PCs, 4 with 586 PCs)	1 (486 PC) 1 (586 PC)	–	4

Choosing Testing Hardware/Software

The hardware and software should be powerful and fast enough to emulate the required number of virtual users.

To decide on the number of machines and correct configuration, consider the following:

- It is advisable to run the LoadRunner Controller on a separate machine.
- Each GUI Vuser requires a separate Windows-based machine; several GUI Vusers can run on a single UNIX machine.
- Configuration of the test machine for GUI Vusers should be as similar as possible to the actual user's machine.

Refer to the following tables to estimate the required hardware for each LoadRunner testing component. These requirements are for optimal performance.

Windows Configuration Requirements

Requirement	Controller with Online Monitors	Virtual Vuser Generator	Virtual Users	Analysis Module
Computer/Processor	Pentium 350 MHz or higher	Pentium 350 MHz or higher	Pentium 1 GHz or higher	Pentium 350 MHz or higher
Operating System	Windows NT® service pack 6a or later Windows 2000 Windows XP	Windows NT® service pack 6a or later Windows 2000 Windows XP	Windows NT® service pack 6a or later Windows 2000 Windows XP HP UX 11.x or higher, Solaris 2.6 or higher, AIX 4.3.3 or higher, Linux Red Hat 6.0 or higher	Windows NT® service pack 6a or later Windows 2000 Windows XP
Memory	128 MB or more	128 MB or more	At least 1 MB RAM for non-multi-threaded Vuser or at least 512 KB multi-threaded Vuser	128 MB or more
Swap Space	Twice the total physical memory	Twice the total physical memory	Twice the total physical memory	Twice the total physical memory
Free Hard Disk Space	200 MB	200 MB	Minimum 500 MB	Minimum 500 MB
Browser	IE 5.x or higher Netscape Navigator 4.x, 6.x	IE 5.x or higher Netscape Navigator 4.x, 6.x	N/A	IE 5.x or higher Netscape Navigator 4.x, 6.x

Note: The results file requires a few MB of disk space for a long scenario run with many transactions. The load generator machines also require a few MB of disk space for temporary files if there is no NFS. See Chapter 10, “Configuring a Scenario” for more information about run-time file storage.

Note: Refer to <http://www.mercuryinteractive.com/products/loadrunner/technical/> for the most updated installation requirements.

UNIX Configuration Requirements

Requirement	GUI Vuser (per user)	Vuser (per user)	Web Vuser (per user)
Memory	4-5 MB plus client application requirements	At least 1.5 MB (depends on application)	~0.5 MB
Swap Space	Four times the total physical memory	Four times the total physical memory	Two times the total physical memory
Disk Space	n/a	n/a	n/a
No. of Processes	4	1	1
No. of pty's	n/a	n/a	n/a
1 CPU supports x users	30-50 or more	200-300 or more	300-400 or more

Note: The results file requires a few MB of disk space for a long scenario run with many transactions. The load generator machines also require a few MB of disk space for temporary files if there is no NFS. Refer to Chapter 10, "Configuring a Scenario" for more information about run-time file storage.


Examining Load Testing Objectives

Your test plan should be based on a clearly defined testing objective. This section presents an overview of common testing objectives:

- Measuring end-user response time
- Defining optimal hardware configuration
- Checking reliability
- Checking hardware or software upgrades
- Evaluating new products
- Identifying bottlenecks
- Measuring system capacity

Measuring End-user Response Time

Check how long it takes for the user to perform a business process and receive a response from the server. For example, suppose that you want to verify that while your system operates under normal load conditions, the end users receive responses to all requests within 20 seconds. The following graph presents a sample load vs. response time measurement for a banking application:



Defining Optimal Hardware Configuration

Check how various system configurations (memory, CPU speed, cache, adaptors, modems) affect performance. Once you understand the system architecture and have tested the application response time, you can measure the application response for different system configurations to determine which settings provide the desired performance levels.

For example, you could set up three different server configurations and run the same tests on each configuration to measure performance variations:

- Configuration 1: 200MHz, 64MB RAM
- Configuration 2: 200MHz, 128MB RAM
- Configuration 3: 266MHz, 128MB RAM

Checking Reliability

Determine the level of system stability under heavy or continuous work loads. You can use LoadRunner to create stress on the system: force the system to handle extended activity in a compressed time period to simulate the kind of activity a system would normally experience over a period of weeks or months.

Checking Hardware or Software Upgrades

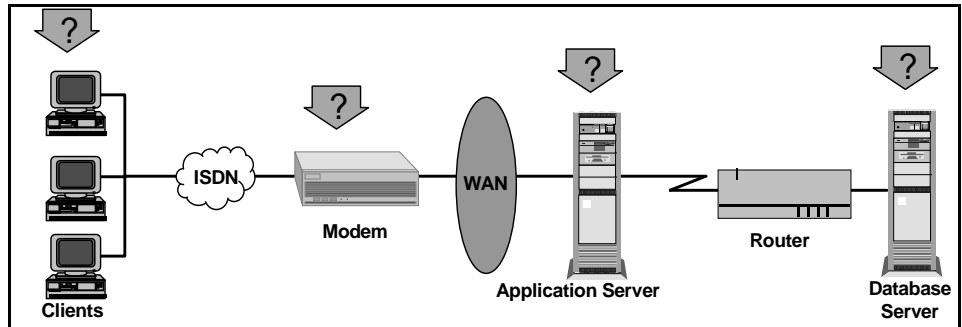
Perform regression testing to compare a new release of hardware or software to an older release. You can check how an upgrade affects response time (benchmark) and reliability. Application regression testing does not check new features of an upgrade; rather it checks that the new release is as efficient and reliable as the older release.

Evaluating New Products

You can run tests to evaluate individual products and subsystems during the planning and design stage of a product's life cycle. For example, you can choose the hardware for the server machine or the database package based on evaluation tests.


Identifying Bottlenecks

You can run tests which identify bottlenecks on the system to determine which element is causing performance degradation, for example, file locking, resource contention and network overload. Use LoadRunner in conjunction with the new network and machine monitoring tools to create load and measure performance at different points in the system. For more information, see Part IV, “Monitoring a Scenario.”.



Measuring System Capacity

Measure how much excess capacity the system can handle without performance degradation. To check capacity, you can compare performance versus load on the existing system, and determine where significant response-time degradation begins to occur. This is often called the “knee” of the response time curve.



Once you determine the current capacity, you can decide if resources need to be increased to support additional users.

4

The LoadRunner Controller at a Glance

This chapter introduces the Controller window and explains how to perform basic scenario operations.

This chapter describes:

- Opening the Controller
- Introducing the LoadRunner Controller
- Managing Scenario Files
- Running a Scenario

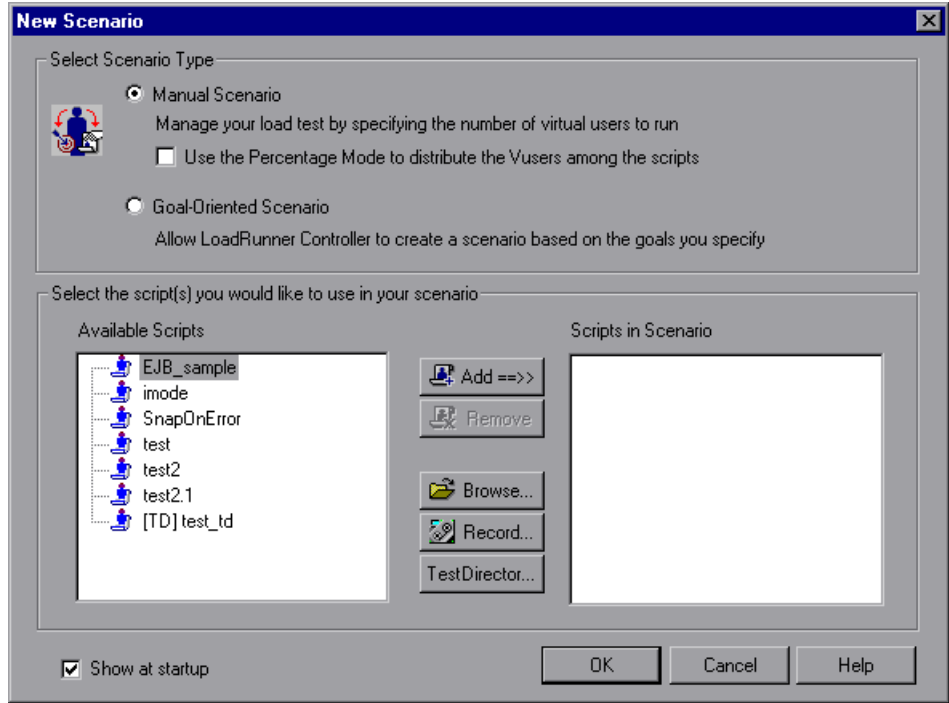
Opening the Controller

Set up the LoadRunner environment according to the instructions in the *LoadRunner Installation Guide*.

To open the Controller:



Select **Start > Programs > LoadRunner > Controller**. The Controller opens with the New Scenario dialog box inside.



You can select one of two methods to create a scenario: **Manual Scenario** or **Goal-Oriented Scenario**. In a manual scenario, you create the scenario yourself by defining the number of Vuser groups you want to run, and building a schedule for LoadRunner to run these groups. You can also create a manual scenario by defining the total number of Vusers to be used in the scenario, and assigning a percentage of the total number of Vusers to each script. If you want to create a scenario using the Percentage Mode, select **Use the Percentage Mode to distribute the Vusers among the scripts**.

In a goal-oriented scenario, you define the goals you want your test to achieve, and LoadRunner automatically builds a scenario for you, based on these goals.

For instructions on creating a manual scenario, see Chapter 5, “Creating a Manual Scenario.” For instructions on creating a manual scenario using the Percentage Mode, see Chapter 6, “Creating a Manual Scenario Using the Percentage Mode.”

For instructions on creating a goal-oriented scenario, see Chapter 7, “Creating a Goal-Oriented Scenario.”

To select the script or scripts that you want to use in your scenario:


- 1 Select a script from the Available Scripts list. By default, the list displays the fifty most recently used scripts.

Note: You can change the maximum number of scripts displayed in the Available Scripts list by modifying the following registry key:
 HKEY_CURRENT_USER\Software\Mercury Interactive\RecentScripts\
 max_num_of_scripts

You can also click the **Browse** button to locate the script you want to use. To view the directory path of a script listed in the Available Scripts list, right-click the script and select **Show Paths**.

To select a script saved in the TestDirector database, click the **TestDirector** button. To record a new script using VuGen, click **Record**.

Note: To select a VB Vuser script, browse to locate the *.usr* file.



- 2 Click the **Add** button to copy the script you selected to the Scripts in Scenario list.

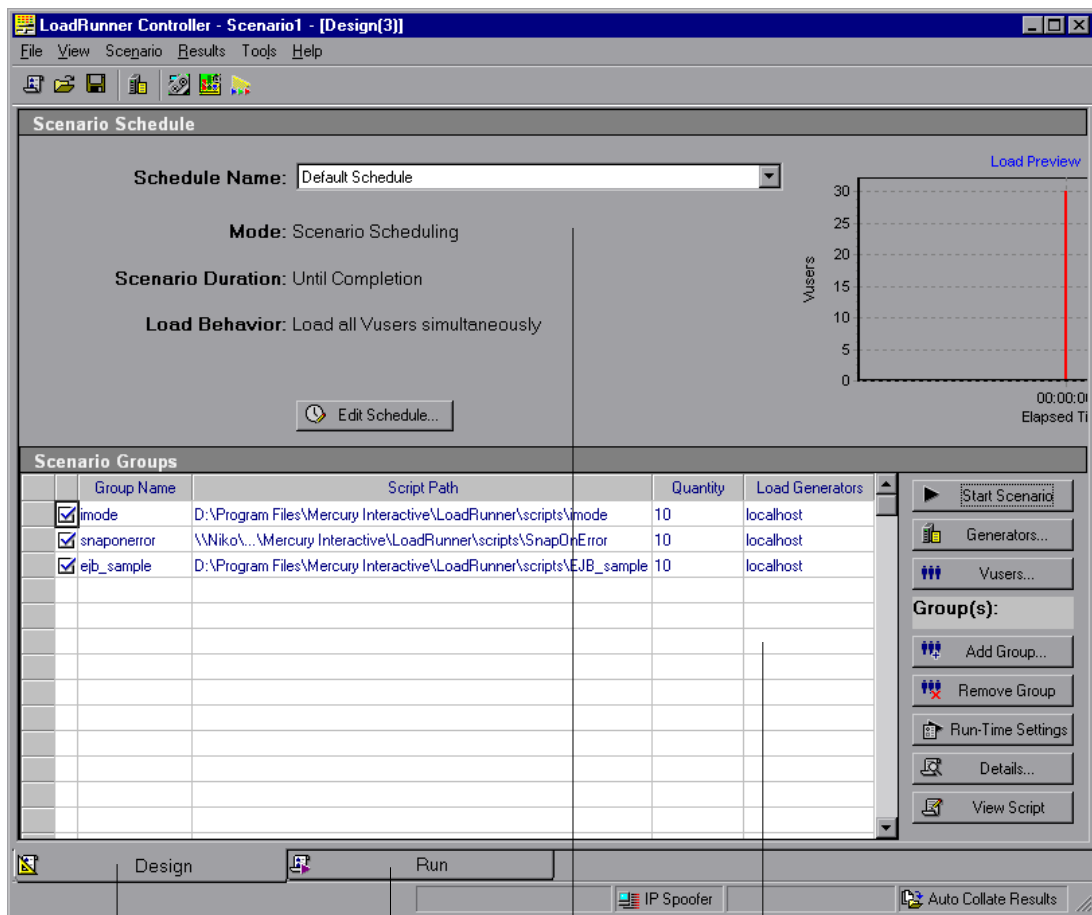


- 3 Click the **Remove** button to remove a script from the Scripts in Scenario list.
- 4 To bypass this dialog box the next time you create a new scenario, clear the **Show at startup** check box. You will be able to add scripts later on, while building your scenario.
- 5 Click **OK** to close the dialog box.

Introducing the LoadRunner Controller

The LoadRunner Controller window contains the following elements:

Title bar	Displays the name of the scenario on which you are currently working.
Menu bar	Displays the menus from which you select commands.
Toolbar	Provides shortcuts for selecting commands. Clicking a button executes a command.
Status Bar	Displays tool tips for the Controller menu items, as well as the following, if they are enabled: TestDirector Connection, IP Spoofer, Auto Collate Results, Auto Load Analysis, and WAN Emulator.



Design tab

Run tab

Scenario Groups pane
(Manual Scenario)

Scenario Schedule pane
(Manual Scenario)

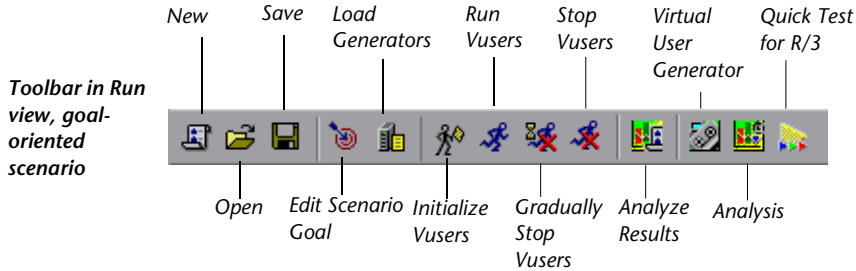
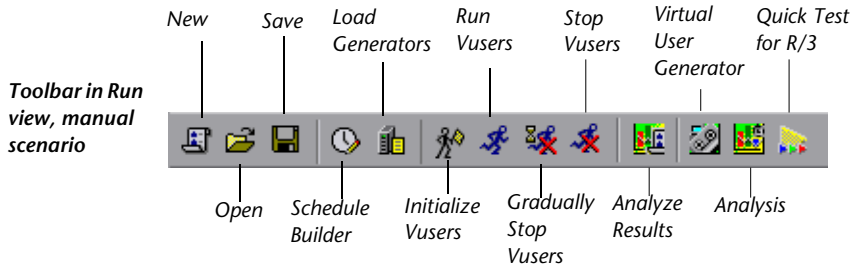

The Controller window has two tabs which correspond to two views:

- | | |
|--------------------|---|
| Design view | This view displays a list of all the Vuser groups/scripts in a scenario, and the load generator machines, and number of Vusers assigned to each group/script. This view also displays basic information about the scenario schedule (manual scenario) or goal (goal-oriented scenario). |
| Run view | Displays information on the running Vusers and Vuser groups, as well as online monitor graphs. |

In addition, if you select **View > Show Output**, the Controller opens the **Output** window which displays error, warning, notification, debug, and batch messages generated during scenario execution.

Choosing Commands from the Toolbar

You can execute many LoadRunner commands by clicking a button on the toolbar in the LoadRunner Controller. There are some variations in the buttons the toolbar displays, depending on whether you are in Design view or Run view, and depending on whether you are creating a manual scenario or a goal-oriented scenario.



Managing Scenario Files

A scenario describes the events that occur during each load testing session. You create a scenario using the Design view of LoadRunner Controller.

After you create the scenario, LoadRunner saves the information in a scenario file (.lrs). You use the commands in the File menu to create, open, save, and close scenario files. Some of these commands are available from the toolbar.

Creating a New Scenario

The New command creates a completely new scenario. Note that the New command clears all the information displayed in the Controller windows.



To create a new scenario, choose **File > New**, or click the **New** button on the Controller toolbar.

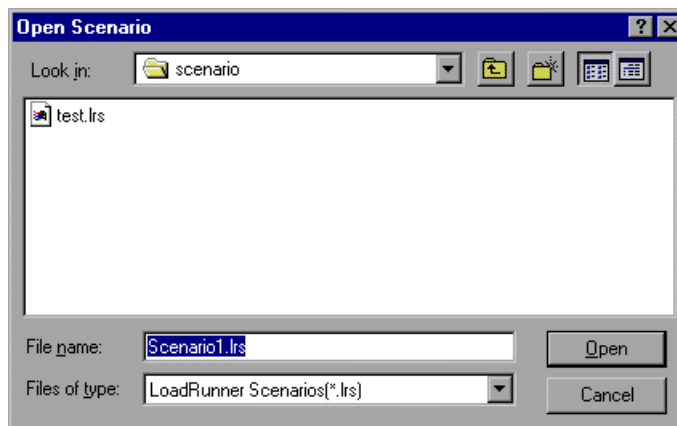
Opening an Existing Scenario

The Open command opens any existing scenario.

To open an existing scenario:



- 1 Choose **File > Open**, or click the **Open** button. The Open Scenario dialog box opens.



- 2 Click a file in the File Name list or type a file name in the File Name box.

- 3 Click **Open**. The File Open dialog box closes and the scenario appears in the LoadRunner Controller.

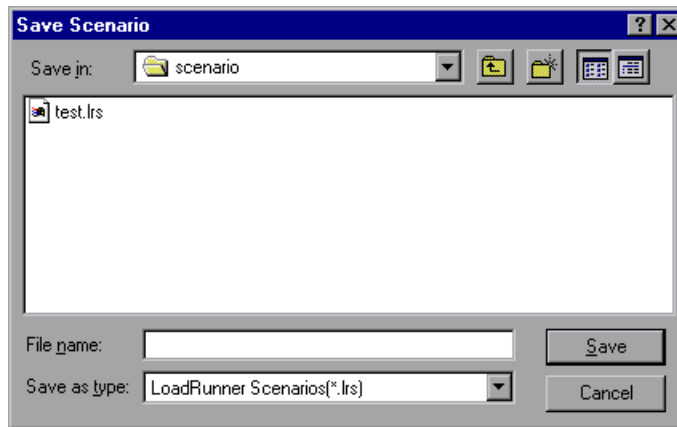
Saving a Scenario

The Save command saves the current scenario.

To save a scenario:



- 1 Choose **File > Save**, or click the **Save** button. The Save Scenario dialog box opens the first time you save a scenario.



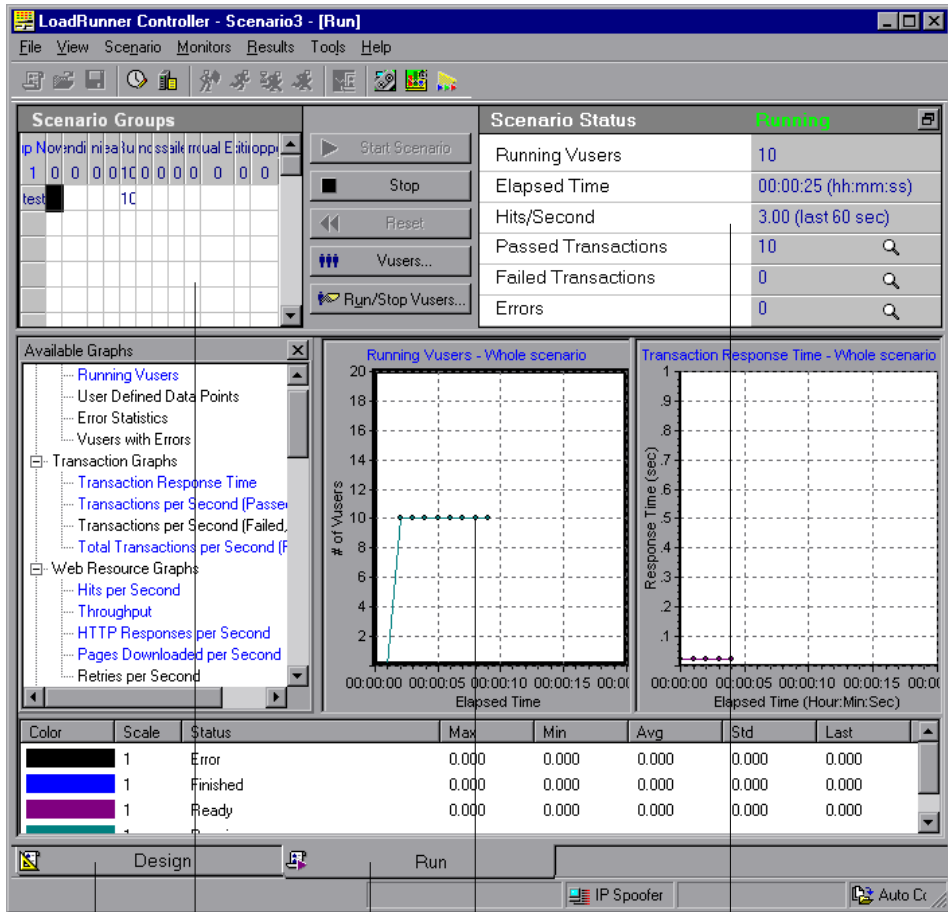
- 2 Type a scenario name in the File Name text box. Note that by default scenario files have the extension *.lrs*.
- 3 Click **Save**. The scenario is saved in the location you specified.

Closing a Scenario

Closing a scenario closes all the Controller windows. To close the scenario, choose **File > Close**. If you made changes to the scenario, a Save Changes message appears. Choose **Yes** to save the changes you made. All open windows and icons in the Controller close.

Running a Scenario

Once you have designed your scenario, you are ready to run it. You can control the Vusers and Vuser groups and monitor their performance online using the Run view of the LoadRunner Controller.



Design tab

Run tab

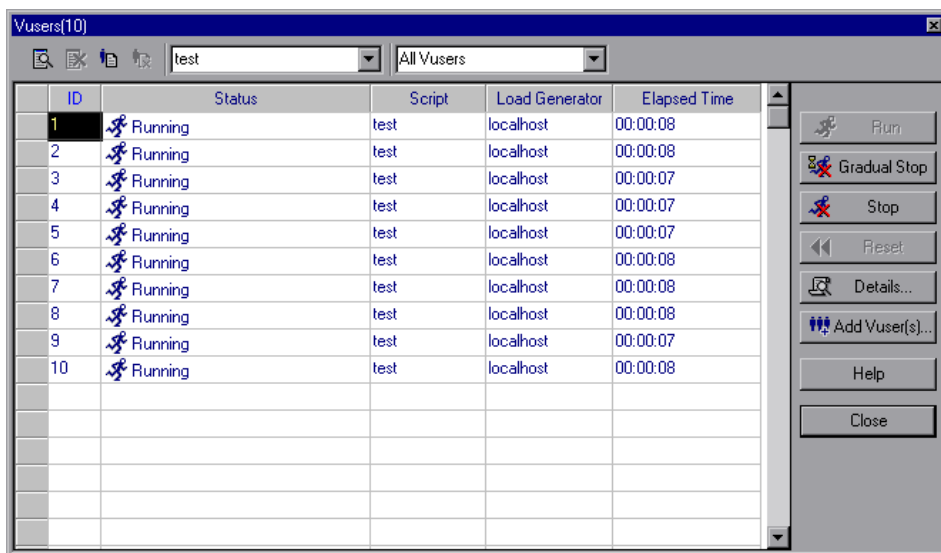

Scenario Groups pane

Online Monitor
Graphs

Scenario Status window

During scenario execution, you use the Scenario Groups pane in the Run view to monitor the actions of all the Vusers and Vuser groups in the scenario. The Status field of each Vuser group displays the current state of each Vuser in the group.

You can also manipulate individual Vusers within the Vuser groups you have defined by selecting a group and clicking the **Vusers** button. The Vusers dialog box appears, with a list of the ID, Status, Script, Load Generator, and Elapsed Time (since the beginning of the scenario) for each of the Vusers in the group.

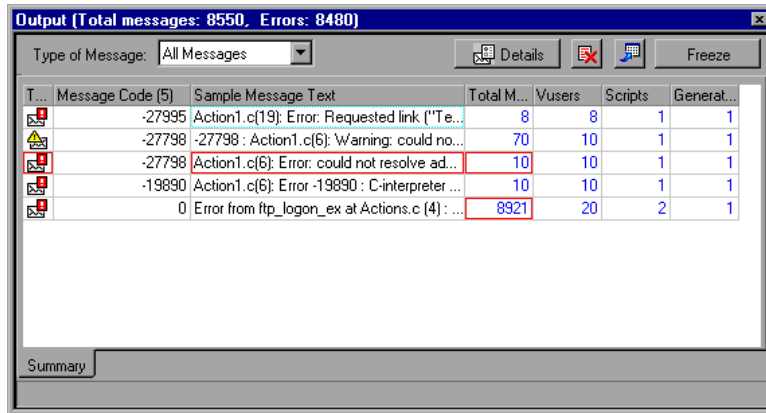


In addition, you can view a synopsis of the running scenario in the box at the top right-hand corner of the Run view.

Scenario Status		Running
Running Vusers	10	
Elapsed Time	00:00:38 (hh:mm:ss)	
Hits/Second	1.48 (last 60 sec)	
Passed Transactions	10	🔍
Failed Transactions	0	🔍
Errors	0	🔍

Note that you can detach the Scenario Status window from the Run view, thereby enlarging the Scenario Groups pane.

While the scenario runs, the Vusers and load generators send error, notification, warning, debug, and batch messages to the Controller. You can view these messages in the Output window (**View > Show Output**).



For more information on the Output window, see “Viewing the Output Window” on page 187.

You use the online monitors and online monitor graphs to monitor Vuser status, transactions, system resources, database server resources, network delay, streaming media resources, firewall server resources, ERP server resources, and Java performance, while running a scenario. For more information on online monitors, see Chapter 16, “Online Monitoring.”

Part II

Designing a Scenario

5

Creating a Manual Scenario

You can build a manual scenario by creating groups and specifying the scripts, the load generators, and the number of Vusers included in each group. You can also build a manual scenario using the Percentage Mode, which allows you to define the total number of Vusers to be used in the scenario, and assign load generators and a percentage of the total number of Vusers to each script.

This chapter describes how to create a manual scenario using the Vuser Group Mode. For information on creating a manual scenario using the Percentage Mode, see Chapter 6, “Creating a Manual Scenario Using the Percentage Mode.”

This chapter discusses:

- Creating Vuser Groups
- Configuring Vusers in a Vuser Group
- Configuring Vuser Run-Time Settings
- Configuring Load Generators
- Configuring Load Generator Settings
- Configuring WAN Emulation Settings
- Configuring Scripts
- Using Relative Paths for Scripts

About Creating a Scenario

To test your system with LoadRunner, you must create a scenario—a file with information about the test session. The scenario is the means by which you emulate a real-life user. The scenario contains information about how to emulate real users: the groups of virtual users (Vusers), the test scripts the Vusers will run, and the load generator machines upon which to run the scripts.

If you chose to create a regular manual scenario, each script you selected in the New Scenario dialog box is assigned to a Vuser group. To each Vuser group you then assign a number of virtual users. You can instruct all Vusers in a group to run the same script on the same load generator machine, or you can assign different scripts and load generators to the various Vusers in a group.

Once you create your Vuser groups, you select or build a schedule for your scenario. See Chapter 8, “Scheduling a Scenario” for more information on creating a scenario schedule.

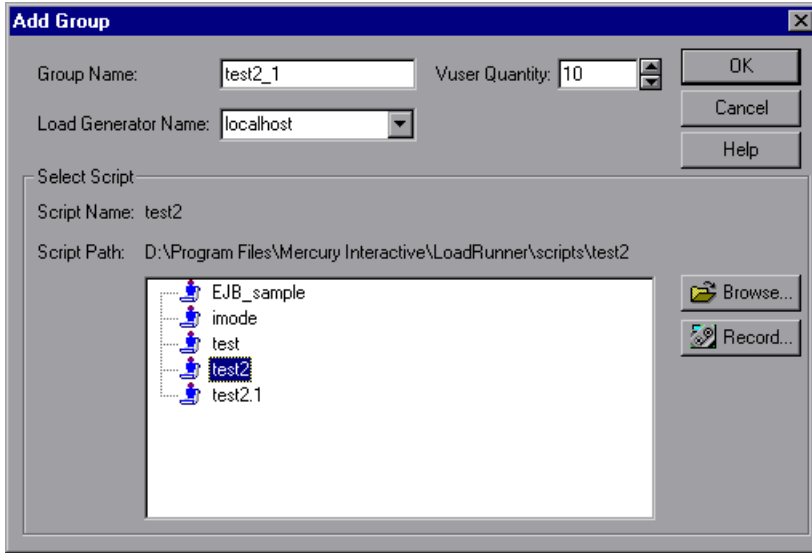
Creating Vuser Groups

A scenario consists of groups of Vusers which emulate human users interacting with your application. When you run a scenario, the Vusers generate load on the server, and LoadRunner monitors the server and transaction performance.

To create Vuser Groups:



- 1 Click the **Add Group** button to the right of the Scenario Groups pane. The Add Group dialog box opens:




- 2 In the Group Name box, enter a name for the Vuser group.

Note: The name used in the Group Name box is limited to a maximum of 55 characters.

- 3 From the Vuser Quantity box, select the number of Vusers that you want to create in the group.
- 4 Select a load generator from the Load Generator Name list. The Load Generator list contains all load generators that you previously added to the scenario.

To use a load generator that does not appear, select **Add** from the Load Generator Name list. 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.

By default, LoadRunner stores temporary files on the load generator during scenario 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.

To allow the load generator to take part in the scenario, check **Enable load generator to take part in the scenario**.

Click **More** to expand the dialog box and show the Add Load Generator tabs. For information on configuring settings for each load generator, see “Configuring Load Generator Settings” on page 64.

Click **OK** to close the Add Load Generator dialog box.

- 5 Select a script from the script list. The script list contains all scripts that you previously added to the scenario.

To use a script that does not appear, click the **Browse** button. Browse to select the path and file name of the new script. To use a VB Vuser script, select the *.usr* file.


Note: When you specify the location of a script, you can specify a location that is relative to the current scenario directory. For details, see “Using Relative Paths for Scripts” on page 83.

- 6 Click **OK** to close the Add Group dialog box. The new group's properties appear in the Scenario Groups pane.

Disabling a Vuser Group

By default, all the Vuser groups displayed in the Scenario Groups pane are enabled to run in the scenario. To disable a Vuser group, click the box to the left of the Vuser group name. The color of the group changes to gray, indicating that the group will not take part in the scenario. To re-enable the Vuser group, click the same box again.

Deleting a Vuser Group



To delete a Vuser group, click the **Remove Group** button to the right of the Scenario Groups pane, or right-click the Vuser group and select **Remove Group**.

Modifying a Vuser Group

You can modify the script, vuser quantity, and load generator for a Vuser group directly from the Scenario Groups pane of the Controller, or by using the Group Information dialog box.

To modify a Vuser Group directly from the Scenario Groups pane:

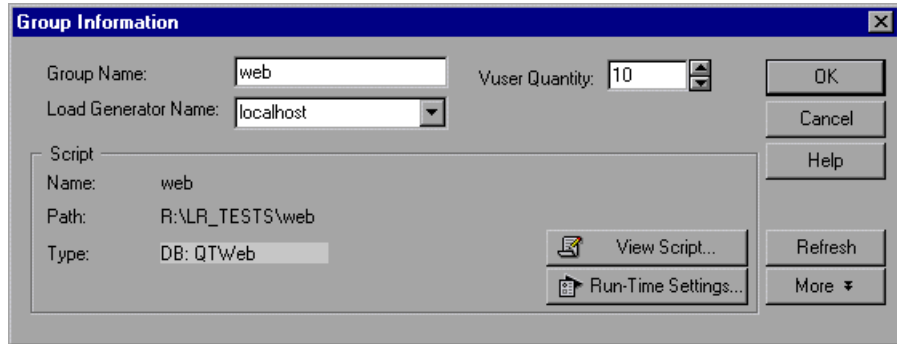
- 1** Select the **Group Name**, **Script Path**, **Quantity**, or **Load Generator** you want to modify.
- 2** Enter or select another name or number for the property.
- 3** To modify a Vuser group script's run-time settings, click the **Run-Time Settings** button to the right of the Scenario Groups pane. For details, see "Configuring Vuser Run-Time Settings", on page 59.



- 4** To edit a Vuser group's script, click the **View Script** button to the right of the Scenario Groups pane. LoadRunner's script generation tool, VuGen, opens. For more information on editing scripts, see the *LoadRunner Creating Vuser Scripts User's Guide*.

To modify a Vuser Group using the Group Information dialog box:

- 1 Click the **Details** button to the right of the Scenario Groups pane, or right-click the property you want to modify and select **Details**. The Group Information dialog box opens.



- 2 In the Group Name box, enter the Vuser group name.
- 3 From the Vuser Quantity box, select the number of Vusers that you want to run in the group.
- 4 Select a load generator from the Load Generator Name list. To use a load generator that does not appear, select **Add** from the Load Generator Name list and add a new load generator using the Add Load Generator dialog box.
- 5 To modify the run-time settings you specified while recording a script using VuGen, click **Run-Time Settings**. For more information about run-time settings, see “Configuring Scripts” on page 80.
- 6 To edit a Vuser group’s script, click **View Script**. LoadRunner’s script generation tool, VuGen, opens. For more information on editing scripts, see “Configuring Scripts” on page 80.
- 7 Click **OK** to close the Group Information dialog box.

Sorting Vuser Groups in the Scenario Groups Pane

Once you have created your Vuser groups, you can sort them by group name, script name, load generator name, or quantity of Vusers assigned to the group.

To sort Vuser groups:

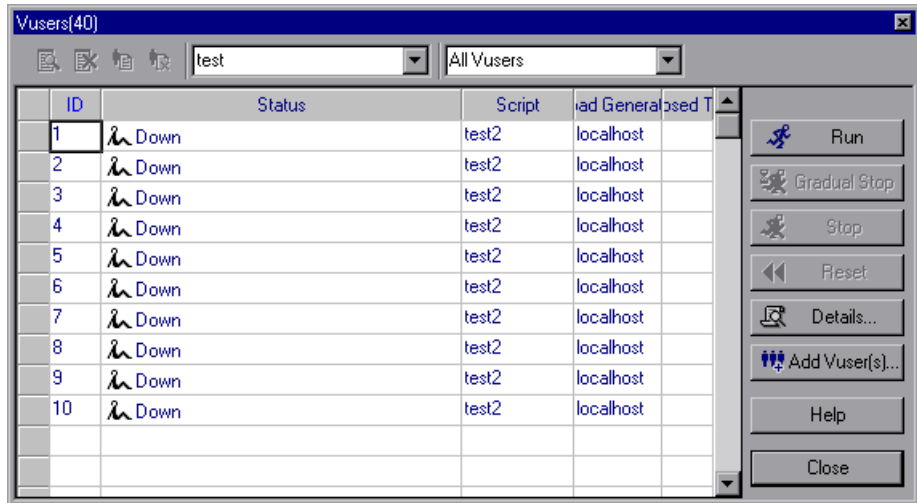
- Select the column by which you want to sort the groups. Click the column header.
- Alternatively, you can right-click anywhere within the column you want to sort, and select **Sort Groups**. Choose to sort by name, path, quantity, or generator.
- To instruct the Controller to automatically sort a new Vuser group entry, right-click the entry and select **Auto Sort**.

Configuring Vusers in a Vuser Group

You can define properties for individual Vusers within the Vuser groups you have defined using the Vusers dialog box. To each Vuser you can assign a different script and/or load generator machine.

To define properties for individual Vusers:

- 1 Select the Vuser group whose Vusers you want to modify, and click the **Vusers** button to the right of the Scenario Groups pane. The Vusers dialog box opens.



- 2 To change the script for an individual Vuser, select a different script in the Script column. Alternatively, you can click the **Details** button, and select a different script from the script list in the Vuser Information dialog box.
- 3 To change the load generator on which a Vuser runs, select a different load generator in the Load Generator column. Alternatively, you can click the **Details** button, and select a different load generator from the Load Generator Name list in the Vuser Information dialog box.


To use a load generator that does not appear, select **Add** from the Load Generator Name list and add a new load generator using the Add Load Generator dialog box.

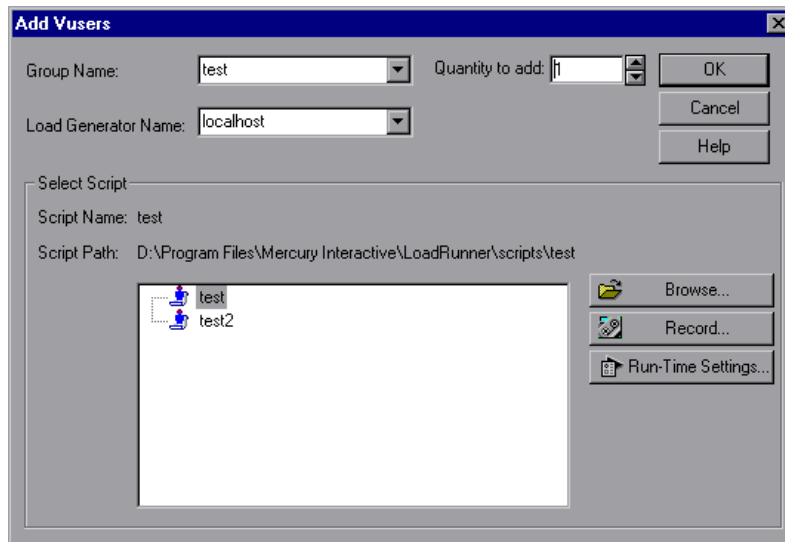
Adding Vusers to a Vuser Group

You add Vusers to a Vuser group and define their properties using the Add Vusers dialog box.

Note: You can activate additional Vusers while the scenario is running, using the Run/Stop Vusers dialog box. For more information, see “Manually Adding Vusers to a Running Scenario” on page 176.

To add Vusers to a Vuser group:

-  **1** In the Vusers dialog box, click the **Add Vuser(s)** button. The Add Vusers dialog box opens.



- 2** From the Group Name box, select the name of the Vuser group.
- 3** From the Quantity to add box, select the number of Vusers that you want to add to the group.
- 4** Select a load generator from the Load Generator Name list. To use a load generator that does not appear, select **Add** from the Load Generator Name list and add a new load generator using the Add Load Generator dialog box.

- 5 Select a script from the script list. The script list contains all scripts that you previously added to the scenario.

To use a script that does not appear, click the **Browse** button. Browse to select the path and file name of the new script. To use a VB Vuser script, select the *.usr* file.

Note: When you specify the location of a script, you can specify a location that is relative to the current scenario directory. For details, see “Using Relative Paths for Scripts” on page 83.

- 6 To modify the run-time settings you specified while recording a script using VuGen, click **Run-Time Settings**. For details, see “Configuring Vuser Run-Time Settings”, on page 59.
- 7 Click **OK** to close the Add Vusers dialog box. The new Vuser’s properties appear in the Vusers dialog box.

Configuring Vuser Run-Time Settings

You set the script’s run-time settings to customize the way the Controller executes a Vuser script. There are a few ways to display the script run-time settings:

- ▶ In the Group Information dialog box, click **Run-Time Settings**.
- ▶ In the Controller’s **Scenario Groups** pane, highlight a group or groups, and click **Run-Time Settings**.

The Run-Time Settings dialog box displays 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 Controller settings. Note that several protocols, such as Web and Java, have specific settings.

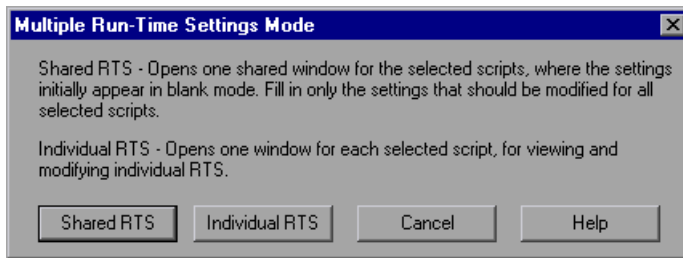
For information on each specific run-time setting, refer to the *LoadRunner Creating Vuser Scripts User’s Guide*.

Modifying the run-time settings for the new Vuser will modify the run-time settings for all the Vusers in the group. If a group contains more than one Vuser type, then you can modify the shared run-time settings, as described in “Modifying Run-Time Settings for Multiple Scripts” on page 60.

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

Modifying Run-Time Settings for Multiple Scripts

When you choose to modify script run-time settings, and you have selected multiple scripts, or a group with multiple scripts, The Controller displays the option of modifying the shared run-time settings:



Note: If one of the selected scripts does not support shared run-time settings, then you will only have the option of modifying each script's individual run-time settings. Shared RTS mode will be disabled for GUI or Astra LoadTest Vusers

Shared RTS: Opens one window containing all of the run-time settings in blank mode. In this mode, you set only the options that you would like to modify for all selected scripts. All other run-time settings remain unchanged.

Individual RTS: Opens a separate window for each selected script. In this mode, you modify each script's settings individually.

Modifying Shared Run-Time Settings

Any settings that you change in shared mode will be applied to all selected scripts. Any other settings remain unchanged. For example, if a dialog has checkboxes, they appear in disabled mode, meaning that they are neither selected nor cleared. If you select or clear a checkbox, then the change will apply to all selected scripts.

Some run-time settings cannot be modified in shared mode. These settings will not appear. To modify them, open the run-time settings for each individual script.

All Run-Time Setting buttons will be disabled, for example the **Change** and **Advanced** buttons in the Browser Emulation node.

The following nodes will not appear in shared mode:

- ▶ the **Java Environment Settings:Classpath** node
- ▶ the **Internet Protocol:ContentCheck** node
- ▶ the **Run Logic** node - for protocols which support the **Run Logic** node, the **Iterations** box will appear in the **Pacing** node.
- ▶ the nodes with tables in the format **Property, Value** for the protocols: Citrix ICA, Oracle NCA, and WAP, for example, the **Oracle NCA:Client Emulation** node

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 scenario, use LoadRunner's Options dialog box. For more information, see Chapter 10, "Configuring a Scenario." To set properties specific for each load generator, use the Load Generators dialog box as described below.

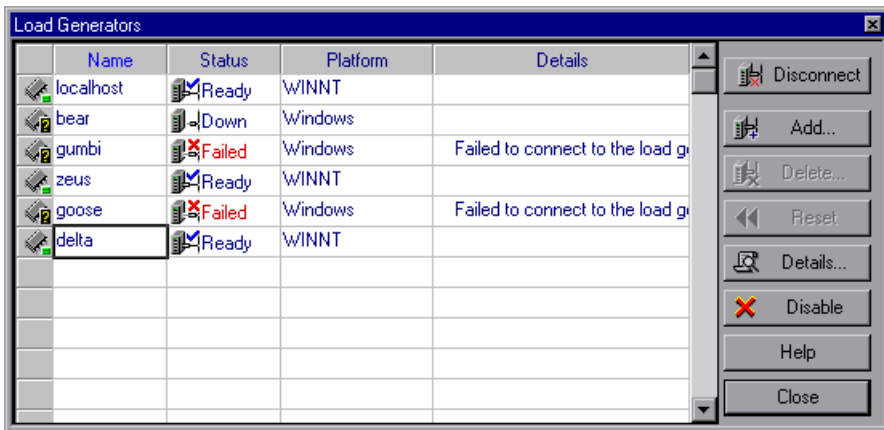
You can also indicate which load generators will run Vusers in the scenario. For example, if a load generator is unavailable for a particular scenario run, you can exclude it temporarily instead of removing it entirely from your list of load generators.

You select which load generators will take part in the scenario 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 **Scenario > 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.

By default, LoadRunner stores temporary files on the load generator during scenario 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.

To allow the load generator to take part in the scenario, check **Enable load generator to take part in the scenario**.

Click **More** to expand the dialog box and show the Add Load Generator tabs. For information on configuring these settings, see “Configuring Load Generator Settings” on page 64.

- 6 To remove a load generator, click **Delete**.
- 7 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.

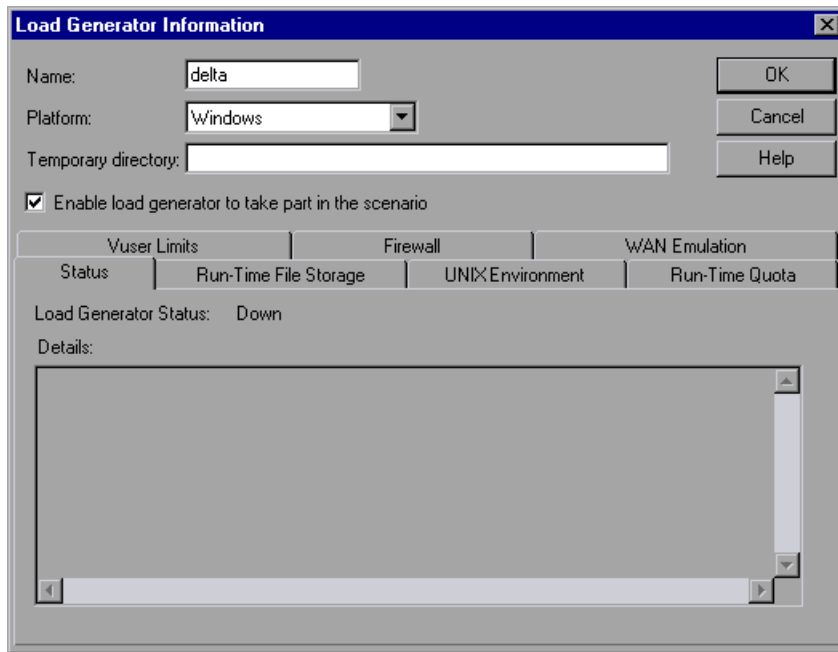
Configuring Load Generator Settings

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: Status, Run-Time File Storage, UNIX Environment, Run-Time Quota, Vuser Limits, Connection Log (Expert mode), Firewall, and WAN Emulation.

You can configure global settings for all load generators participating in the scenario, using the Options dialog box. For more information, see Chapter 10, “Configuring a Scenario.”

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, WAN Emulation, and Firewall (when the load generator is not the localhost) tabs.



- 2 Select the **Status** tab to display details of the Load Generator Status.

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

Load Generator Information

Name:

Platform:

Temporary directory:

Enable load generator to take part in the scenario

User Limits	Firewall	WAN Emulation	
Status	Run-Time File Storage	UNIX Environment	Run-Time Quota

Scripts and results stored:

- As defined in Tools->Options->Run-Time File Storage
- In temporary directory on delta
- On a shared network drive

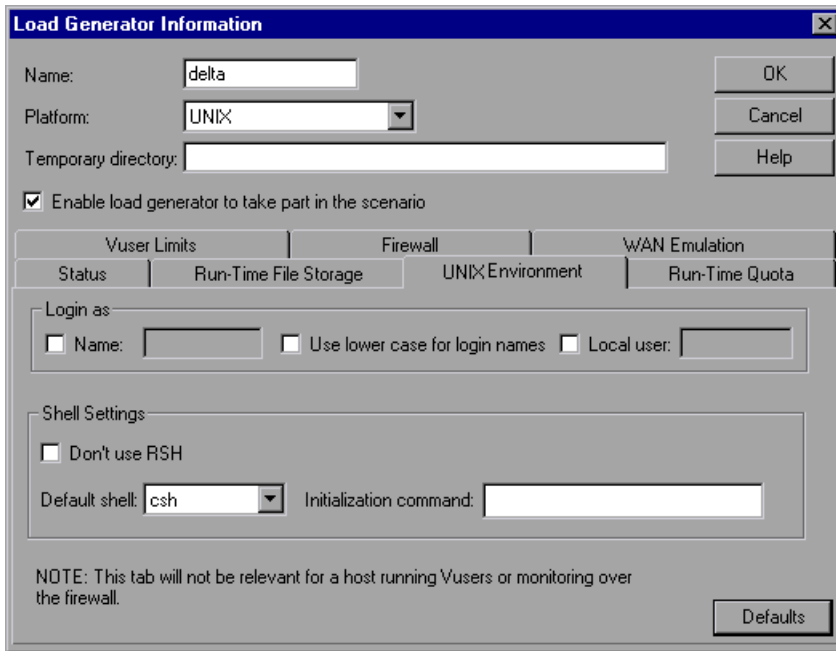
NOTE: This tab will not be relevant for a host monitoring over the firewall

Note: If you are monitoring over the firewall, the Run-Time File Storage settings are not relevant.

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 scenario scripts or results on a shared network drive, click **On a shared network drive**. To set the network location for the results, see Chapter 11, “Preparing to Run a Scenario.”

Note: If the load generator is *localhost*, then LoadRunner 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.



Note: If you are monitoring or running Vusers over the firewall, the UNIX Environment settings are not relevant.

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 589.

From the Default Shell box, select **cs**h (C Shell—the default), **bs**h (Bourne Shell), or **ks**h (Korn Shell).

To allow LoadRunner to run your application under the Korn shell, you first need to make sure that the *.profile* file contains all of the LoadRunner 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 LoadRunner environment settings.

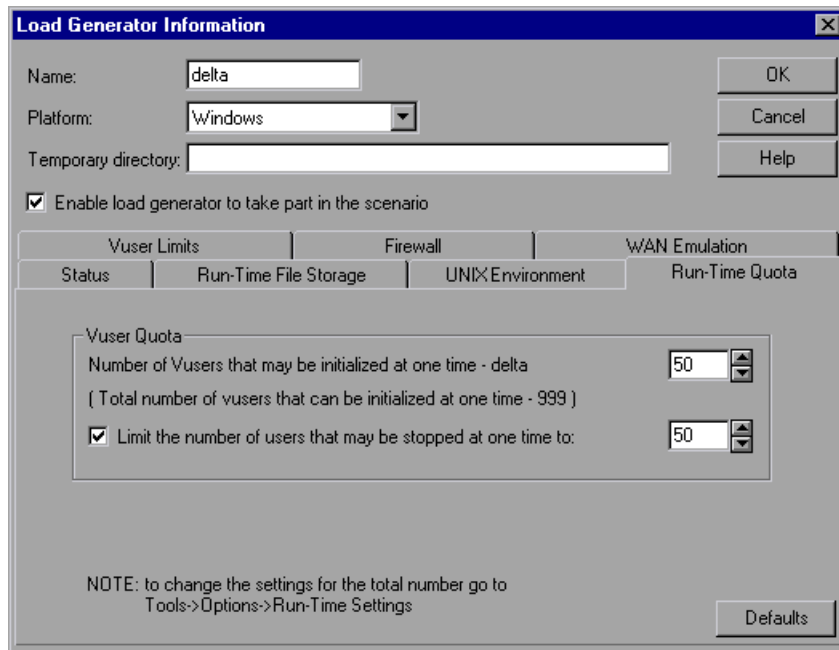
Note: If you are using a Korn shell (*ksh*), you must delete all LoadRunner settings from the *.cshrc* file (e.g. `M_LROOT`) before executing the scenario.

In the Initialization Command box, enter any command line options that LoadRunner 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.



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 scenario using the Run-Time Settings tab in the Options dialog box. For information on setting quotas globally for an entire scenario, see Chapter 10, “Configuring a Scenario.”

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

Load Generator Information

Name:

Platform:

Temporary directory:

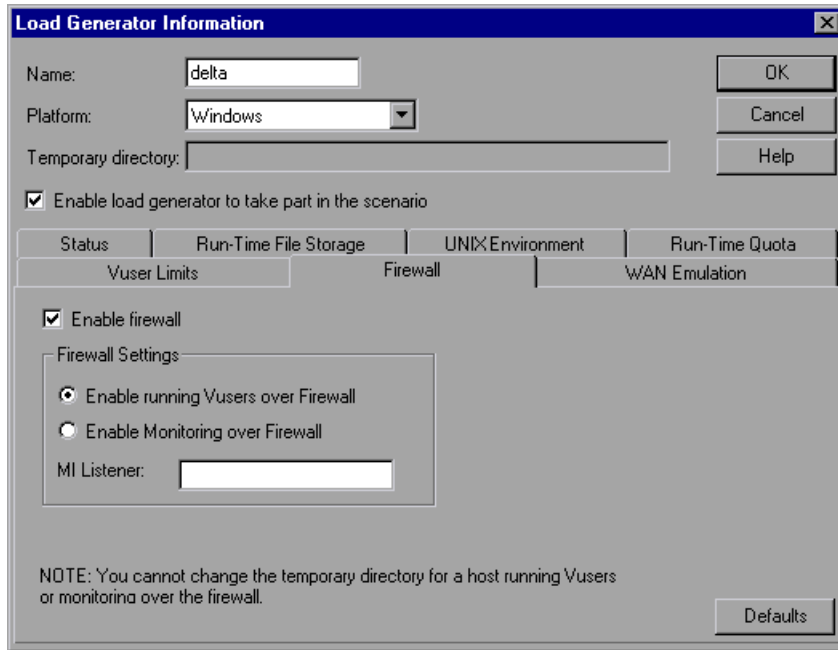
Enable load generator to take part in the scenario

Status	Run-Time File Storage	UNIX Environment	Run-Time Quota
Vuser Limits		Firewall	WAN Emulation
Available Types:		Maximum Active	
<input type="checkbox"/> GUI/WinRunner:	<input type="text" value="1"/>		
<input type="checkbox"/> RTE:	<input type="text" value="1000"/>		
<input checked="" type="checkbox"/> Other Vusers:	<input type="text" value="1000"/>		

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 LoadRunner**.

- 7 Select the **Firewall** tab to enable monitoring to occur through a firewall.



Note: You cannot change firewall settings while a load generator is connected. To disconnect a load generator, select the load generator in the Load Generators dialog box, and click **Disconnect**. The load generator status changes to **DOWN**, and you can change the settings.

The **Firewall** tab is disabled if the load generator is *Localhost*.

To obtain information for the monitors configured inside the firewall, or to run Vusers inside the firewall, select the **Enable firewall** check box and specify the Firewall Settings.

To enable LoadRunner to run Vusers over the firewall, click **Enable running Vusers over Firewall**. To enable LoadRunner to monitor the load generator machine over the firewall, click **Enable Monitoring over Firewall**. In the MI Listener box, specify the name of MI Listener the load generator is using.

Note: If you monitor or run Vusers over the firewall, the Temporary Directory option for storing files is disabled. Any location in the Temporary Directory box is erased.

- 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.

Status	Run-Time File Storage	UNIX Environment	Run-Time Quota
Vuser Status	Vuser Limits	Firewall	WAN Emulation
	Pending	Initializing	Active
GUI/WinRunner:	0	0	0
RTE:	0	0	0
Other Vusers:	0	0	0
<hr/>			
Totals:	0	0	0

Note: For information on the Connection Log tab available in Expert mode, see “Working in Expert Mode” on page 589.

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

Configuring WAN Emulation Settings

You can emulate the behavior of a wide variety of network infrastructures in your load testing scenario using the Shunra WAN Emulator.

About WAN Emulation

WAN emulation enables you to accurately test point-to-point performance of WAN-deployed products under real-world network conditions in your test environment. By introducing highly probable WAN effects such as latency, packet loss, link faults, and dynamic routing effects over your LAN, you can characterize many aspects of the WAN cloud and efficiently control your emulation in a single network environment. You can observe the effects of your emulation settings on the network performance in the WAN emulation monitoring reports.

For an explanation of the WAN emulation parameters and their default settings, see “WAN Emulation Parameters,” on page 78.

Note: WAN Emulation is only available for load generators running on a Windows platform. The WAN Emulation tab is disabled for load generators running on a UNIX platform.

Setting up the WAN Emulator

To use the Shunra WAN Emulator, you must first install the WAN Emulator Driver on the Load Generator machine from the LoadRunner Controller 7.6 CD. For instructions, refer to the *LoadRunner Controller Installation Guide*.

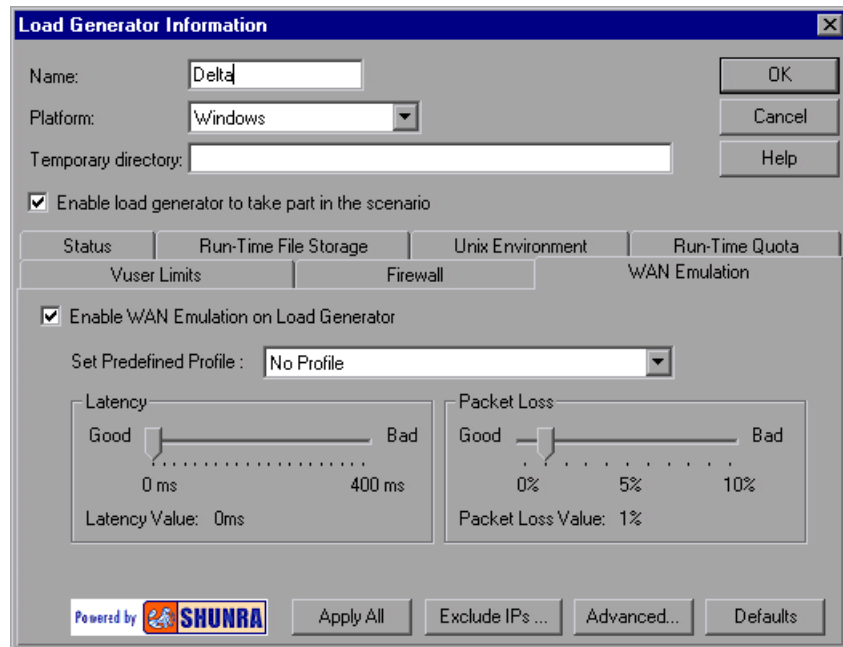
Note: WAN Emulation requires a special license. Contact Mercury Interactive's Customer Support Web site (<http://support.mercuryinteractive.com>) for licensing information.

Configuring the WAN Emulator

Once you have set up the Shunra WAN Emulator on the load generator machine, you can configure the settings from the Controller machine.

To configure WAN emulation settings:

- 1 From the Load Generator Information dialog box, select the **WAN Emulation** tab.



Note: WAN emulation is disabled if the load generator is *Localhost*.

- 2** Select the **Enable WAN Emulation on Load Generator** check box to enable WAN emulation to start automatically on scenario execution.

Note: You cannot change WAN emulation settings while a load generator is connected. To disconnect a load generator, select the load generator in the Load Generators dialog box, and click **Disconnect**. The load generator status changes to DOWN, and you can change the settings.

- 3** Set the Latency and Packet Loss parameters. To set a profile value manually, move the **Latency** and **Packet Loss** slider to the desired setting. To use a predefined WAN profile, select a profile from the **Set Predefined Profile** drop-down list.

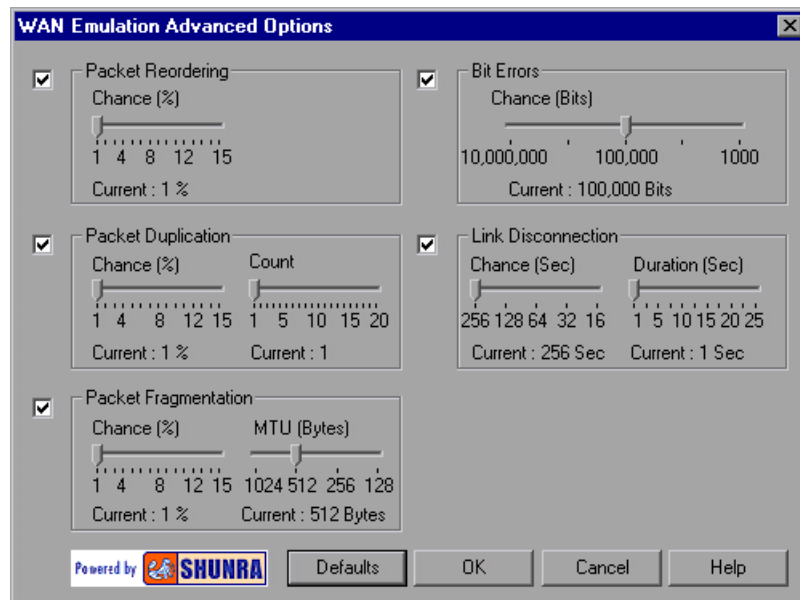
The following profile settings are available:

Profile	Description	Parameter Values
No Profile	This is the default setting. No profile has been selected, or a predefined profile has been changed manually.	Latency Value: 0ms Packet Loss Value: 1%
Metropolitan Area Network link	Emulates a metropolitan area network link.	Latency Value: 20ms Packet Loss Value: 1%
Mainland Low Congestion link (Terrestrial)	Emulates a mainland terrestrial link with low network traffic congestion.	Latency Value: 40ms Packet Loss Value: 1%
Mainland Congested link (Terrestrial)	Emulates a mainland terrestrial link with high network traffic congestion.	Latency Value: 100ms Packet Loss Value: 3%

Profile	Description	Parameter Values
Transatlantic Low Congestion link (Terrestrial)	Emulates an overseas terrestrial link with low network traffic congestion.	Latency Value: 60ms Packet Loss Value: 1%
Transatlantic Congested link (Terrestrial)	Emulates an overseas terrestrial link with high network traffic congestion.	Latency Value: 120ms Packet Loss Value: 3%
Transatlantic Low Congestion link (Satellite)	Emulates a satellite link with low network traffic congestion.	Latency Value: 280ms Packet Loss Value: 1%
Transatlantic Congested link (Satellite)	Emulates a satellite link with high network traffic congestion.	Latency Value: 400ms Packet Loss Value: 3%

The profile values are displayed beneath the setting range. To restore the default settings, click the **Defaults** button.

- To set the advanced options, click the **Advanced** button. The WAN Emulation Advanced options dialog box opens.



- 5 To enable an advanced option, select the option's check box. By default, all the options are enabled. To adjust an option setting, move the slider to the desired value. The profile values are displayed beneath the setting range.

Click the **Defaults** button to restore the default settings. To disable an option, clear the option's check box. A disabled option does not take part in the WAN emulation. Click **OK** to accept the advanced settings and close the WAN Emulation Advanced options dialog box.

- 6 To apply the WAN emulation settings to all load generators listed in the Load Generators dialog box, click the **Apply All** button.
- 7 Click **OK** to close the Add Load Generator or Load Generator Information dialog box and save your settings.

Excluding IP Addresses from WAN Emulation

In some situations you may want to exclude certain IP addresses from the WAN emulation. You can do this by setting the WAN emulator to refrain from affecting traffic to specified IP addresses. Network traffic which is not affected by the emulation will not suffer any WAN effects and will not be included in the WAN emulation monitoring reports.

Note: You do not need to exclude the Controller machine and the network file server (if you have the Network Installation configuration), as they are automatically excluded from the emulated WAN.

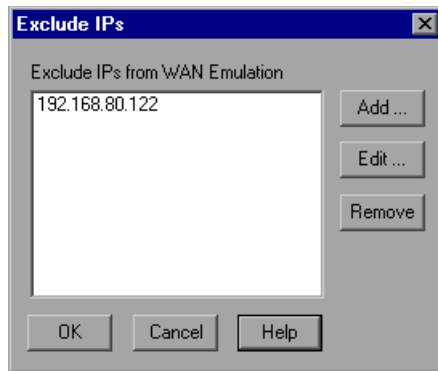
Examples of situations where you might choose to exclude IP addresses from an emulated WAN include the following:

- ▶ In a Multiprotocol scenario that includes a Web server and a database server; where information from the database server is not required as a part of the load test.
- ▶ Where a user runs and stores scripts on a shared network drive.
- ▶ Where the Controller is running or monitoring Vusers over a firewall using the TCP configuration. If the MI Listener is on a different machine to the Controller, the MI Listener machine should be excluded.

- Where the Controller is running or monitoring Vusers over a firewall using the HTTPS configuration. The IP address of the proxy server should be excluded.

To exclude an IP address from WAN emulation:

- 1 Select **Exclude IPs** from the **WAN Emulation** tab in the Load Generator Information dialog box. The Exclude IPs dialog box opens.



- 2 To exclude an IP address from WAN emulation, click the **Add** button. The Add IP dialog box opens.



- 3 Type the name or IP address of the machine you want to exclude from WAN emulation. Click **OK** to add the IP address to the Exclude IPs from WAN Emulation list.

Note: If you type the name of a machine, LoadRunner resolves the name and replaces it with the machine's IP address in the Exclude IPs list.

You can edit an IP address by selecting it from the Exclude IPs list and clicking the **Edit** button. In the Edit Machine dialog box, edit the IP address and click **OK** to save your changes and close the dialog box.

You can remove an IP address from the list by selecting it and clicking the **Remove** button.

Stopping and Restarting WAN Emulation

You can stop and restart WAN emulation at any time during a scenario run.

To stop or restart WAN emulation during scenario execution:

- 1** To stop WAN emulation select **Scenario > Stop WAN Emulation**.
- 2** To restart WAN emulation, select **Scenario > Restart WAN Emulation**.

WAN Emulation Parameters

You can set values for the following WAN emulation parameters:

Parameter	Description	Default Setting
Latency	The time, in milliseconds, it takes an IP packet to cross the WAN. This is usually effected by the geographical distance, the available bandwidth, the network load on the route between the two ends, and whether this is a terrestrial link or not.	The default setting is 0ms.
Packet Loss	The chance of losing IP packets while data travels through a WAN. Packets can get lost due to link faults or due to extreme network load.	The default setting is 1%.

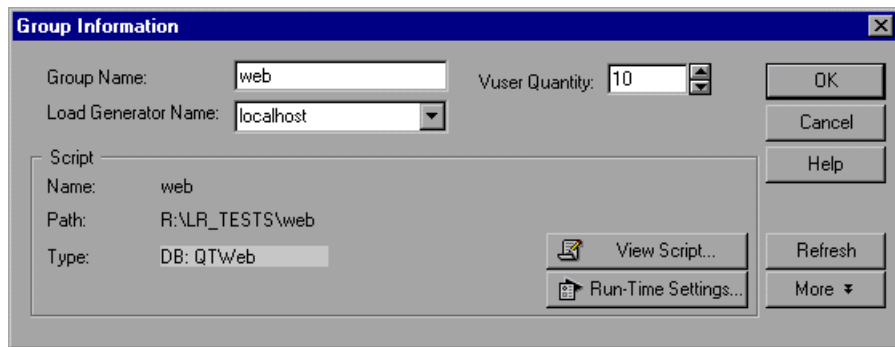
Parameter	Description	Default Setting
Packet Reordering	The chance of a packet order changing while it travels through the WAN cloud.	The default setting is 1%.
Packet Duplication	The chance of a packet duplication occurring while it travels through the WAN cloud. Count is the number of copies of each packet that will be created when duplication occurs.	Default Chance setting is 1%. Default Count setting is 1.
Packet Fragmentation	The chance of a packet fragmentation occurring (due to a short Maximum Transmission Unit) while it travels through the WAN cloud. MTU is the largest size packet or frame (specified in bytes), that can be sent in a packet- or frame-based network such as the internet.	Default Chance setting is 1%. Default MTU setting is 512 Bytes.
Bit Errors	The chance of bit errors occurring while a packet travels through the WAN cloud.	Default Chance setting is 100,000 Bits.
Link Disconnection	The chance (average frequency) of a network disconnection occurring while a packet travels through the WAN cloud, and the disconnection time span.	Default Chance setting is one disconnection every 256 seconds. Default Duration is 1 second.

Configuring Scripts

Once you have selected a script for a Vuser or Vuser group, you can edit the script, or view the details of the script you selected, from the Vusers or Group Information dialog boxes.

To edit and view the details of a script used by a Vuser group:

- 1 Select the Vuser group whose script you want to modify, and click the **Details** button to the right of the Scenario Groups pane, or right-click the Vuser group and select **Details**. The Group Information dialog box opens displaying the Name, Path, and Type of the script.



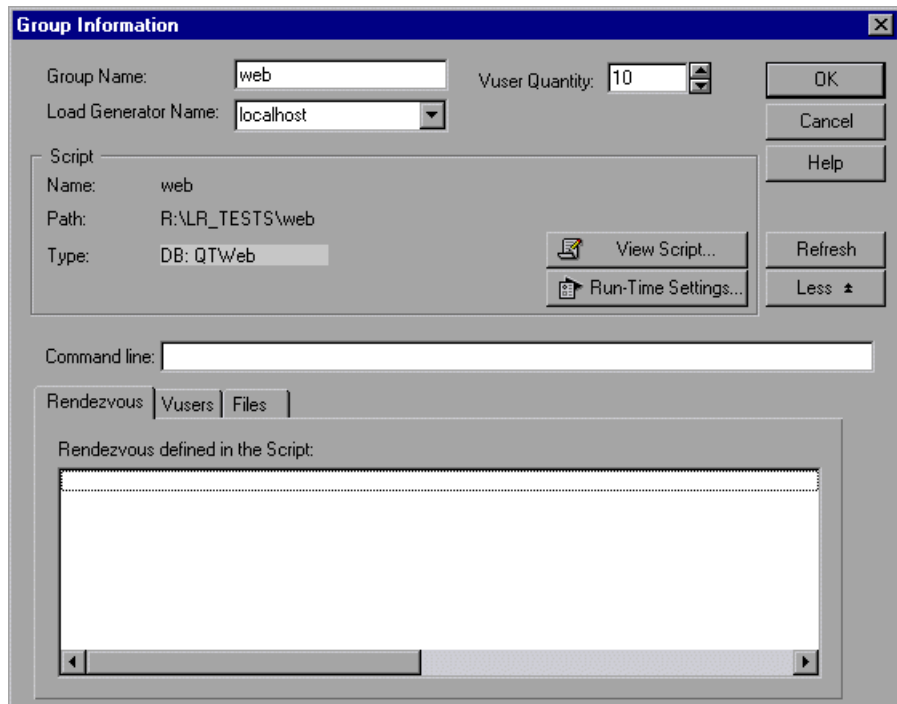
- 2 Click **Run-Time Settings** to set the script's run-time settings (optional). For details, see "Configuring Vuser Run-Time Settings", on page 59.

Note: If you modify the run-time settings from the Controller, LoadRunner 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, see the *LoadRunner Creating Vuser Scripts User's Guide*.

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

- 4 Click **More** to expand the Group Information dialog box and view additional script information.



- 5 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 *LoadRunner Creating Vuser Scripts User's Guide*.

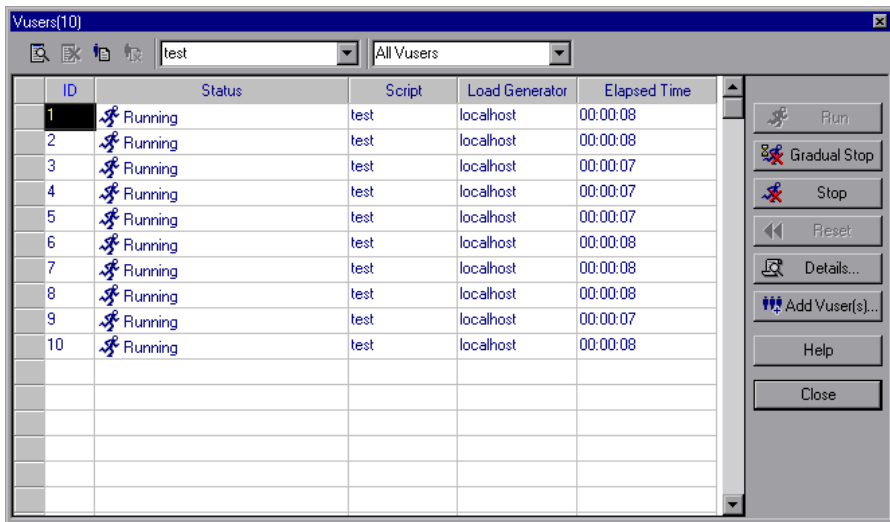
- 6 To see the rendezvous points included in the selected script, click the **Rendezvous** tab.

- 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 **OK** to close the Group Information dialog box.

To edit and view the details of a script used by an individual Vuser:



- 1 Click the **Vusers** button to the right of the Scenario Groups pane. The Vusers dialog box opens.



To view details of a script, click **Details**. The script's name and path are displayed in the Vuser Information dialog box. To select a different script, click the **Browse** button and select the path and file name of the new script. To select a VB Vuser script, browse to locate the *.usr* file.

Note: When you specify the location of a script, you can specify a location that is relative to the current scenario directory. For details, see “Using Relative Paths for Scripts” on page 83.

- 2** To edit a script, right-click the script in the Vusers dialog box, and select **View Script**. The script generation tool, VuGen, opens. For more information on editing scripts, see the *LoadRunner Creating Vuser Scripts User's Guide*.
- 3** To modify the run-time settings you specified while recording a script using VuGen, right-click the script in the Vusers dialog box, and select **Run-Time Settings**. Note that modifying the run-time settings for one Vuser will modify the run-time settings for all the Vusers in the group that are using the same script.

If you highlight more than one script, you can modify the run-time settings in shared mode, as described in “Modifying Run-Time Settings for Multiple Scripts” on page 60

For details about each run-time setting, see the *LoadRunner Creating Vuser Scripts User's Guide*.

Using Relative Paths for Scripts

When you specify the location of a script, you can specify a relative location. The location can be relative to the current scenario directory, or the LoadRunner installation directory.

You can specify a path relative to the current scenario 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 scenario directory. |
| .. \ | indicates that the path is relative to the location of the parent directory of the scenario directory. |

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

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

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

```
%\scripts\user1
```

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

When you run a scenario, 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 Controller to store the script on a shared network drive (see Chapter 10, “Configuring a Scenario.”) If you configure the Controller 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 Vuser scripts and their paths. A script’s path is based on the Controller 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 Controller load generator’s mapping to the Vuser load generator’s mapping. For more information see Appendix B, “Performing Path Translation.”

6

Creating a Manual Scenario Using the Percentage Mode

You build a manual scenario in the Percentage Mode by defining the total number of Vusers to be used in the scenario, and assigning load generators and a percentage of the total number of Vusers to each script. This chapter describes how to create a manual scenario in the Percentage Mode.

This chapter discusses:

- Defining the Total Number of Vusers
- Assigning Properties to Scripts
- Configuring Scripts
- Converting a Scenario to the Vuser Group Mode

About Creating a Manual Scenario Using the Percentage Mode

When you design a regular manual scenario, you create Vuser groups, assign them scripts, load generator machines, and virtual users. When you work in the Percentage Mode, you define the total number of Vusers to be used in the scenario, and assign load generators and a percentage of the total number of Vusers to each script.


When you create a new scenario, you can access the Percentage Mode directly by selecting **Use the Percentage Mode to distribute the Vusers among the scripts** in the New Scenario dialog box. You can also convert a scenario created in the Vuser Group Mode to the Percentage Mode by selecting **Scenario > Convert Scenario to the Percentage Mode**.

When converting your scenario from Vuser Group Mode to Percentage Mode, note the following:

- ▶ If you defined multiple scripts for a Vuser group, the number of Vuser scripts created in the Percentage Mode will equal the number of scripts defined for the group.
- ▶ <All Load Generators> will be assigned to all Vuser scripts created in the Percentage Mode. If you defined multiple load generators for a Vuser group, the Vusers you assign to the script(s) in the Percentage Mode will be distributed evenly among the load generators you previously assigned to the group.
- ▶ All Vuser group schedule settings will be lost. All profiles will contain scenario schedule settings only.

Defining the Total Number of Vusers

When you build a scenario in the Percentage Mode, you define a total number of Vusers to be used in the scenario, rather than a number of Vusers per script. You enter this number in the Scenario Schedule pane.

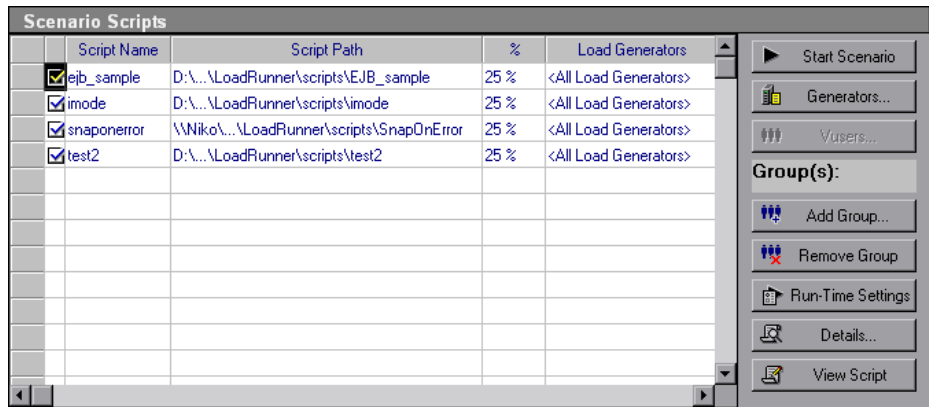


The screenshot shows the 'Scenario Schedule' dialog box. It has a title bar 'Scenario Schedule'. Inside, there is a 'Schedule Name' dropdown menu with 'Default Schedule' selected. Below that, the 'Mode' is set to 'Scenario Scheduling'. The 'Scenario Duration' is 'Until Completion'. The 'Load Behavior' is 'Load all Vusers simultaneously'. At the bottom, there is a 'Total Number of Vusers' text box containing the number '100'. Below the text box is an 'Edit Schedule...' button with a clock icon.

For information on creating a scenario schedule, see Chapter 8, “Scheduling a Scenario.” Note that the Vuser Group settings are not available for the Percentage Mode.

Assigning Properties to Scripts

The Scenario Scripts window displays a list of scripts you selected in the New Scenario dialog box, or defined in the Vuser Group Mode.



The % column indicates the percentage of the total number of Vusers that is automatically distributed to each Vuser script. During scenario execution, each script runs the percentage of Vusers assigned to it. The Load Generators column automatically contains <All Load Generators> for each Vuser script.

Note: If you defined multiple load generators for a Vuser group, the Vusers you assign to the script(s) in the Percentage Mode will be distributed evenly among the load generators you previously assigned to the group.

For each script you can modify:


- the percentage of the total number of Vusers assigned to the script
- the load generator(s) on which the Vusers will execute the script

To modify the percentage of Vusers assigned to a script:

In the script's % column, enter a percentage of the total number of Vusers you defined in the Scenario Schedule pane. The percentages assigned to the other scripts change to create a total of 100 percent for all of the Vuser scripts.

To modify the load generator(s) for a script:

- 1** In the script's Load Generators column, select one or more machine(s) from the Load Generator Name list, and click **OK**. If you select multiple machines, the Vusers assigned to the script are distributed evenly among the load generators.
- 2** Alternatively, you can choose **Add** to add a load generator to the list. 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.

By default, LoadRunner stores temporary files on the load generator during scenario 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.

To allow the load generator to take part in the scenario, check **Enable load generator to take part in the scenario**.

Click **More** to expand the dialog box and show the Add Load Generator tabs. For information on configuring settings for each load generator, see “Configuring Load Generator Settings” on page 64.

Click **OK** to close the Add Load Generator dialog box. LoadRunner adds the new load generator to the Load Generator Name list. To include the new load generator in your scenario, select it from the Load Generator Name list, and click **OK**.

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

Note: The Controller monitors a Windows load generator machine’s CPU usage and automatically stops loading Vusers on the overloaded load generator, and distributes them among other load generators taking part in the scenario. For more information, see “Load Balancing,” on page 107. 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 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. You can also use the Load Generators dialog box to indicate which load generators will run Vusers in the scenario. For example, if a load generator is unavailable for a particular scenario run, 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 61. To configure additional load generator settings, see “Configuring Load Generator Settings” on page 64.

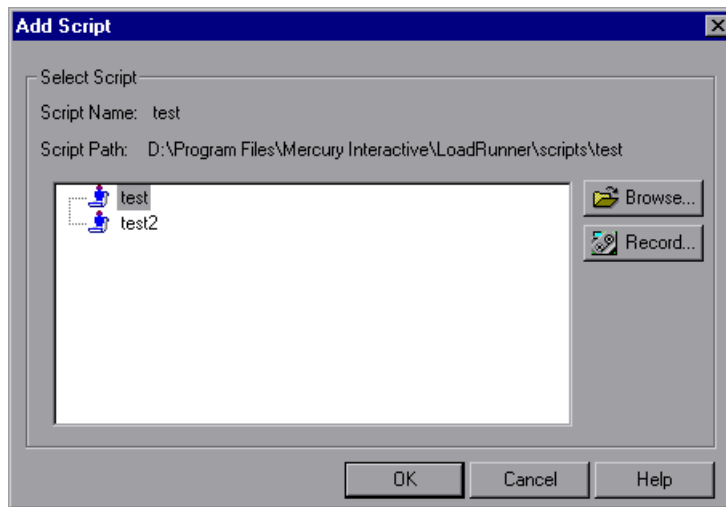
To configure global settings for all load generators participating in the scenario, use LoadRunner’s Options dialog box. For more information, see Chapter 10, “Configuring a Scenario.”

Configuring Scripts

You can add a script to the Scenario Scripts list using the Add Script dialog box. Once you have added the script to the list, you can view the details of the script you selected, edit the script, enable/disable it, or change its run-time settings.

To add a script:

- 1 Click the **Add Script** button to the right of the Scenario Scripts window, or right-click within a column and select **Add Script**. The Add Script dialog box opens.



- 2 Click the **Browse** button to the right of the path box. The Open Test dialog box opens.

Select the path and file name of the new script. To select a VB Vuser script, browse to locate the *.usr* file.


Note: When you specify the location of a script, you can specify a location that is relative to the current scenario directory. For details, see “Using Relative Paths for Scripts” on page 83.

- 3 Click **Open** to select the files. The Open Test dialog box closes, and the new script name appears in the Add Script dialog box.
- 4 Click **OK** to close the Add Script dialog box and enter the new script information in the Scenario Scripts window.

To view script details:



- 1 Select the script and click the **Details** button to the right of the Scenario Scripts window, or right-click the script and select **Details**. The Script Information dialog box opens, displaying the Path, Name, and Type of the script you selected.



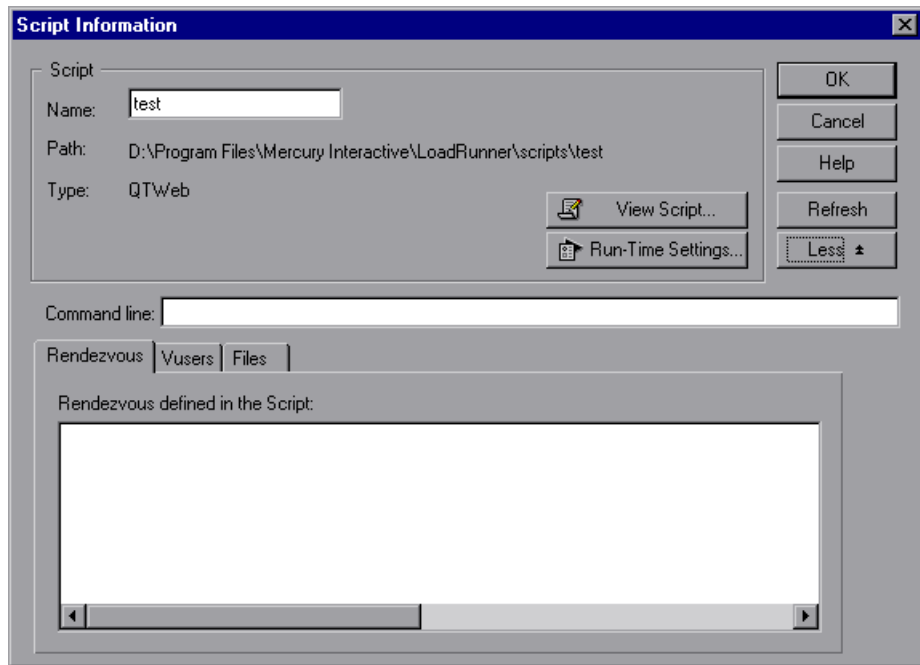
- 2 To modify the run-time settings you specified while recording a script using VuGen, click **Run-Time Settings**. For details, see “Configuring Vuser Run-Time Settings”, on page 59.

Note: If you modify the run-time settings from the Controller, LoadRunner 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, see the *LoadRunner Creating Vuser Scripts User's Guide*.

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

- 4 Click **More** to expand the Script Information dialog box and view additional script information.




- 5 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 *LoadRunner Creating Vuser Scripts User's Guide*.

- 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.

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

To delete a script:



Select the script and click the **Remove Script** button to the right of the Scenario Scripts window, or right-click the script and select **Remove Script**.

To disable a script

Click the box to the left of the Vuser script name. The color of the script entry changes to gray, indicating that the script will not take part in the scenario. To re-enable the Vuser script, click the same box again.

Converting a Scenario to the Vuser Group Mode

You can convert a scenario created in the Percentage Mode to the Vuser Group Mode by selecting **Scenario > Convert Scenario to the Vuser Group Mode**.

When converting your scenario from Percentage Mode to Vuser Group Mode, note the following:

- Each script will be converted to a Vuser group.
- If you defined multiple load generators for a Vuser script, the Vuser group that is created when converting the scenario will also contain multiple load generators.
- All schedule settings will be kept.

7

Creating a Goal-Oriented Scenario

You build a goal-oriented scenario for an application by defining the goals you want your test to achieve. This chapter describes how to create a goal-oriented scenario.

This chapter discusses:

- Defining Scenario Goals
- Assigning Properties to Scripts
- Configuring Scripts

About Planning a Goal-Oriented Scenario

In a goal-oriented scenario, you define the goals you want your test to achieve, and LoadRunner automatically builds a scenario for you, based on these goals. You can define five types of goals in a goal-oriented scenario: the number of virtual users, the number of hits per second (Web Vusers only), the number of transactions per second, the number of pages per minute (Web Vusers only), or the transaction response time you want your scenario to reach. You define one of these scenario goals using the Edit Scenario Goal dialog box. For more information on this dialog box, see “Defining Scenario Goals” on page 100.

Note: To run a Transactions per Second or Transaction Response Time goal type, your script must contain transactions. For each of these goal types, you define the transaction in your script that you want to test.

Virtual Users Goal Type

If you want to test how many Vusers your application can run simultaneously, it is recommended that you define a Virtual Users goal type. Running this type of goal-oriented scenario is similar to running a manual scenario. For more information on defining this goal type, see “Defining Scenario Goals” on page 100.

Pages per Minute and Hits/Transactions per Second Goal Types

If you want to test the strength of your server, it is recommended that you define a Hits per Second, Pages per Minute, or Transactions per Second goal type. Specify a minimum-maximum range of Vusers for LoadRunner to run, and a Transaction Name for the Transactions per Second goal type.

The Controller attempts to reach the goal you defined using a minimum number of Vusers. If it cannot reach this goal using the minimum number of Vusers, the Controller increases the number of Vusers until the maximum number you defined is reached. If your goal cannot be reached with the maximum number of Vusers you specified, increase this number and execute your scenario again. For more information regarding the formula used by the Controller in running the Pages per Minute and Hits/Transactions per Second goal types, see page 104.

Transaction Response Time Goal Type

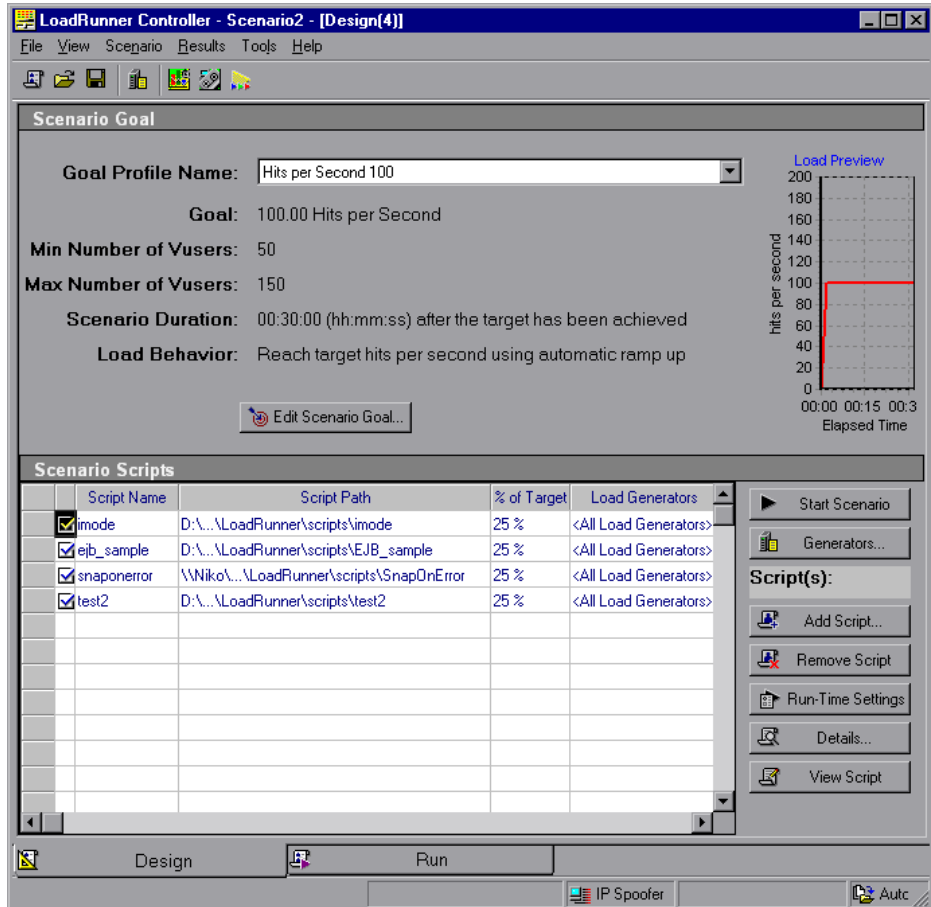
If you want to test how many Vusers can be run simultaneously without exceeding a desired transaction response time, it is recommended that you define a Transaction Response Time goal type. Specify the name of the transaction in your script that you want to test, and a minimum-maximum range of Vusers for LoadRunner to run. The transaction response time you specify should be a pre-defined threshold value. For example, if you do not want a customer to wait more than five seconds to log in to your e-commerce site, specify a maximum acceptable transaction response time of five seconds. Set the minimum and maximum number of Vusers to the minimum-maximum range of customers you want to be able to serve simultaneously.

If the scenario does not reach the maximum transaction response time that you defined, your server is capable of responding within a reasonable period of time to the number of customers you want to be able to serve simultaneously. If the defined response time is reached after only a portion of the Vusers has been executed, or if you receive a message that the defined response time will be exceeded if the Controller uses the maximum number of Vusers defined, you should consider revamping your application and/or upgrading your server software and hardware.

Note: In order for a Transaction Response Time goal-oriented scenario to be effective, you must choose your transaction carefully, ensuring that it performs effective hits on the server.

Defining Scenario Goals

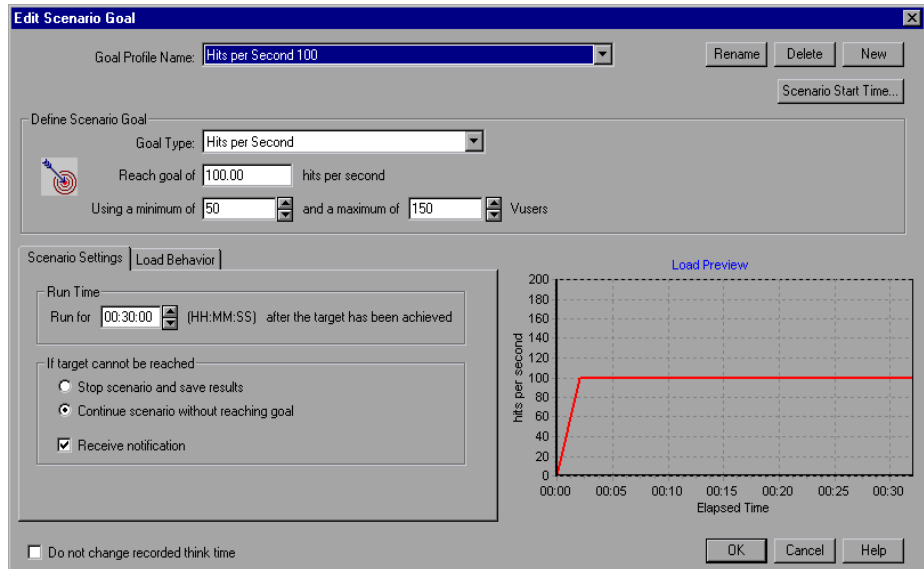
When you choose to create a goal-oriented scenario, the Controller displays the Scenario Goal and Scenario Scripts windows.



The Scenario Goal window contains basic scenario information, as defined in the Edit Scenario Goal dialog box.

To define a scenario goal:

- 1 Click the **Edit Scenario Goal** button. The Edit Scenario Goal dialog box opens.



- 2 Select a **Goal Profile Name**. To enter a new name, click **New**, type the new goal profile name in the New Goal Profile dialog box, and click **OK**. The new goal profile name appears in the selector.

To rename a goal profile, click **Rename** and enter the new goal profile name in the New Goal Profile dialog box.

To delete a goal profile, select it and click **Delete**.

- 3 In the Define Scenario Goal box, select a **Goal Type**.
 - If you select **Virtual Users**, enter a target number of virtual users that you would like your scenario to reach.
 - If you select **Hits per Second**, enter a target number of hits per second (HTTP requests per second) that you would like your scenario to reach, and select a minimum and maximum number of Vusers for the scenario.

- If you select **Transactions per Second**, enter a target number of transactions per second that you would like your scenario to reach, and select a minimum and maximum number of Vusers for the scenario. In addition, select a static script transaction for your scenario to test, or enter the name of an automatic script transaction that you have recorded in the Transaction Name box.

Note: VuGen automatically defines each *Init*, *Action*, and *End* unit as a transaction. In addition, you can insert a static transaction in your script using the Start Transaction and End Transaction functions.

- If you select **Transaction Response Time**, enter a target transaction response time that you would like your scenario to reach, and select a minimum and maximum number of Vusers for the scenario. In addition, select a static script transaction for your scenario to test, or enter the name of a dynamic script transaction that you have recorded in the Transaction Name box.
 - If you select **Pages per Minute**, enter a target number of downloaded pages per minute that you would like your scenario to reach, and select a minimum and maximum number of Vusers for the scenario.
- 4 In the Scenario Settings tab, select the amount of time that you want your scenario to run after your target has been reached.
 - 5 Choose whether you want to stop the scenario and save the scenario results or continue the scenario, if LoadRunner does not succeed in reaching the target that you defined. If you want LoadRunner to send you an error message indicating that your target was not reached, select **Receive notification**.

- 6 Select the **Load Behavior** tab. If you selected the Transactions per Second or Transaction Response Time goal types, choose whether you want LoadRunner to reach your target by automatically running a default number of Vusers in every batch, or after a certain period of the scenario has elapsed. If you selected the Pages per Minute, Virtual Users, or Hits per Second goal types, choose whether you want LoadRunner to reach your target by automatically running a default number of Vusers in every batch, after a certain period of the scenario has elapsed, or by gradation (x number of Vusers/pages/hits every x amount of time).
-

Note: The Load Preview graph visually displays the load behavior you define.

- 7 Select **Do not change recorded think time** if you want LoadRunner to run the scenario using the think time recorded in your script. Note that if you select this option, you may need to increase the number of Vusers in your scenario in order to reach your target.
 - 8 Click **OK** to close the Edit Scenario Goal dialog box. The scenario target information you entered appears in the Scenario Goal window.
-

Note: When you run a goal-oriented scenario, the goal you defined is displayed in the appropriate graph, along with the scenario results. This enables you to compare the results with your target goal.

Understanding the Hits/Transactions per Second and Pages per Minute Goal Types

When you define a Pages per Minute or Hits/Transactions per Second goal type, the Controller divides the target you defined by the minimum number of Vusers you specified, and determines the target number of hits/transactions per second or pages per minute that each Vuser should reach. The Controller then begins loading the Vusers according to the load behavior settings you defined, as follows:

- If you selected to run the Vusers automatically, LoadRunner loads fifty Vusers in the first batch. If the maximum number of Vusers defined is less than fifty, LoadRunner loads all of the Vusers simultaneously.
- If you chose to reach your target after a certain period of the scenario elapses, LoadRunner attempts to reach the defined target within this period of time. It determines the size of the first batch of Vusers based on the time limit you defined and the calculated target number of hits, transactions, or pages per Vuser.
- If you chose to reach your target by gradation (x number of pages/hits every x amount of time), LoadRunner calculates the target number of hits or pages per Vuser and determines the size of the first batch of Vusers accordingly.

Note: The last load behavior option is not available for the Transactions per Second goal type.

After running each batch of Vusers, LoadRunner evaluates whether the target for the batch was achieved. If the batch target was not reached, LoadRunner recalculates the target number of hits, transactions, or pages per Vuser, and readjusts the number of Vusers for the next batch in order to be able to achieve the defined goal. Note that, by default, a new batch of Vusers is released every two minutes.

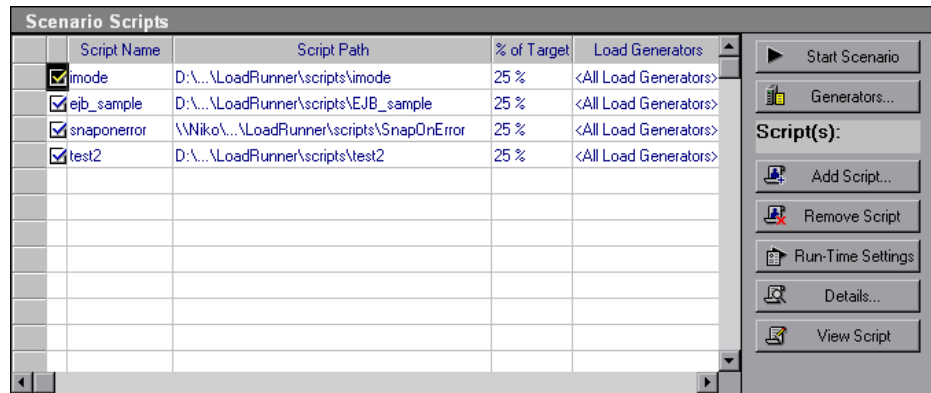
If the goal has not been reached once the Controller has launched the maximum number of Vusers, LoadRunner attempts to reach the defined target once more by recalculating the target number of hits, transactions, or pages per Vuser, and running the maximum number of Vusers simultaneously.

A Pages per Minute or Hits/Transactions per Second goal-oriented scenario is assigned a "Failed" status if:

- ▶ the Controller has twice attempted to reach the goal using the maximum number of Vusers specified, and the goal could not be reached.
- ▶ no pages per minute or hits/transactions per second were registered after the first batch of Vusers was run.
- ▶ the number of pages per minute or hits/transactions per second did not increase after the Controller ran a certain number of Vuser batches.
- ▶ all the Vusers that were run failed.
- ▶ there were no available load generators for the type of Vusers you attempted to run.

Assigning Properties to Scripts

The Scenario Scripts window displays a list of scripts you selected for the scenario.




The % of Target column indicates the percentage of the overall target number of Vusers, pages per minute, hits per second, transactions per second, or transaction response time that is automatically distributed to each Vuser script. The Load Generators column automatically contains <All Load Generators> for each Vuser script.

To modify the percentage of Vusers assigned to a script:

In the script's % of Target column, enter a percentage of the total target number of Vusers, pages per second, hits per second, transactions per second, or transaction response time you want LoadRunner to reach during the scenario. During the scenario, LoadRunner attempts to reach the percentage of the total that you stipulate for each script in the scenario.

To modify the load generator(s) for a script:

- 1 In the script's Load Generators column, select one or more machine(s) from the Load Generator Name list, and click **OK**. If you select multiple machines, the Vusers assigned to the script are distributed evenly among the load generators.
- 2 Alternatively, you can choose **Add** to add a load generator to the list. 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.

By default, LoadRunner stores temporary files on the load generator during scenario 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.

To allow the load generator to take part in the scenario, check **Enable load generator to take part in the scenario**.

Click **More** to expand the dialog box and show the Add Load Generator tabs. For information on configuring settings for each load generator, see "Configuring Load Generator Settings" on page 64.

Click **OK** to close the Add Load Generator dialog box. LoadRunner adds the new load generator to the Load Generator Name list. To include the new

load generator in your scenario, 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 scenario.

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. You can also use the Load Generators dialog box to indicate which load generators will run Vusers in the scenario. For example, if a load generator is unavailable for a particular scenario run, 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 61. To configure additional load generator settings, see "Configuring Load Generator Settings" on page 64.

To configure global settings for all load generators participating in the scenario, use LoadRunner's Options dialog box. For more information, see Chapter 10, "Configuring a Scenario."

Load Balancing

Load balancing evenly distributes the load generated by Vusers among the requested load generator machines, ensuring an accurate load test.

When a Windows load generator machine's CPU usage becomes overloaded, the Controller stops loading Vusers on the overloaded load generator, and automatically distributes them among load generators taking part in the scenario. Only where there are no other load generators in the scenario does the Controller stop loading Vusers.

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.

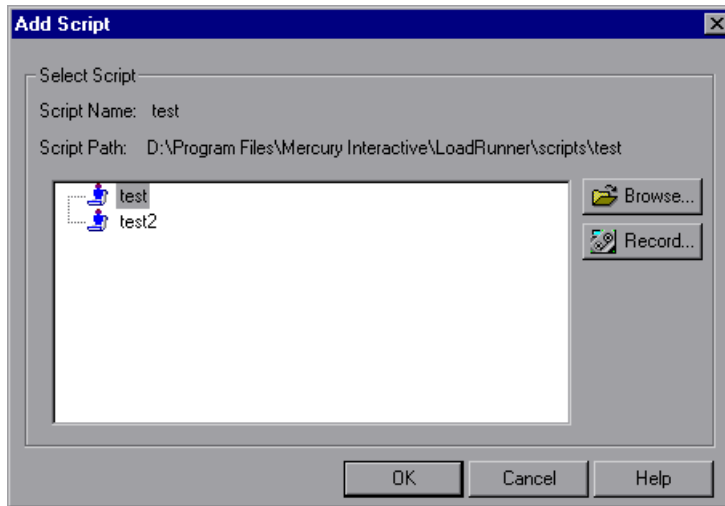
Note: Load balancing is only available in goal-oriented scenarios and manually controlled scenarios in the Percentage Mode.

Configuring Scripts

You can add a script to the Scenario Scripts list using the Add Script dialog box. Once you have added the script to the list, you can view the details of the script you selected, edit the script, or change its run-time settings.

To add a script:

- 1 Click the **Add Script** button to the right of the Scenario Scripts window, or right-click within a column and select **Add Script**. The Add Script dialog box opens.



- 2 Click the **Browse** button to the right of the path box. The Open Test dialog box opens.

Select the path and file name of the new script. To select a VB Vuser script, browse to locate the *.usr* file.

Note: When you specify the location of a script, you can specify a location that is relative to the current scenario directory. For details, see “Using Relative Paths for Scripts” on page 83.


- 3 Click **Open** to select the files. The Open Test dialog box closes, and the new script name appears in the Add Script dialog box.
 - 4 Click **OK** to close the Add Script dialog box and enter the new script information in the Scenario Scripts window.
-

Note: A script’s rendezvous points are disabled in a goal-oriented scenario.

To view script details:



- 1 Click the **Details** button to the right of the Scenario Scripts window, or right-click a script and select **Details**. The Script Information dialog box opens, displaying the Path, Name, and Type of the script you selected.



- 2** Click **Run-Time Settings** to set the script's run-time settings (optional), which allow you to customize the way the Controller executes a Vuser 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 Controller settings. Note that several protocols, such as Web and Java, have specific settings.

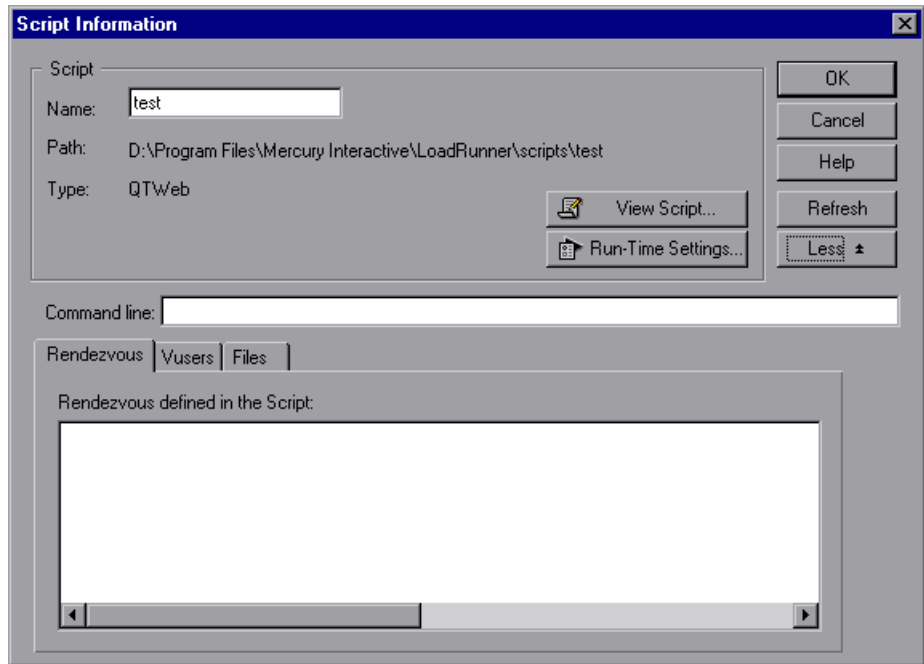
For information on configuring the run-time settings, refer to the *LoadRunner Creating Vuser Scripts User's Guide*.

Note: If you modify the run-time settings from the Controller, LoadRunner 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, see the *LoadRunner Creating Vuser Scripts User's Guide*.

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

- 4 Click **More** to expand the Script Information dialog box and view additional script information.



- 5 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 *LoadRunner Creating Vuser Scripts User's Guide*.
- 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 **OK** to close the Script Information dialog box.

To delete a script:



Click the **Remove Script** button to the right of the Scenario Scripts window, or right-click the script and select **Remove Script**.

To disable a script:

Click the box to the left of the Vuser script name. The color of the script entry changes to gray, indicating that the script will not take part in the scenario. To re-enable the Vuser script, click the same box again.

8

Scheduling a Scenario

After you create a scenario, you can set the time at which the scenario will begin running. In addition, for a manual scenario, you can set the duration time of the scenario or of the Vuser groups within the scenario. You can also select to gradually run and stop the Vusers within the scenario or within a Vuser group.

Note: The Vuser group settings are not applicable to the Percentage Mode.

This chapter describes:

- ▶ Delaying the Start of a Scenario
- ▶ Selecting a Schedule
- ▶ Scheduling a Scenario
- ▶ Scheduling Vuser Groups
- ▶ Adding Vusers to a Scheduled Scenario

About Scenario Scheduling

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

Using the Scenario Start dialog box, you can instruct LoadRunner to begin executing a scenario with a delay. You can specify either the number of minutes you want LoadRunner to wait from the time a *Run* command is issued, or the specific time at which you want the scenario to begin.

Using the Schedule Builder, you can set the timing aspect of a manual scenario, limiting the execution duration of the scenario or of a Vuser group within the scenario. You limit the execution time duration by specifying the number of minutes a scenario or Vuser group should be in the `RUNNING` state. When the scenario or group reaches its time limitation, it finishes.

Note: The Vuser group settings are not applicable to the Percentage Mode.

For manual scenarios, you can also stipulate how many Vusers LoadRunner starts and stops within a certain time frame. You specify whether LoadRunner should start or stop all Vusers in a scenario or Vuser group simultaneously, or start/stop only a certain number of Vusers within a specified amount of time.

The schedule you define is visually displayed in the Load Preview graph.

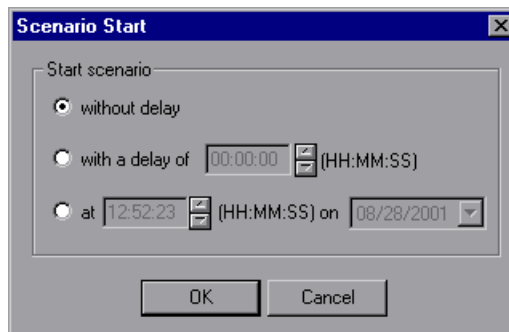
Note: Rendezvous points in a Vuser script interfere with a scheduled scenario. If your script contains rendezvous points, your scenario will not run as scheduled.

Delaying the Start of a Scenario

For both manual and goal-oriented scenarios, you can instruct LoadRunner to start running the scenario at a later point in time. You can specify either the number of minutes you want LoadRunner to wait from the time a *Run* command is issued, or the specific time at which you want the scenario to begin.

To delay the start of the scenario:

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



- 2 Select **with a delay of X (HH:MM:SS)** and enter the amount of time (in hours:minutes:seconds format) by which you want to delay the start of your scenario.

Alternatively, you can select **at X (HH:MM:SS) on X** and specify the time (in hours:minutes:seconds format) and date for the start of the scenario.

- 3 Click **OK** to close the dialog box and save your settings.


Selecting a Schedule

You select the schedule you want to use for your manual scenario from the Scenario Name box in the Scenario Schedule pane. You can select one of the existing schedules—Slow Ramp Up or Ramp Up—or New Schedule, if you want to create a schedule with new properties using the Schedule Builder.

Note that you can also change the properties for one of the three existing schedules using the Schedule Builder.

To create a new schedule:

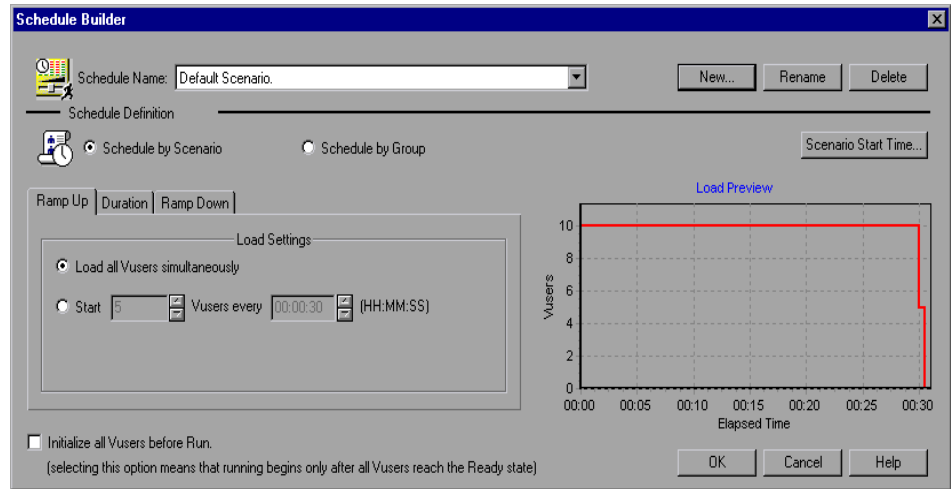
- 1 Select <new schedule> from the Scenario Name box in the Scenario Schedule pane. The New Schedule dialog box opens.



- 2 In the **Name** text box, type the name of the New Schedule and click **OK**. The Schedule Builder dialog box opens.

To modify the properties of an existing schedule:

- 1 Select **Slow Ramp Up** or **Ramp Up** from the Scenario Name box in the Scenario Schedule pane of the Design tab.
- 2 Select **Scenario > Schedule Builder**, or click the **Edit Schedule** button. The Schedule Builder dialog box opens.



To rename a schedule, click **Rename**. Enter the new name you want to use in the dialog box that opens. To delete a schedule, click **Delete**.

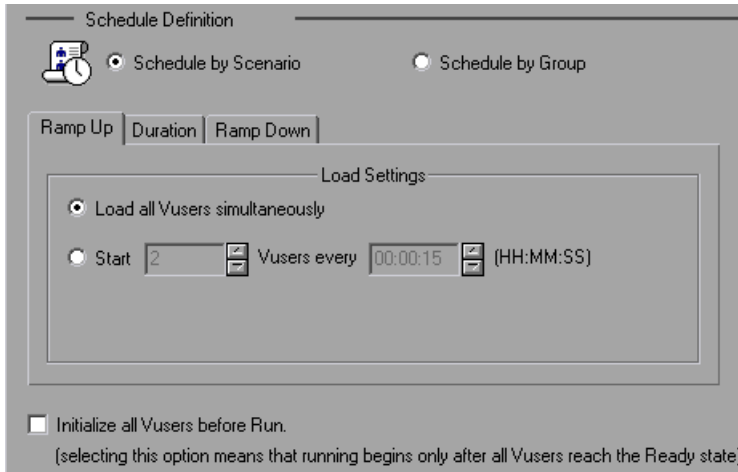
Scheduling a Scenario

Using the Schedule Builder, you can control the execution of your scenario by:

- limiting the scenario duration
- gradually running Vusers within a scenario
- gradually stopping Vusers within a scenario

To set the scheduling options for a scenario:

- 1** Select the **Schedule by Scenario** option.




- 2** To determine how to start the scenario, click the **Ramp Up** tab. Choose one of the following options:
 - **Load all Vusers simultaneously:** Starts all the Vusers in the scenario at once.
 - **Start X Vusers every X (HH:MM:SS):** Begins running the specified number of Vusers concurrently, and waits the specified time between Vuser ramp ups.

Note: While a scenario is running, you can add Vuser groups/scripts to a scenario and enable them. In the gradual ramp up mode, if you add a Vuser group/script after all the Vusers in the scenario have been ramped up, the new group/script will start loading immediately.

- 3** To initialize all Vusers before beginning to load them, select the **Initialize all Vusers before Run** check box. The running of Vusers begins only after all Vusers reach the READY state.

- 4 To set the duration of the scenario, click the **Duration** tab.



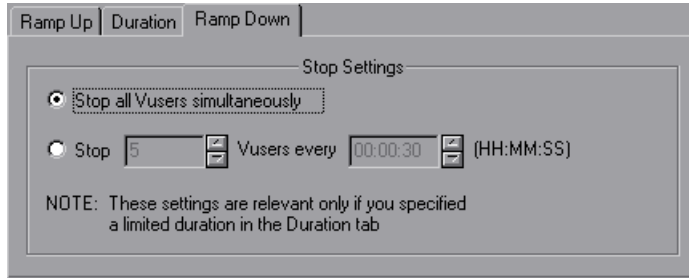
Choose one of the following options:

- **Run until completion**
- **Run for X (HHH:MM:SS) after the ramp up has been completed:** Runs the scenario for a specified amount of time, once all the Vusers have been ramped up.
- **Run indefinitely**

Note: The duration setting overrides the Vuser iteration settings. This means that if the duration is set to 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.

In a scenario of limited duration, the duration time begins to run from the start of the scenario. Vusers that take a long time to initialize may not reach the RUNNING state before the scenario ends. To ensure that all Vusers run in the scenario, check the **Initialize all Vusers before Run** check box.

- 5 To determine how to stop the scenario, click the **Ramp Down** tab.



Choose one of the following options:

- ▶ **Stop all Vusers simultaneously:** Stops all the Vusers in the scenario at once.
- ▶ **Stop X Vusers every X (HH:MM:SS):** Stops a certain number of Vusers within a specified time frame.

Note: The Ramp Down tab settings will only be applied if you select the second option in the Duration tab.

- 6 To instruct LoadRunner to initialize Vusers before beginning to load them, select **Initialize all Vusers before ramp up**. Note that LoadRunner will only begin to load the Vusers once they have all reached the `READY` state.
- 7 Click **OK** to close the Schedule Builder and save your settings.

Scheduling Vuser Groups

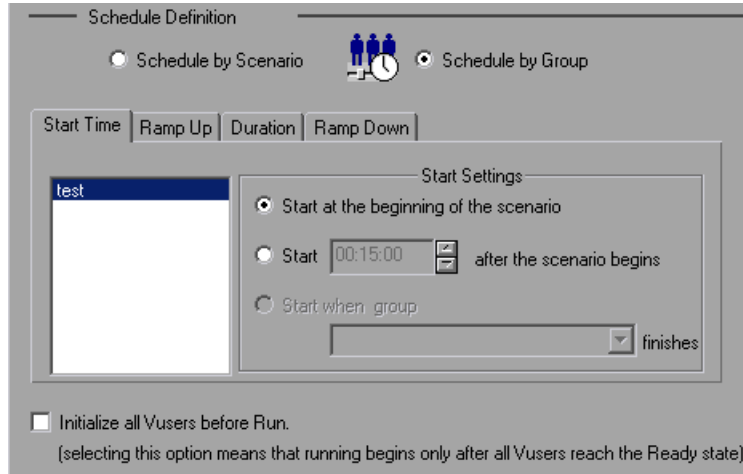
After you create a Vuser group, you can schedule the group's script execution by setting:

- ▶ the amount of time after the start of the scenario that the group must wait before it starts running
- ▶ the number of Vusers that will run within a specified period of time
- ▶ the number of Vusers that will be stopped within a specified period of time
- ▶ the amount of time the group will run

Note: The Vuser group settings are not applicable to the Percentage Mode.


To schedule a Vuser Group:

- 1** Select the **Schedule by Group** option.



- 2** Select a group from the box on the left.

- 3 To set the start time for the group, click the **Start Time** tab. Choose one of the following three options:
 - **Start the group at the beginning of the scenario**
 - **Start X after the scenario begins:** Waits the specified amount of time before running the group.
 - **Start when group X finishes:** Begins running the group after the specified group has finished running.
- 4 To set the ramp up for the group, click the **Ramp Up** tab.

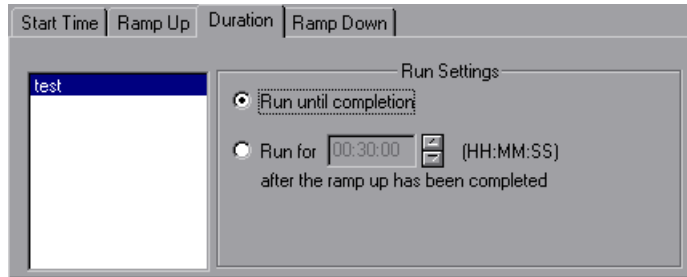


Choose one of the following options:

- **Load all of the Vusers simultaneously:** Starts all the Vusers in the scenario at once.
- **Start X Vusers every X (HH:MM:SS):** Begins running the specified number of Vusers concurrently, and waits the specified time between Vuser ramp ups.

Note: While a scenario is running, you can add Vuser groups to a scenario and enable them. In the gradual ramp up mode, if you add a Vuser group after all the Vusers in the scenario have been ramped up, the new group will start loading immediately.

- 5 To initialize all Vusers before beginning to load them, select the **Initialize all Vusers before Run** check box. The running of Vusers begins only after all Vusers reach the READY state.
- 6 To set the duration of the group, click the **Duration** tab.



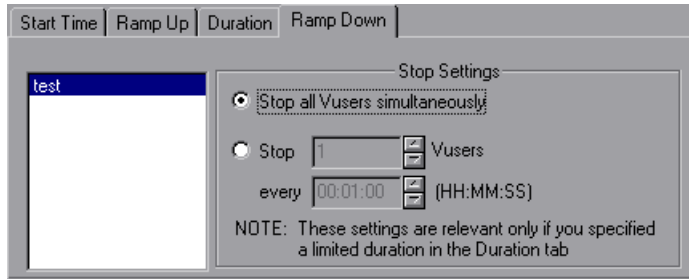
Choose one of the following options:

- **Run until completion**
- **Run for X (HH:MM:SS) after the ramp has been completed:** Runs the group for the specified amount of time, once all the Vusers have been ramped up.

Note: The duration setting overrides the Vuser iteration settings. This means that if the duration is set to 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.

In a scenario of limited duration, the duration time begins to run from the start of the scenario. Vusers that take a long time to initialize may not reach the RUNNING state before the scenario ends. To ensure that all Vusers run in the scenario, check the **Initialize all Vusers before Run** check box.

- 7 To determine how to stop the Vuser group, click the **Ramp Down** tab.



Choose one of the following options:

- ▶ **Stop all Vusers simultaneously:** Stops all the Vusers in the group at once.
- ▶ **Stop X Vusers every X (HH:MM:SS):** Stops a certain number of Vusers within a specified time frame.

Note: The Ramp Down tab settings will only be applied if you select the second option in the Duration tab.

- 8 To instruct LoadRunner to initialize Vusers before beginning to load them, select **Initialize all Vusers before ramp up**. Note that LoadRunner will only begin to load the Vusers once they have all reached the READY state.
- 9 Click **OK** to close the Schedule Builder and save your settings.

Adding Vusers to a Scheduled Scenario

If you are running a scenario or Vuser group using Schedule Builder settings, these settings will be applied to all Vusers that are manually added to the scenario or Vuser group during the scenario run. For example, if a running scenario or Vuser group has a set duration of five minutes, all Vusers subsequently added to the scenario or Vuser group will only run for the remaining part of this time period.

Vusers added to a scheduled scenario or Vuser group which has finished running will not be affected by Schedule Builder settings and will run according to the scenario run-time settings.

For more information on manually controlled Vusers, see “Manually Adding Vusers to a Running Scenario,” on page 176.

9

Using Rendezvous Points

LoadRunner allows you to check your system's response under specific load. To do this, you can use *rendezvous points* to cause multiple Vusers to perform tasks at exactly the same time, thereby creating intense user load on the server.

This chapter describes:

- Setting the Rendezvous Attributes
- Setting the Rendezvous Policy
- Disabling and Enabling Rendezvous Points
- Disabling and Enabling Vusers at Rendezvous Points
- Viewing Rendezvous Information

About Using Rendezvous Points

During a scenario run you can instruct multiple Vusers to perform tasks simultaneously by using rendezvous points. A rendezvous point creates intense user load on the server and enables LoadRunner to measure server performance under load.

Suppose you want to measure how a web-based banking system performs when ten Vusers simultaneously check account information. In order to emulate the required user load on the server, you instruct all the Vusers to check account information at exactly the same time.

You ensure that multiple Vusers act simultaneously by creating a *rendezvous point*. When a Vuser arrives at a rendezvous point, it is held there by the Controller. The Controller releases the Vusers from the rendezvous either when the required number of Vusers arrives, or when a specified amount of time has passed. For details on the release criteria, see “Setting the Rendezvous Policy,” on page 131.

You define rendezvous points in the Vuser script. For information about inserting rendezvous points into Vuser scripts, refer to the *LoadRunner Creating Vuser Scripts User's Guide*.

Using the Controller, you can influence the level of server load by selecting:

- ▶ which of the rendezvous points will be active during the scenario
- ▶ how many Vusers will take part in each rendezvous

For example, to test a bank server, you could create a scenario that contains two rendezvous points. The first rendezvous ensures that one thousand Vusers simultaneously deposit cash. The second rendezvous ensures that another thousand Vusers simultaneously withdraw cash. If you want to measure how the server performs when only five hundred Vusers deposit cash, you can deactivate (disable) the “withdraw” rendezvous, and instruct only five hundred Vusers to participate in the “deposit” rendezvous.

The following procedure outlines how to control load peaks on the server:

- 1 Create the Vuser scripts, inserting the necessary rendezvous points.**
- 2 Create a scenario.**

When you add a Vuser group to a scenario, LoadRunner scans the group's associated script for the names of the rendezvous points and adds them to the list in the Rendezvous Information dialog box (**Scenario > Rendezvous**). If you create another Vuser group that runs the same script, the Controller adds the new Vusers to the rendezvous and updates the list.

3 Set the level of emulated user load.

You determine the exact level of load by selecting the rendezvous points that will take part in the scenario, and how many Vusers will participate in each rendezvous.

4 Set the attributes for the rendezvous (optional).

For each rendezvous you can set *Policy* attributes. For more information, see “Setting the Rendezvous Policy,” on page 131.


5 Run the scenario.

Setting the Rendezvous Attributes

You can set the following rendezvous attributes from the Rendezvous Information dialog box (**Scenario > Rendezvous**):

- ▶ Rendezvous Policy
- ▶ Enabling and Disabling of Rendezvous Points
- ▶ Enabling and Disabling of Vusers

In addition, the dialog box displays general information about the rendezvous point: which script is associated with the rendezvous, and release history.



For information on manipulating the Vusers during scenario execution using the Release command, see Chapter 13, “Running a Scenario.”

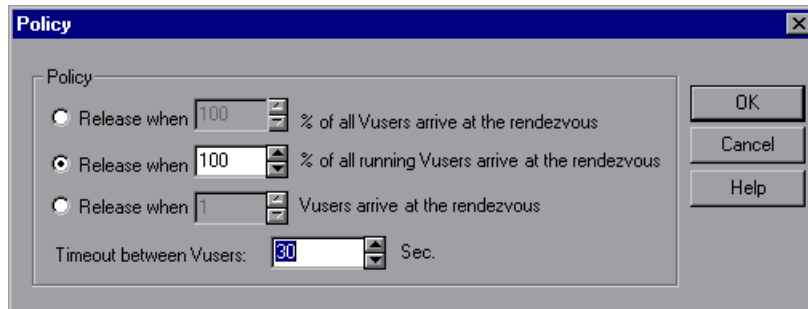
Setting the Rendezvous Policy

Setting the rendezvous policy determines how the Vusers handle a rendezvous point. You set the following policy attributes for each rendezvous:

- | | |
|-----------------------|--|
| release policy | sets how many Vusers will be released from a rendezvous at a time. |
| timeout | how long the Controller waits before releasing Vusers from a rendezvous. |

To set the rendezvous policy attributes:

- 1 Choose **Scenario > Rendezvous**. The Rendezvous Information dialog box opens.
- 2 Select a rendezvous from the Rendezvous box, and click the **Policy** button. The Policy dialog box opens.



3 In the Policy section, select one of the following three options:

- ▶ **Release when X% of all Vusers arrive at the rendezvous:** Releases the Vusers only when the specified percentage of all Vusers arrives at the rendezvous point.

Note: This option interferes with the scheduling of your scenario. If you select this option, therefore, your scenario will not run as scheduled.

- ▶ **Release when X% of all running Vusers arrive at the rendezvous:** Releases the Vusers only when the specified percentage of all Vusers running in the scenario arrives at the rendezvous point.

- ▶ **Release when X Vusers arrive at the rendezvous:** Releases the Vusers only when the specified number arrives at the rendezvous point.

4 Enter a timeout value in the Timeout between Vusers box. After each Vuser arrives at the rendezvous point, LoadRunner waits up to the maximum *timeout* period you set for the next Vuser to arrive. If the next Vuser does not arrive within the *timeout* period, the Controller releases all the Vusers from the rendezvous.

Each time a new Vuser arrives, the timer is reset to zero. The default *timeout* is thirty seconds.

5 Click **OK** to save your settings and close the Policy dialog box.

Disabling and Enabling Rendezvous Points

You can temporarily disable a rendezvous and exclude it from the scenario. By disabling and enabling a rendezvous, you influence the level of server load.

You use the Disable Rendezvous/Enable Rendezvous button in the Rendezvous Information dialog box, to change the status of a rendezvous.

To disable a rendezvous:

- 1 In the Rendezvous box, select the rendezvous you want to disable.
- 2 Click the **Disable Rendezvous** button. The button changes to **Enable Rendezvous** and the rendezvous becomes disabled.

To enable a rendezvous:

- 1 In the Rendezvous box, select the disabled rendezvous that you want to enable.
- 2 Click the **Enable Rendezvous** button. The button changes to **Disable Rendezvous** and the rendezvous becomes enabled.

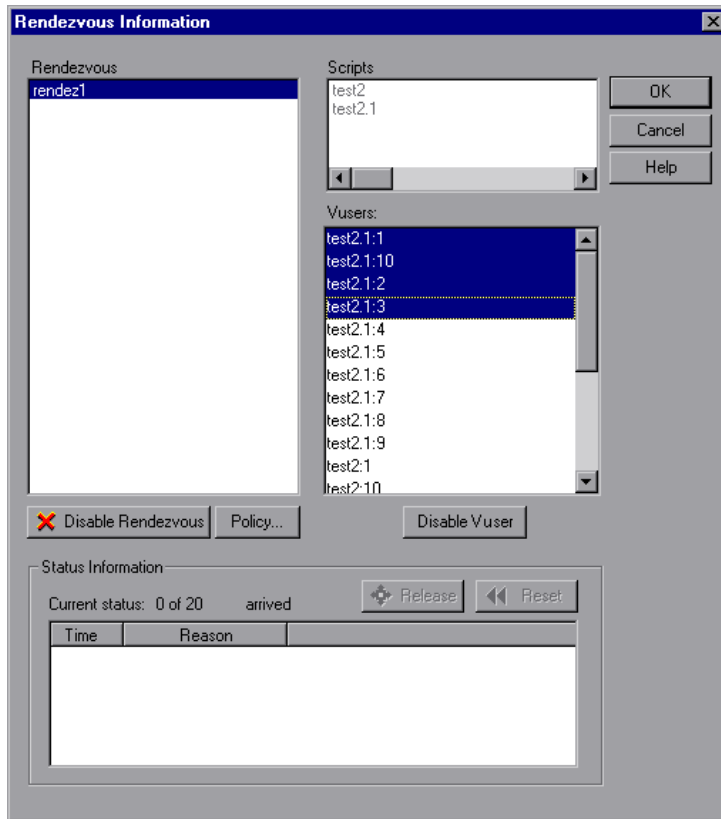
Disabling and Enabling Vusers at Rendezvous Points

In addition to disabling a rendezvous point for all Vusers in a scenario, LoadRunner lets you disable it for specific Vusers. By disabling Vusers at a rendezvous, you temporarily exclude them from participating in the rendezvous. Enabling disabled Vusers returns them to the rendezvous. You use the Disable and Enable commands to specify which Vusers will take part in a rendezvous.

To disable a Vuser in a rendezvous:

- 1 In the Rendezvous box, select the rendezvous for which you want to disable Vusers.

- 2 In the Vusers box, select the Vuser(s) you want to exclude from the rendezvous. Select multiple Vusers using the CTRL key.



- 3 Click the **Disable Vuser** button below the Vusers box. The disabled Vusers change from black to gray and will not take part in the rendezvous.
To enable a Vuser, select it and click **Enable Vuser**.

Viewing Rendezvous Information

During and after a scenario, you can view the rendezvous status in the Rendezvous Information dialog box. The following information is provided:

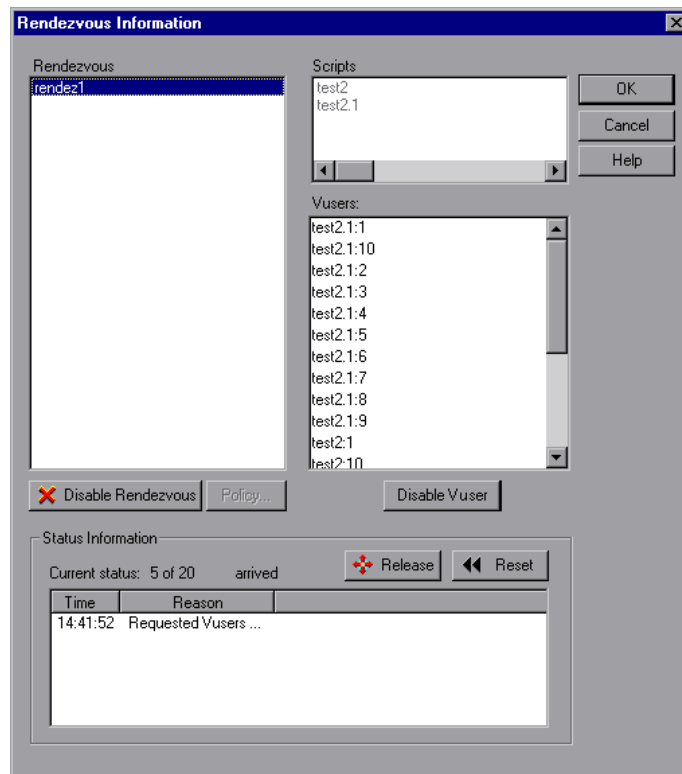
Time: The time at which the Vusers at the rendezvous point were released.

Reason: The reason the Vusers at the rendezvous point were released. The possible reasons are *Timeout* or *Arrived*.

Current Status: The number of Vusers that arrived at the rendezvous point, out of the total number of Vusers assigned to the rendezvous.

To view rendezvous information:

Select the rendezvous whose information you want to view. The rendezvous status is displayed in the Status Information section.



10

Configuring a Scenario

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

This chapter describes:

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

About Configuring a Scenario

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

The settings described in this chapter apply to all the load generators in a scenario. To change the settings for individual load generator machines, refer to Chapter 5, “Creating a Manual Scenario.” If the global scenario 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, see the *LoadRunner Creating Vuser Scripts User's Guide*.

For information on setting the options for online monitors, see Chapter 16, "Online Monitoring."

The LoadRunner Expert mode allows you to configure additional settings for the LoadRunner agent and other LoadRunner components. For more information, see Appendix C, "Working in Expert Mode."

Configuring Scenario Run-Time Settings

The scenario run-time settings relate to:

- Vuser Quotas
- Stopping Vusers
- Random Sequence Seed

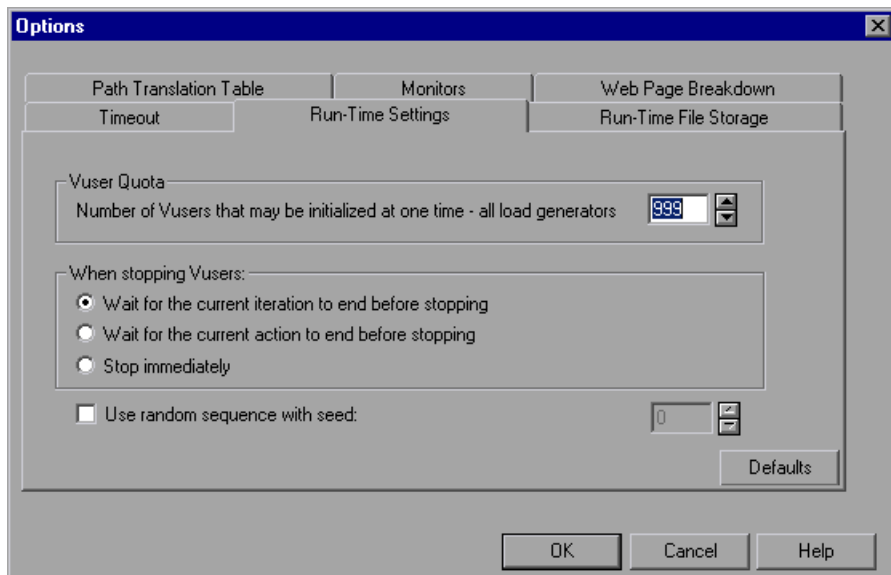
Vuser quotas: To prevent your system from overloading, you can set quotas for Vuser activity. The Vuser quotas apply to Vusers on all load generators. You can limit the number of Vusers initialized at one time (when you send an Initialize command).

Stopping Vusers: LoadRunner lets you control the way in which Vusers stop running when you click the Stop button. You can instruct LoadRunner 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: LoadRunner 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 scenario. This setting applies to parameterized Vuser 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 scenario 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 LoadRunner 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

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

The command timeouts are the maximum time limits for various LoadRunner commands. When a command is issued by the Controller, 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 Controller 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*.

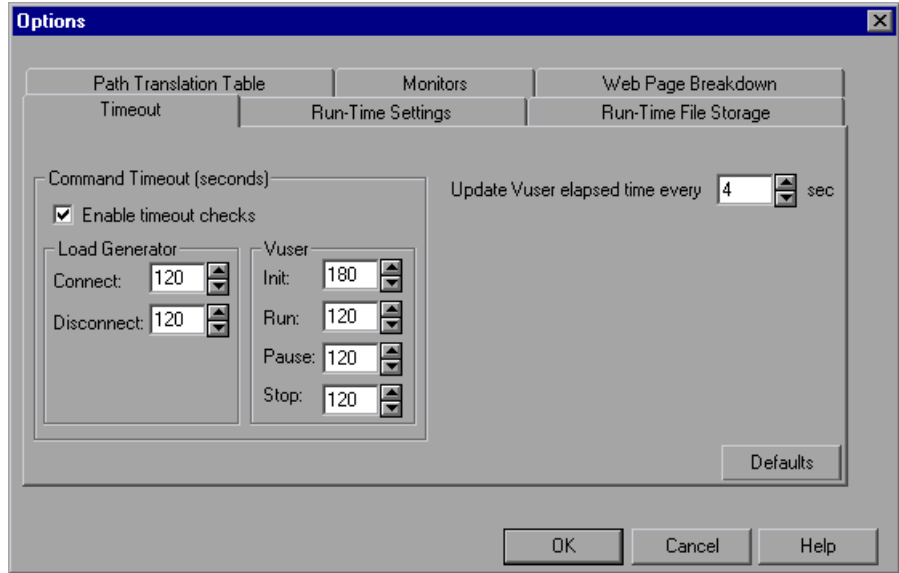
For example, the default *Init* timeout is 180 seconds. If you select a Vuser and click the **Initialize** button, LoadRunner checks whether the Vuser reaches the READY state within 180 seconds; if it does not, the Controller issues a message indicating that the *Init* command timed out.

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

Note: LoadRunner'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. LoadRunner 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. LoadRunner 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 LoadRunner updates the Elapsed time, in the **Update Vuser elapsed time every** box.

Setting the Run-Time File Location

When you run a scenario, by default the run-time files are stored locally on each Vuser load generator (the machine running the Vuser 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 chapter apply to all the load generators in a scenario. You can change the settings for individual load generator machines as described in “Configuring Load Generators” on page 61.

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

Script files:

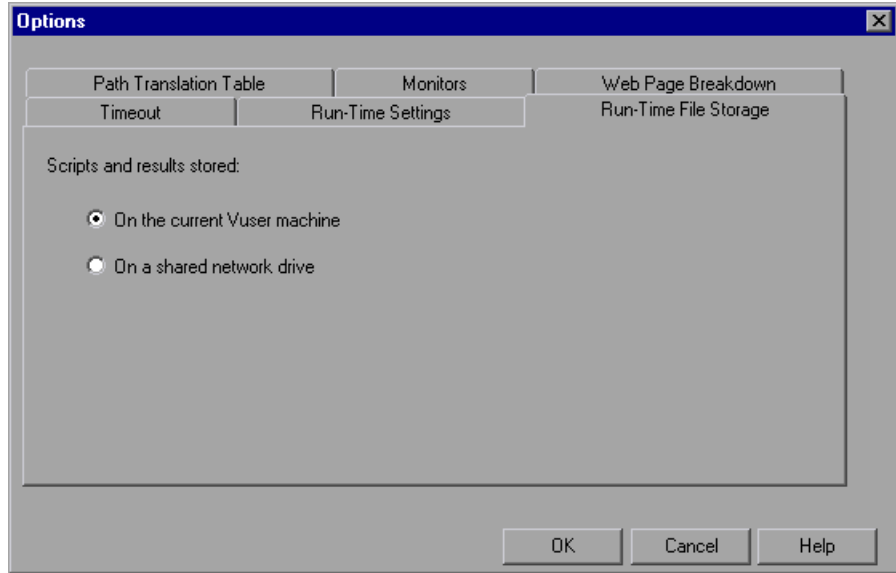
When you run a Vuser, the Controller sends a copy of the associated Vuser script to the Vuser load generator. The script is stored in the load generator's temporary run-time directory.

Result files:

While you run a scenario, the participating Vusers write their results to the temporary run-time file directory. After scenario 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 13, “Running a Scenario.” After collating the results, the temporary run-time directory is deleted.

To specify where LoadRunner 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 (Controller 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 11, “Preparing to Run a Scenario.”

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 B, “Performing Path Translation.”

If you specify that all Vusers access their Vuser 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 Vuser script files are large and therefore take a long time to transfer. Remember that Vuser script files are transferred only once during a scenario.

This alternate method often necessitates path translation. For details, see Appendix B, “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 LoadRunner 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 scenario 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 142), you may need to perform *path translation*. Path translation is a mechanism used by LoadRunner to convert a remote path names. A typical scenario may contain several load generator machines that map the shared network drive differently. For more information, see the Appendix B, “Performing Path Translation.”

11

Preparing to Run a Scenario

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

This chapter describes:

- ▶ Specifying a Results Location
- ▶ Results Directory File Structure
- ▶ Collating Results
- ▶ Setting Scenario Summary Information

About Preparing to Run a Scenario

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

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


For information on one-time configuration settings such as timeout, output, and quotas, see Chapter 10, “Configuring a Scenario.”

Specifying a Results Location

When you run a scenario, by default the run-time files are stored locally on each load generator. After the scenario, the results are collated together and processed on the Controller machine. Alternatively, you can instruct LoadRunner 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 10, “Configuring a Scenario.”

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

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

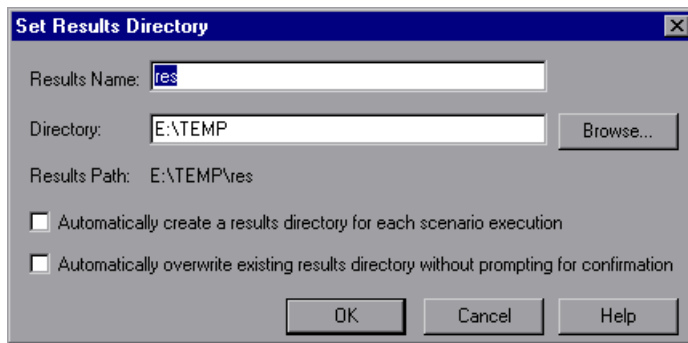


For more details on cross result graphs, see the *LoadRunner Analysis User's Guide*.

Note: You can also use Mercury Interactive’s Web-based test management program, TestDirector, to store results to a project. For information, see Chapter 12, “Managing Scenarios Using TestDirector.”

To specify where results are stored:

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



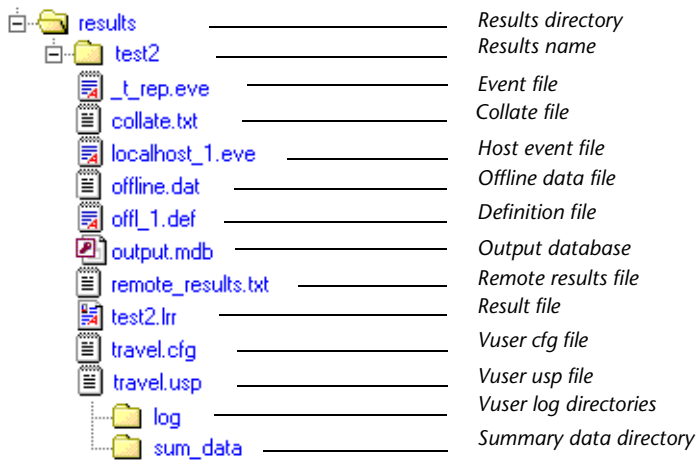
- 2 In the Results Name box, enter a name for the results. Avoid using the same name with different paths, since the names will appear identical on the graphs.
- 3 In the Directory box, type the full path of the results directory. If you are using the default file storage setting (local machine), specify a directory in which to store all of the collated results after the scenario run. If you specified a shared network drive as the file storage method, specify the directory to which Vuser groups should write during scenario execution.

Using the results name from step 2, the Controller creates a subdirectory within the results directory. All results are saved within this subdirectory.
- 4 Select the appropriate check box for subsequent executions: **Automatically create a results directory for each scenario execution** or **Automatically overwrite existing results directory without prompting for confirmation**.
- 5 Click **OK** to save the results directory setting.

Results Directory File Structure

When you set the results directory, you also specify a results name. LoadRunner 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 scenario in a result file (*.lrr*) and an event (*.eve*) file.

During scenario execution, LoadRunner creates a directory for each group in the scenario 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.
- *collate.txt* contains the file paths of the result files, and Analysis collation information.
- *local_host.eve* contains information from each agent host.
- *offline.dat* contains sample monitor information.

- **.def* are definition files for graphs that describe the online and other custom monitors.
- *output.mdb* is the database created by the Analysis (from the results files) that stores the output information.
- *remote_results.txt* contains the file paths for the host event files.
- *results_name.lrr* is the LoadRunner Analysis document file.
- **.cfg* files contain a listing of the script's run-time settings as defined in the VuGen application (think time, iterations, log, web).
- **.usp* files contain the script's run logic, including how the actions sections run.
- *Log* directory contains output information generated during replay for each Vuser. A separate directory exists for each Vuser group that runs in the scenario. Each group directory consists of Vusers subdirectories.
- *Sum data* directory. A directory containing the graph summary data (.dat) files.

When you generate analysis graphs and reports, the LoadRunner Analysis engine copies all of the scenario result files (*.eve* and *.lrr*) to a database. Once the database is created, the Analysis works directly with the database and does not use the result files.

For information on LoadRunner Analysis, see the *LoadRunner Analysis User's Guide*.

Collating Results

When you run a scenario, by default all Vuser information is stored locally on each load generator. After scenario 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 146.

Note: If you have selected to store all the scenario results directly to a shared network drive, then collation of the results is not required. See “About Configuring a Scenario,” on page 137 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 591.

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

If collation fails due to a lack of disk space, select **Results > Collate Results > Recollate**. LoadRunner attempts to collate the results again, without compressing the *.eve* file.

Before generating the analysis data, LoadRunner 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.

Setting Scenario Summary Information

The Controller allows you to provide a detailed description of the scenario. In addition, you can specify the author's name and a subject title for the scenario. Whenever you open this scenario, the summary information is available to you.

To set the scenario summary information:

- 1** Choose **Scenario > Summary Information**. The Summary Information box opens.



Summary Information

Scenario path: E:\Program Files\Mercury Interactive\LoadRunner\users\tutorial\lesson8.lrs

Author: john

Subject: Flight Reservation

Description:

This scenario creates 14 flight reservations ranging from 1 to 10 seats.
The script, reserve, is executed 10 times with 200 Vusers.

OK Cancel

- 2** In the Author box, enter the name of the author.
- 3** In the Subject box, enter a subject name or short title for the scenario.
- 4** In the Description box, enter a detailed description about the scenario.
- 5** Click **OK** to close the dialog box.

12

Managing Scenarios Using TestDirector

LoadRunner's integration with TestDirector lets you manage LoadRunner scenarios using TestDirector. TestDirector helps you organize and manage your scripts, scenarios, and results.

This chapter describes:

- ▶ Connecting to and Disconnecting from TestDirector
- ▶ Opening Scenarios from a TestDirector Project
- ▶ Saving Scenarios to a TestDirector Project
- ▶ Saving Results to a TestDirector Project
- ▶ Adding Vuser Scripts from a TestDirector Project

About Managing Scenarios Using TestDirector

LoadRunner works together with TestDirector, Mercury Interactive's Web-based test management tool. TestDirector provides an efficient method for storing and retrieving scenarios and collecting results. You store scenarios and results in a TestDirector project and organize them into unique groups.

In order for LoadRunner to access a TestDirector project, you must connect it to the Web server on which TestDirector is installed. You can connect to either a local or remote Web server.

For more information on working with TestDirector, refer to the *TestDirector User's Guide*.

Connecting to and Disconnecting from TestDirector

If you are working with both LoadRunner and TestDirector, LoadRunner can communicate with your TestDirector project. You can connect or disconnect LoadRunner from a TestDirector project at any time during the testing process.


Connecting LoadRunner to TestDirector

The connection process has two stages. First, you connect LoadRunner to a local or remote TestDirector Web server. This server handles the connections between LoadRunner and the TestDirector project.

Next, you choose the project you want LoadRunner to access. The project stores scenarios and results for the application you are testing. Note that TestDirector projects are password protected, so you must provide a user name and a password.

To connect LoadRunner to TestDirector:

- 1 In the Controller, choose **Tools > TestDirector Connection**. The TestDirector Connection dialog box opens.



- 2 In the Server box, type the URL address of the Web server on which TestDirector is installed.

Note: You can choose a Web server accessible via a Local Area Network (LAN) or a Wide Area Network (WAN).

- 3 Click **Connect**. Once the connection to the server is established, the server's name is displayed in read-only format in the Server box.
- 4 From the Project box in the Project Connection section, select a TestDirector project.
- 5 In the User Name box, type a user name.
- 6 In the Password box, type a password.
- 7 Click **Connect** to connect LoadRunner to the selected project.


Once the connection to the selected project is established, the project's name is displayed in read-only format in the Project box.

- 8 To automatically reconnect to the TestDirector server and the selected project on startup, select the **Reconnect on startup** check box.
- 9 If you select **Reconnect on startup**, you can save the specified password to reconnect on startup. Select the **Save password for reconnection on startup** check box.

If you do not save your password, you will be prompted to enter it when LoadRunner connects to TestDirector on startup.

- 10 Click **Close** to close the TestDirector Connection dialog box.

The status bar indicates that LoadRunner is currently connected to a TestDirector project.




Disconnecting LoadRunner from TestDirector

You can disconnect LoadRunner from a selected TestDirector project and Web server.

To disconnect LoadRunner from TestDirector:

- 1 In the Controller, choose **Tools > TestDirector Connection**. The TestDirector Connection dialog box opens.



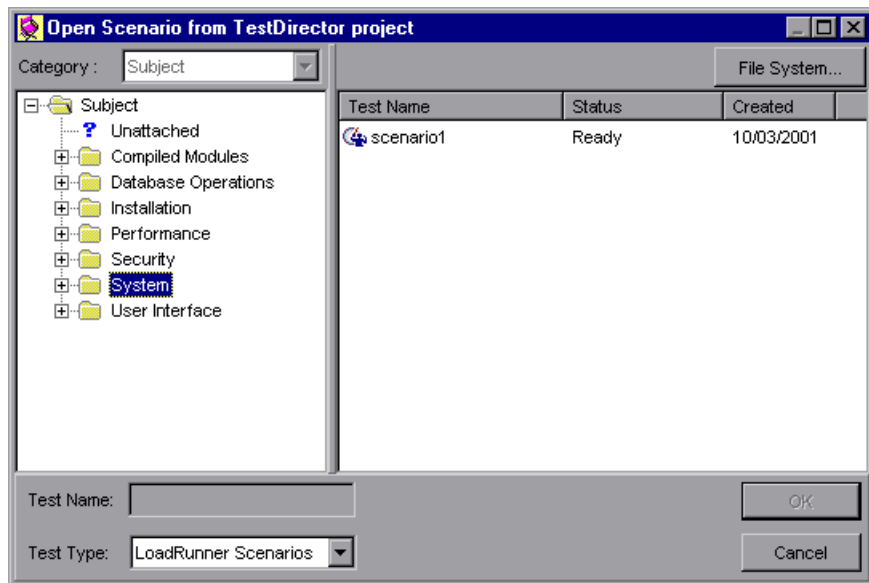
- 2 To disconnect LoadRunner from the selected project, click **Disconnect** in the Project Connection section.
- 3 To disconnect LoadRunner from the selected server, click **Disconnect** in the Server Connection section.
- 4 Click **Close** to close the TestDirector Connection dialog box.

Opening Scenarios from a TestDirector Project

When LoadRunner is connected to a TestDirector project, you can open your scenarios from TestDirector. You locate tests according to their position in the test plan tree, rather than by their actual location in the file system.

To open a scenario from a TestDirector project:

- 1** Connect to the TestDirector server (see “Connecting LoadRunner to TestDirector” on page 154).
- 2** In the Controller, choose **File > Open** or click the **File Open** button. The Open Scenario from TestDirector Project dialog box opens and displays the test plan tree.



To open a scenario directly from the file system, click the **File System** button. The Open Scenario dialog box opens. (From the Open Scenario dialog box, you may return to the Open Scenario from TestDirector Project dialog box by clicking the TestDirector button.)

- 3 Click the relevant subject in the test plan tree. To expand the tree and view sublevels, double-click closed folders. To collapse the tree, double-click open folders.

Note that when you select a subject, the scenarios that belong to the subject appear in the Test Name list.

- 4 Select a scenario from the Test Name list. The scenario appears in the read-only Test Name box.
- 5 Click **OK** to open the scenario. LoadRunner loads the scenario. The name of the scenario appears in the Controller's title bar. The Design tab shows the scripts, the load generators, and the Vusers and Vuser groups in the scenario.

Note: You can also open scenarios from the recent scenarios list in the File menu. If you select a scenario located in a TestDirector project, but LoadRunner is currently not connected to that project, the TestDirector Connection dialog box opens. Enter your user name and password to log in to the project, and click **OK**.

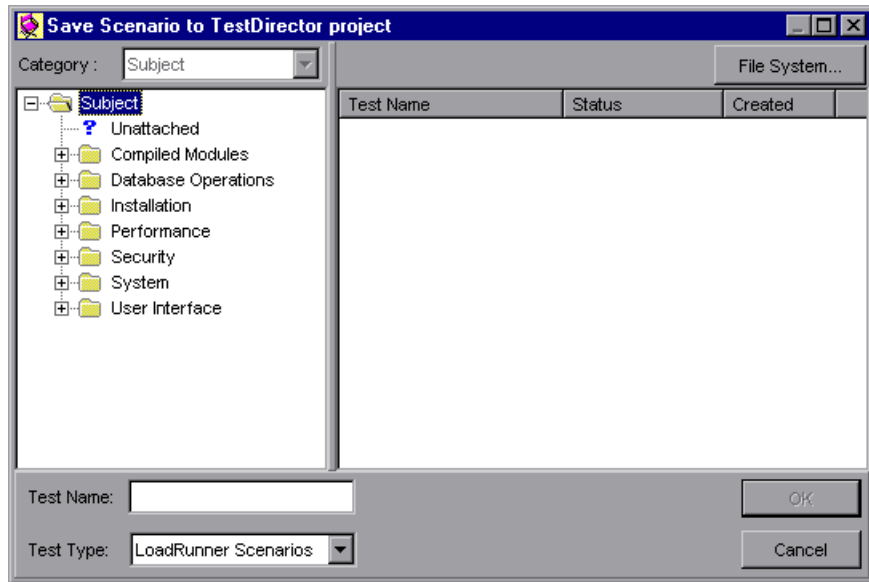
Saving Scenarios to a TestDirector Project

When LoadRunner is connected to a TestDirector project, you can create new scenarios in LoadRunner and save them directly to your project. To save a scenario, you give it a descriptive name and associate it with the relevant subject in the test plan tree. This helps you to keep track of the scenarios created for each subject and to quickly view the progress of test planning and creation.

To save a scenario to a TestDirector project:

- 1 Connect to the TestDirector server (see "Connecting LoadRunner to TestDirector" on page 154).

- 2** In the Controller, choose **File > Save As**. The Save Scenario to TestDirector Project dialog box opens and displays the test plan tree.



To save a scenario directly in the file system, click the **File System** button. The Save Scenario dialog box opens. (From the Save Scenario dialog box, you may return to the Save Scenario to TestDirector Project dialog box by clicking the TestDirector button.)

- 3** Select the relevant subject in the test plan tree. To expand the tree and view a sublevel, double-click a closed folder. To collapse a sublevel, double-click an open folder.
- 4** In the Test Name box, enter a name for the scenario. Use a descriptive name that will help you easily identify the scenario.
- 5** Click **OK** to save the scenario and close the dialog box.

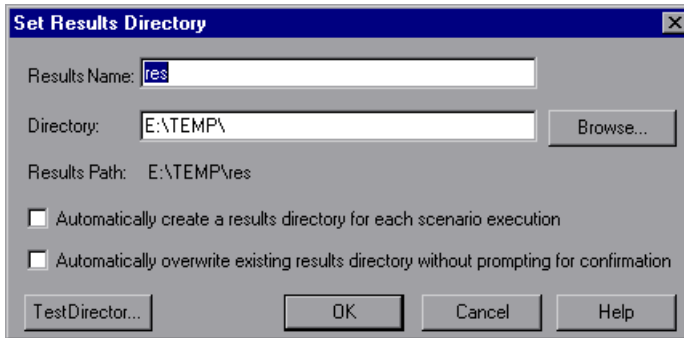
The next time you start TestDirector, the new scenario will appear in TestDirector's test plan tree.

Saving Results to a TestDirector Project

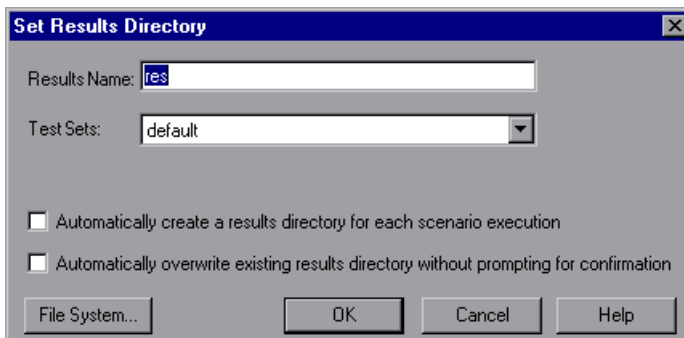
Before you run a scenario, you set the results location. When LoadRunner is connected to a TestDirector project, results are saved to a test set. You can also save the results to disk using the standard file system.

To save results to a TestDirector project:

- 1 Connect to the TestDirector server (see “Connecting LoadRunner to TestDirector” on page 154).
- 2 In the Controller, choose **Results > Results Settings**. The Set Results Directory dialog box opens.



- 3 Click **TestDirector**. The Directory box changes to Test Sets.



- 4 In the Results Name box, enter a name for the results.
- 5 In the Test Sets list, accept the default test set name or select a different name.

- 6 Select the appropriate check box:
 - **Automatically create a results directory for each scenario execution:** Instructs LoadRunner to create a unique results directory for each scenario execution. By default, the result names are res1, res2, res3, etc.
 - **Automatically overwrite existing results directory without prompting for confirmation:** Instructs LoadRunner to automatically overwrite previous result sets, without prompting the user.
- 7 Click **OK** to save the results settings.

Adding Vuser Scripts from a TestDirector Project

You can add Vuser scripts from a TestDirector project to the Controller's script list. You can add the script to a manual or goal-oriented scenario.

Adding a Vuser Script to a Manual Scenario

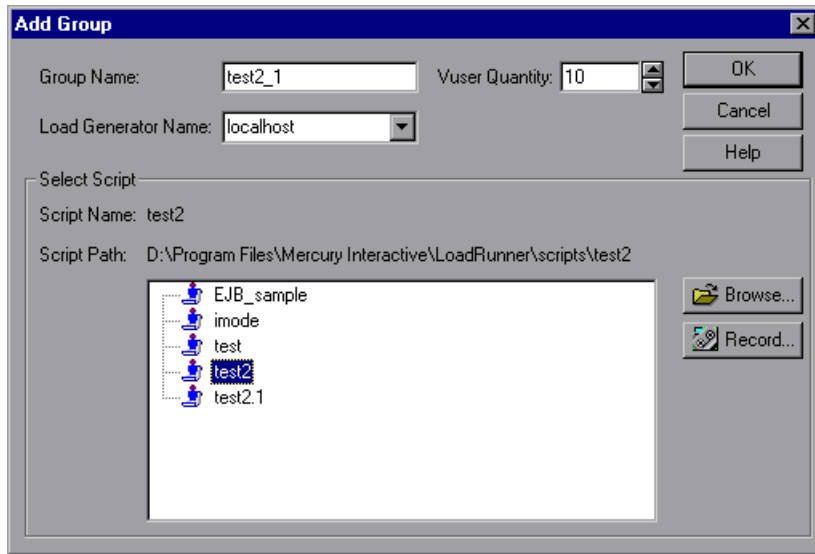
When you create a manual scenario, you use the Add Group dialog box to add Vuser scripts.

To add a Vuser script to a manual scenario:

- 1 Connect to the TestDirector server (see "Connecting LoadRunner to TestDirector," on page 154).



- 2 In the Scenario Groups window, click the **Add Group** button. The Add Group dialog box opens.




- 3 Click the **Browse** button. The Open Test from TestDirector Project dialog box opens and displays the test plan tree.
- 4 Select the script and click **OK**. The Script Path field displays [TD], the full subject path, and the script name. For example:
[TD]\Subject\System\test_td
- 5 Click **OK** to close the Add Group dialog box. The script is displayed in the Scenario Groups pane.

Adding a Vuser Script to a Goal-Oriented Scenario


When you create a goal-oriented scenario, you use the Add Script dialog box to add scripts.

To add a Vuser script to a goal-oriented scenario:

- 1 Connect to the TestDirector server (see “Connecting LoadRunner to TestDirector,” on page 154).



- 2 In the Scenario Scripts window, click the **Add Script** button. The Add Script dialog box opens.



- 3 Click the **Browse** button. The Open Test from TestDirector Project dialog box opens and displays the test plan tree opens.
- 4 Select the script and click **OK**. The Script Path field displays [TD], the full subject path, and the script name. For example:
[TD]\Subject\System\test_td
- 5 Click **OK** to close the Add Script dialog box. The script appears in the Script Path column in the Scenario Scripts window.

Part III

Executing a Scenario

13

Running a Scenario

When you run a scenario, LoadRunner generates load on the application you are testing, and measures the system's performance.

This chapter describes:

- ▶ Running an Entire Scenario
- ▶ Controlling Vuser Groups
- ▶ Controlling Individual Vusers
- ▶ Manually Releasing Vusers from a Rendezvous
- ▶ Manually Adding Vusers to a Running Scenario

About Running a Scenario

When you run a scenario, the Vuser groups are assigned to their load generators and execute their Vuser scripts. During scenario execution, LoadRunner:

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

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

While the scenario 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 LoadRunner 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 Scenario Run-Time Settings” on page 138.

Note: When automatically stopping Vusers in a goal-oriented scenario, LoadRunner stops running the Vusers immediately.

You can also activate additional Vusers while the scenario is running, using the Run/Stop Vusers dialog box. For more information, see “Manually Adding Vusers to a Running Scenario” on page 176.

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

The following procedure outlines how to run a scenario:

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


Running an Entire Scenario

You can run all the Vusers and Vuser groups in a scenario, or you can select the specific Vuser groups and Vusers that you want to run. Note that when you run an entire scenario, LoadRunner does not begin running Vusers until all of them have reached the `READY` state. However, if you run individual groups or Vusers, LoadRunner runs the Vusers as soon as they reach the `READY` state.

The following section describes how to run an entire scenario. “Controlling Vuser Groups,” on page 170 and “Controlling Individual Vusers,” on page 173 describe how to manipulate Vuser groups and individual Vusers.

To run an entire scenario:

- 1 Open an existing scenario or create a new one. Click the **Run** tab. The Scenario Groups pane appears in the top left-hand corner of the screen.
- 2 Choose **Scenario > Start**, or click the **Start Scenario** button. The Controller starts initializing the Vusers and distributing them to their designated load generators—where they begin to execute their Vuser scripts.




Note: The Controller begins running the scenario according to the start time set in the Scenario Start dialog box.

If you have not specified a results directory for the scenario, the Set Results Directory dialog box opens.


During scenario execution you can manipulate individual Vusers and Vuser groups. This is described in “Controlling Vuser Groups,” on page 170 and “Controlling Individual Vusers,” on page 173.

- 3 Select **Scenario > Stop/Resume Ramp Up** to stop the ramp up process. Select it again to resume the ramping up of Vusers.
- 4 Select **Scenario > Stop/Resume Ramp Down** to stop the ramp down process. Select it again to resume the ramping down of Vusers.



- 5 Choose **Scenario > Stop**, or click the **Stop** button to terminate the scenario. If you selected the **Exit immediately** option in the Run-Time Settings tab of the Options dialog box, all of the Vusers in the scenario 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 scenario move to the GRADUAL EXITING status and exit the scenario gradually. To stop the Vusers immediately, click **Stop Now**.



- 6 Choose **Scenario > Reset**, or click the **Reset** button to reset all Vusers to their pre-scenario, DOWN status.


Controlling Vuser Groups

You can run an entire scenario as described above, or you can manipulate individual Vuser groups in the scenario. This section describes how to initialize, run, and stop Vuser groups.

Initializing Vuser Groups

Initializing a Vuser group distributes the Vusers in the group to their designated load generators so that they are ready to execute their script(s). By initializing all of the Vusers in a group before running them, you can ensure that they all begin executing the scenario at the same time.


To initialize a Vuser group:

- 1 Select the Vuser group or groups that you want to initialize.
- 2 Click the **Initialize Vusers** button  , or right-click the Vuser group or groups that you want to initialize and select **Initialize Group/s**. The Vuser group's status changes from DOWN to PENDING to INITIALIZING to READY. If a Vuser group fails to initialize, the Vuser group status changes to ERROR.

Running Vuser Groups

Running a Vuser group tells the Vuser group to execute its script.

To run a Vuser group:

- 1 Select the Vuser group or groups that you want to run.
- 2  Click the **Run Vusers** button, or right-click the Vuser group or groups that you want to run and select **Run Group/s**. The Vuser groups execute their scripts. If you run a Vuser group in the DOWN or ERROR state, LoadRunner initializes and then runs the Vuser group.

Note: You can instruct LoadRunner to randomly run only one Vuser in a group by right-clicking the Vuser group and selecting **Run One Vuser**. A Vuser script log opens, displaying run-time information about the Vuser. For more information about the Vuser log, see “Viewing the Vuser Script Log” on page 190.

Pausing Vuser Groups

Pausing a Vuser group temporarily stops script execution. The Pause command changes the status of a Vuser group from RUNNING to PAUSED.

Note: Pausing a Vuser group will affect its transaction response time.


To pause a Vuser:

- 1 Select the Vuser group or groups that you want to pause.
- 2 Select **Pause** from the right-click menu. The Vusers groups temporarily stop script execution.

Stopping Vuser Groups

Stopping a Vuser group stops script execution. If you stop a Vuser group, the group still appears in the Vuser group list.

To stop a Vuser group:

- 1 Select the Vuser group or groups that you want to stop.
- 2  Click the **Stop Vusers** button, or right-click the Vuser group or groups and select **Stop**. The Vuser groups stop executing their scripts immediately.

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 group in the RUN state, click the Gradual Stop button, or right-click the Vuser group and select **Gradual Stop**. The Vusers in the group move to the GRADUAL EXITING status and exit the scenario gradually.



Note: If the Vusers are not in the RUN state, the Gradual Stop option is disabled.

Resetting Vuser Groups

Resetting causes all of the Vusers in a group to revert to their pre-scenario, DOWN status.

To reset a Vuser group:

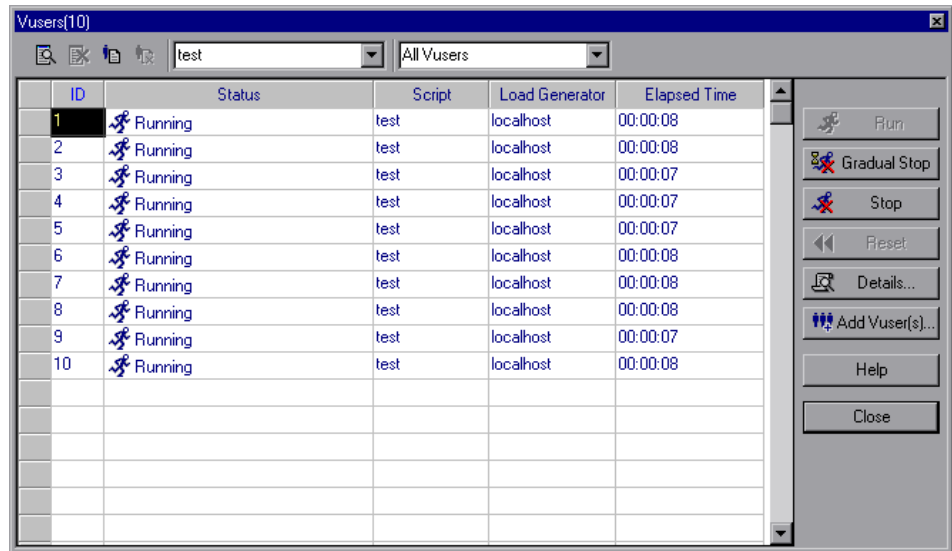
- 1 Select the Vuser group or groups that you want to stop.
- 2 Right-click the Vuser group or groups that you want to stop, and select **Reset**. The Vuser groups revert to their pre-scenario, DOWN status.

Controlling Individual Vusers

You can also manipulate individual Vusers within the Vuser groups you have defined. This section describes how to initialize, run, and stop individual Vusers.

To control an individual Vuser:

- 1 Select a Vuser group and click the **Vusers** button. The Vusers dialog box opens, with a list of the ID, Status, Script, Load Generator, and Elapsed Time (since the beginning of the scenario) for each of the Vusers in the group.



You can control an individual Vuser using the following utilities:

- Select a Vuser and click **Run** to run it.
- Select a Vuser and click **Stop** to stop it immediately from running.

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 the **Gradual Stop** button. The Vuser moves to the GRADUAL EXITING status and exits the scenario gradually.

- To pause a Vuser, right-click it and select **Pause**.

Note: Pausing a Vuser will affect its transaction response time.

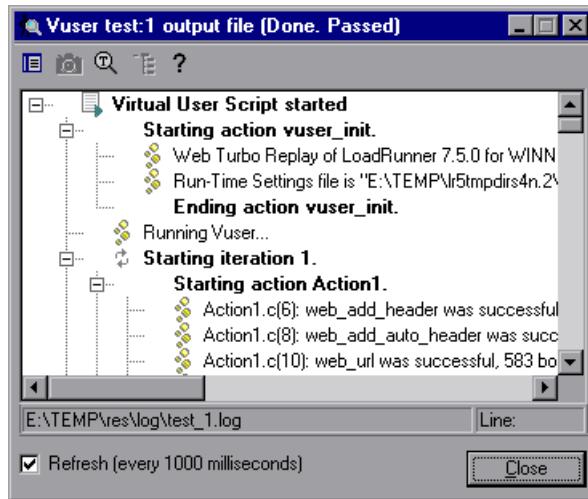
- Select a Vuser and click **Reset** to revert its status to DOWN.
- To initialize a Vuser, right-click it and select **Initialize Vuser/s**.
- To renumber the Vusers in a group, right-click the Vusers you want to renumber and select **Renumber**.
- To filter the Vusers listed, right-click in one of 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 Vusers listed, 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** button. The Run-Time Viewer opens, allowing you to see the Vuser executing the script.



To close the Run-Time Viewer, click the **Hide** button.



- To view the Vuser script log, click the **Vuser log** button. A script log, such as the following, appears.



To close the Vuser script log, click the **Close** button. For more information on the Vuser script log, see page 190.

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

Manually Releasing Vusers from a Rendezvous

While you run a scenario, you can manually release Vusers from a rendezvous before the Controller releases them.

To manually release Vusers from a rendezvous:

- 1 Choose **Scenario > Rendezvous**. The Rendezvous Information dialog box opens.
- 2 Select a rendezvous from the Rendezvous list.
- 3 Click **Release**. The Vusers in the rendezvous are released.




Manually Adding Vusers to a Running Scenario

During a running scenario, you can manually control the addition of new Vusers using the Run/Stop Vusers dialog box. The dialog box differs depending on the scenario mode you are running:

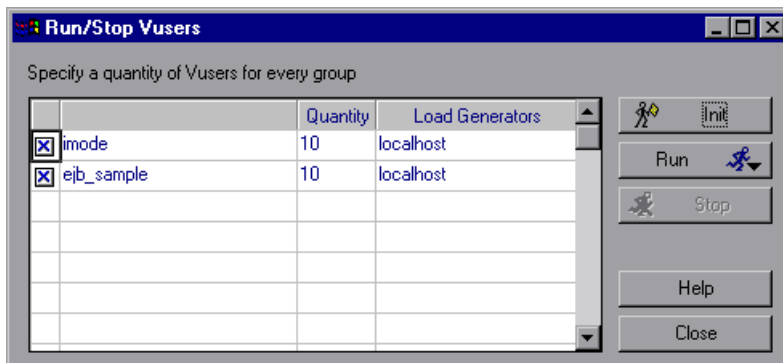
- ▶ If you are running a scenario in the Vuser Group Mode, you control the number of new Vusers that can be added to each Vuser Group, and the load generators on which these additional Vusers will run.
- ▶ If you are running a scenario in the Percentage Mode, you control the number of new Vusers that can be distributed among the Vuser scripts according to the percentage you define, and the load generators on which these additional Vusers will run.

Note: If you are running a scenario or Vuser group using Schedule Builder settings, these settings will be applied to all Vusers that are manually added to the scenario or Vuser group during the scenario run. For more information, see “Adding Vusers to a Scheduled Scenario,” on page 125.

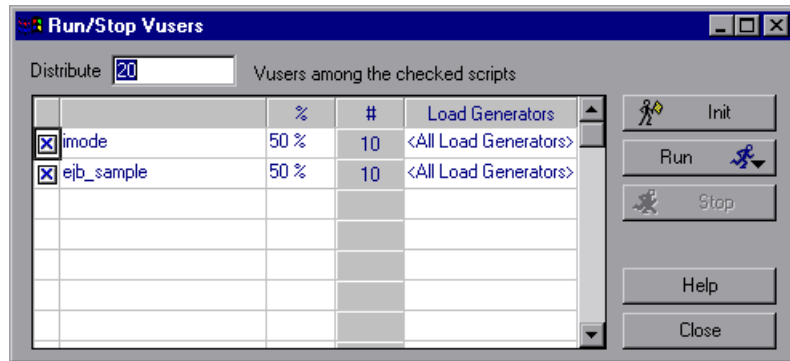
To add Vusers to a running scenario:



- 1 Select **Scenario > Run/Stop Vusers**, or click the **Run/Stop Vusers** button in the Scenario Groups pane of the Run view. The Run/Stop Vusers dialog box opens. If you are in the Vuser Group Mode, the dialog box displays the Vuser groups included in the scenario.



If you are in the Percentage Mode, the Run/Stop Vusers dialog box displays the Vuser scripts included in the scenario.



- 2 If you are in the Vuser Group Mode, enter the number of Vusers you want to run for each group in the Quantity column.

If you are in the Percentage Mode, enter the number of Vusers and the percentage in which you want these Vusers to be distributed among the checked Vuser scripts. LoadRunner automatically distributes the number of Vusers you entered. The % column indicates the percentage of Vusers distributed to each Vuser script. The # column indicates the number of Vusers distributed to each Vuser script.

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

Note: If you disable a Vuser group in the Vuser Group Mode, no Vusers will be distributed to it. If you disable a Vuser script in the Percentage Mode, no Vusers will be distributed to it, and the unused percentage of the Vusers will not be distributed among the remaining scripts, unless you define a zero percent value for the disabled script.

- 4 To change the load generator on which the Vuser group/script will run, select a different load generator from the Load Generator column.

To use a load generator that does not appear, select **Add** from the Load Generator Name list and add a new load generator using the Add Load Generator dialog box.

If you are in the Percentage Mode, you can select more than one load generator to run the Vuser script. From the Load Generator Name list, select the load generator(s) and click **OK**. To use all the load generators in the list, click the **All Generators** button.

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

- 5 Click the **Init** button to initialize the number of Vusers you added.

The Controller first initializes the Vusers in your scenario 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 two options:

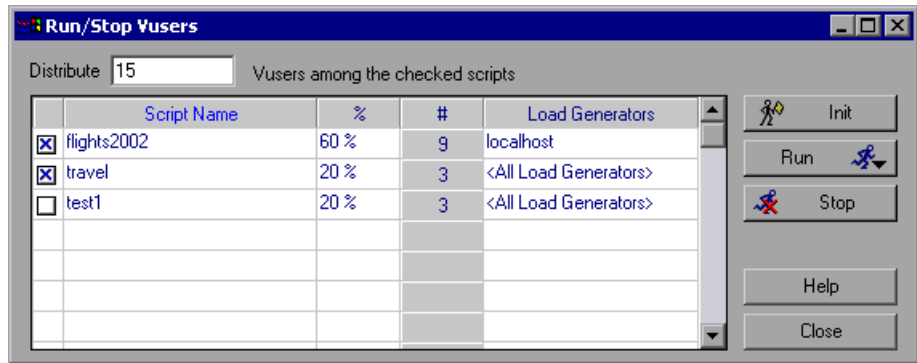
- ▶ **Run Initialized:** Runs the Vusers in the scenario that have already been initialized on the load generators defined in the Run/Stop Vusers dialog box. The Controller runs only those Vusers that have already been initialized, regardless of their quantity.
- ▶ **Run New:** Runs the number of Vusers you specified. The Controller first runs the Vusers in your scenario 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 Controller 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.

Example of a Manually Controlled Scenario

The example below shows the Run/Stop Vusers dialog box from a scenario running in the Percentage Mode.



The number of Vusers that are distributed among the checked scripts is 15. The % column indicates that 60% of these Vusers are distributed to the script *flights2002*, and 20% to both *travel* and *test1*. In accordance with these percentages, the # column indicates that nine Vusers are distributed to *flights2002* and three to *travel* and *test1*.

Note: The unused percentage of the Vusers from the disabled *test1* script are not distributed among the remaining scripts, because a percentage value has been defined for this script.

When an action (**Init**, **Run**, **Stop**) is selected from the Run/Stop Vusers dialog box, the Controller runs only the number of Vusers that appear in the # column. In this example, nine Vuser are initialized, run, or stopped in the *flights2002* script, and three in the *travel* script.

All Vusers distributed to the *flights2002* script are run on the *localhost* load generator. For the *travel* script, Vusers are proportionally distributed among all defined load generators.

Note: Load generator balancing is applied to a manually controlled scenario in the Percentage Mode where there are other load generators assigned to Vuser scripts. For more information, see “Load Balancing,” on page 107.

14

Viewing Vusers During Execution

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

This chapter describes:

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

About Viewing Vusers During Execution

LoadRunner lets you view Vuser activity during a scenario:

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





Monitoring Vuser Status

During scenario execution, you can use the Scenario Groups pane in the Run view to monitor the actions of all the Vusers and Vuser groups in the scenario.

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 scenario.

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 Vuser script is being executed on a load generator.
RENDEZVOUS	The Vuser has arrived at the rendezvous and is waiting to be released by LoadRunner.
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 scenario in the box at the top right-hand corner of the Run view.


Scenario Status		Running 
Running Vusers	10	
Elapsed Time	00:00:38 (hh:mm:ss)	
Hits/Second	1.48 (last 60 sec)	
Passed Transactions	10	
Failed Transactions	0	
Errors	0	

Note: You can detach the Scenario Status window from the Run view by clicking the button in the upper right-hand corner. This allows you to enlarge the Scenario Groups pane.

Status Summary	Description
SCENARIO STATUS	indicates whether the scenario 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 scenario
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

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 Scenario 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.

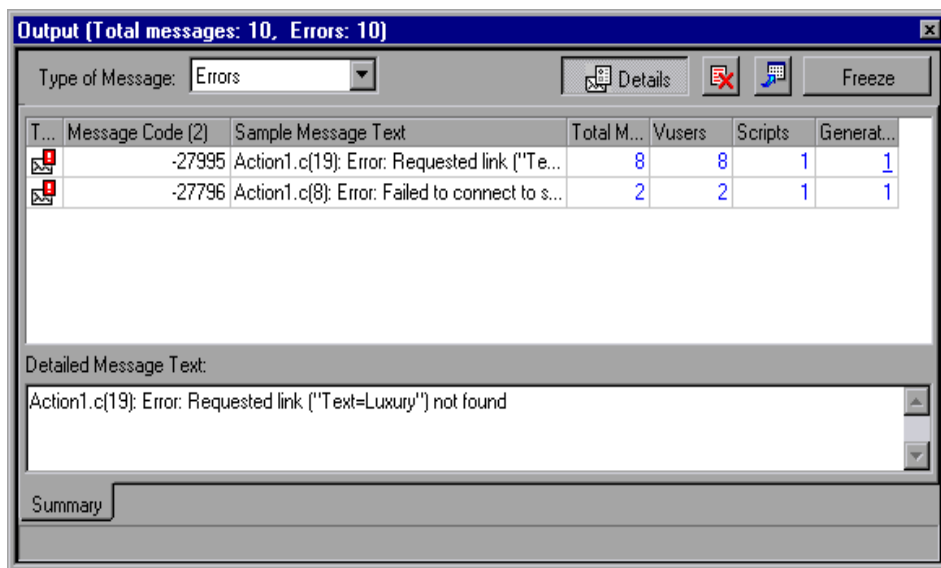


T...	Message Code (2)	Sample Message Text	Total M...	Vusers	Scripts	Generat...
	-27995	Action1.c(19): Error: Requested link ("Te...	8	8	1	1
	-27796	Action1.c(8): Error: Failed to connect to s...	2	2	1	1

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.

- 3 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 sample message text.



Viewing the Output Window

While the scenario runs, the Vusers and load generators send error, notification, warning, debug, and batch messages to the Controller. You can view 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 LoadRunner clears the messages in the Output window at the start of each scenario execution. If you reset a scenario, messages remain in the Output window unless you instruct LoadRunner to delete messages from the Output window upon reset. For more information, see Appendix C, “Options - Output Settings.”



Note: You can also specify the maximum number of Vuser logs that may be displayed simultaneously on the Controller machine. For more information, see Appendix C, “Working in Expert Mode.”

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

Column	Description
TYPE	the type of message sent: Error, Notify, Warning, Debug, or Batch (each represented by a different icon) 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 Controller, 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 Scenario 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 LoadRunner 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 Detailed 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.



Output(10)

Active Filter: Error_Code: -27798 Viewed by: VUser

Vuser	Message	Count	#Messa...	Script	#Lines	Generator
test6:1	Action1.c(6): Error: cou	1	1	test6	1	localhost
test6:10	Action1.c(6): Error: cou	1	1	test6	1	localhost
test6:2	Action1.c(6): Error: cou	1	1	test6	1	localhost
test6:7	Action1.c(6): Error: cou	1	1	test6	1	localhost
test6:8	Action1.c(6): Error: cou	1	1	test6	1	localhost
test6:9	Action1.c(6): Error: cou	1	1	test6	1	localhost

Summary Filtered

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 Vuser Script Log

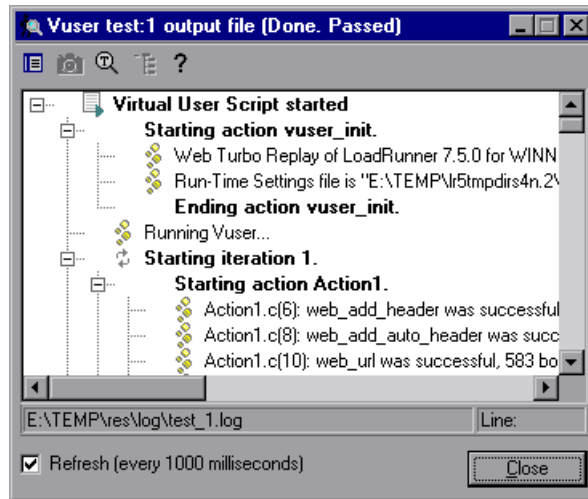
During scenario execution, you can view a log containing run-time information about each running Vuser.

To view the Vuser 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 Vuser 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 594.

Note: If you disabled the logging feature in the Run-Time Settings Log tab, the Vuser script log will contain output only if your script contains the **lr_output_message** or **lr_message** function. If you selected the **Send messages only when an error occurs** option in the Log tab, the Vuser 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: 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 General tab of the Run-Time Settings dialog box before running the scenario.



➤ 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.

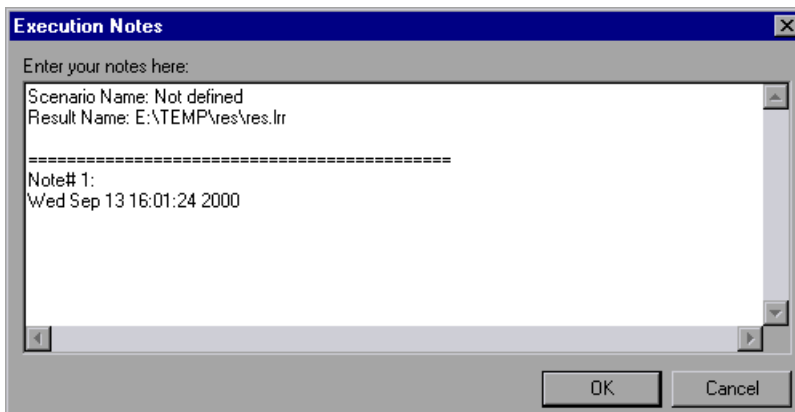
➤ To copy text from the Vuser log, right-click on the selected text in the Vuser log, and click the **Copy** button.

➤ To copy the path of the Vuser log, right-click on the path in the status bar, and click the **Copy path from status bar** button.

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

Logging Execution Notes

The Controller provides you with a dialog box in which you can log comments while a scenario is running.



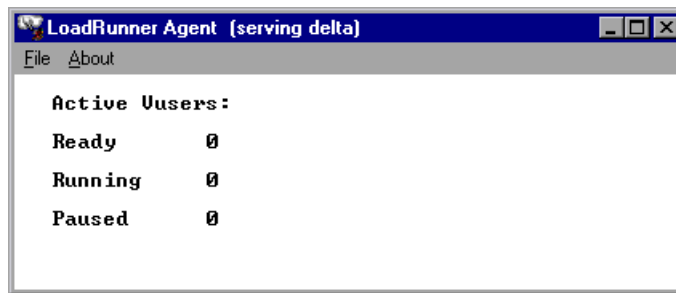
To log execution notes:

- 1** Select **Scenario > Execution Notes**. The Execution Notes dialog box opens.
- 2** Enter the note or notes that you want to log.
- 3** Click **OK** to close the dialog box. LoadRunner saves the note(s) you recorded.

Viewing the Agent Summary

When you run a scenario with non-GUI Vusers, the machine running the Vusers invokes an agent that controls the Vuser execution on that load generator. During scenario execution, the agent displays a summary of the Ready, Running, and Paused Vusers.

The Agent window comes forward at the start of the scenario. You can minimize and restore it at any time.



15

Working with Firewalls

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

This chapter describes:

- ▶ Overview of Running or Monitoring over the Firewall
- ▶ Configuring the LoadRunner Agents in LAN1
- ▶ Configuring the Firewall to Allow Agent Access
- ▶ Installing and Configuring the MI_Listener in LAN2
- ▶ Configuring the Controller to Run or Monitor Vusers over the Firewall

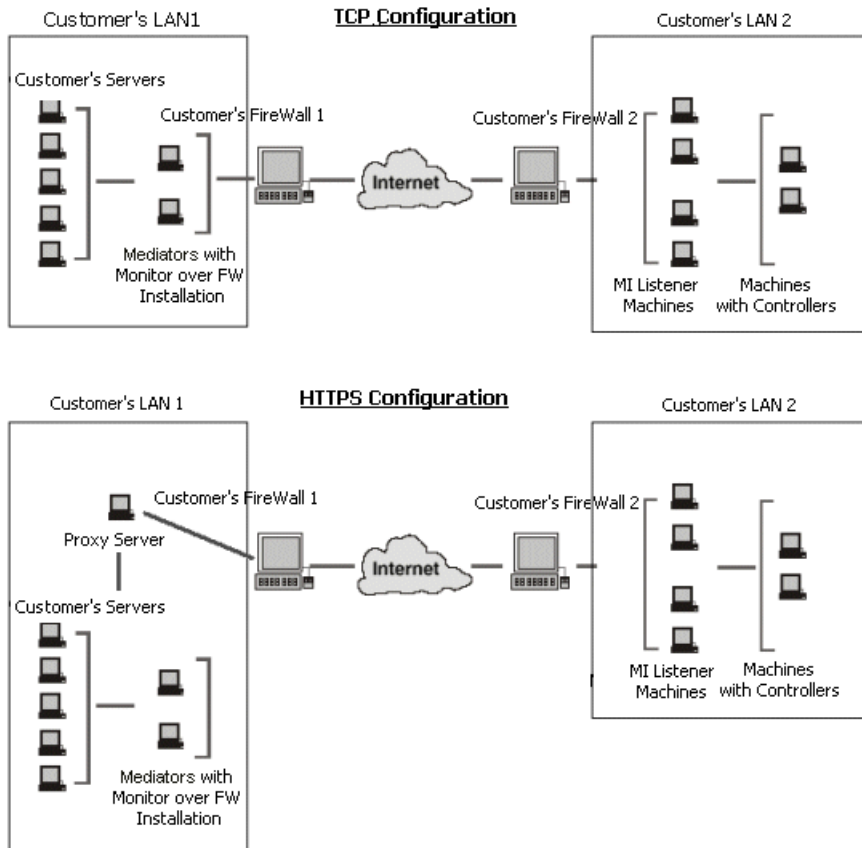
About Using Firewalls in LoadRunner

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 LoadRunner scenario (not over the firewall), the Controller has direct access to the LoadRunner agents running on remote machines. This enables the Controller 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 Controller, 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 LoadRunner 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 Controller and the LoadRunner agent.

In the above configuration, the MI Listener is on a different machine than the Controller. Every LoadRunner agent can behave as a MI Listener, so you

can use the Controller machine as the MI Listener also, and you do not need a separate installation.

TCP Configuration

The TCP configuration requires every LoadRunner 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 LoadRunner 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 Controller to work over the firewall.

Refer to Chapter 17, “Monitoring over a Firewall” for additional information about configuring LoadRunner to monitor servers from outside the firewall.

Note: If you are using the WAN emulator, you should add the IP address of the MI Listener machine or the proxy server to the Exclude IP list. For more information, refer to “Excluding IP Addresses from WAN Emulation,” on page 76.

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 LoadRunner 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 LoadRunner," on page 195.

2 Configure the LoadRunner agent to operate over the firewall.

Configure the LoadRunner agent on the machines running Vusers, or agents acting as mediators for the servers to be monitored. See "Configuring the LoadRunner Agents in LAN1," on page 199 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 207.

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 Controller, and the monitored server. See the diagram, "About Using Firewalls in LoadRunner," on page 195 for information about where to install the Monitoring over the Firewall component, and refer to the *Installing LoadRunner* guide for installation instructions.

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

See the diagram, "About Using Firewalls in LoadRunner," on page 195 for information about where to install the MI Listener, and refer to the *Installing LoadRunner* 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 208.

7 Configure the Controller machine.

Configure the Controller machine to recognize the agent and MI Listener machines. See "Configuring the Controller to Run or Monitor Vusers over the Firewall," on page 209.

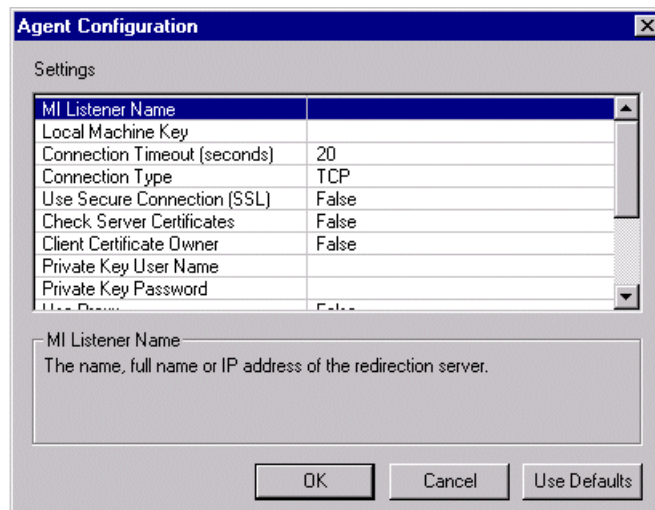
Configuring the LoadRunner Agents in LAN1

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

Configuring and Running the Windows LoadRunner Agent

To configure the LoadRunner agents on Windows machines:

- 1 Stop the LoadRunner agent by right-clicking its icon in the system tray and selecting **Close**.
- 2 Select **Agent Settings** from Start > Programs > LoadRunner > Advanced Settings (or open `<LR 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 > LoadRunner > Advanced Settings, or run `<LR>\launch_service\bin\AgentConfig.exe`.



- 5 Set each option as described in “Agent Configuration Settings” on page 205.

- 6 Click **OK** to save your changes, **Cancel** to cancel them, or **Use Defaults**.
- 7 Restart the LoadRunner agent by double-clicking the shortcut on the desktop, or from Start > Programs > LoadRunner > LoadRunner Agent Service/Process.

Configuring and Running the UNIX LoadRunner Agent

To configure the LoadRunner agents on UNIX machines:

- 1 Open `<LoadRunner 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 `<LoadRunner 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 205.

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 LoadRunner (as described in “Agent Configuration Settings,” on page 205), enter 5.

To start or remove the LoadRunner agent:

- 1** To start the LoadRunner agent, run the command 'm_daemon_setup -install' from the <LoadRunner root folder>/bin directory.
- 2** To remove the LoadRunner agent, run the command 'm_daemon_setup -remove' from the <LoadRunner root folder>/bin directory.

For more information about running the LoadRunner agent, refer to “UNIX Shell” in Appendix D, “Troubleshooting the Controller.”

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 Controller 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 Use Secure Connection option is set to True .
<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 Use Secure Connection option is set to True .
<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 Client Certificate Owner option is set to True .
<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 Client Certificate Owner option is set to True .
<i>Proxy Name</i>	<IE proxy server name> or None	The name of the proxy server. This option is mandatory if the Connection Type option is set to HTTP .
<i>Proxy Port</i>	<IE proxy server port> or None	The proxy server connection port. This option is mandatory if the Connection Type option is set to HTTP .
<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 LoadRunner 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:

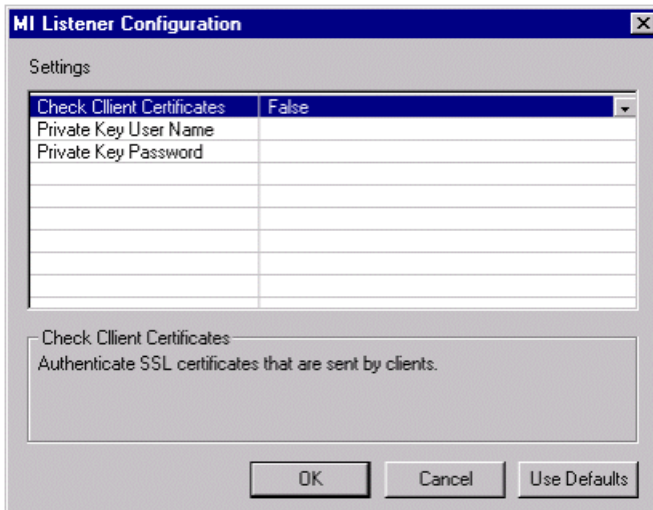
The LoadRunner agent tries to establish a connection with MI Listener using the proxy port specified in the Proxy Port field and at an interval of seconds specified in the Connection Timeout field in the agent configuration. On successful connection, the proxy server connects to MI Listener. To enable 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 *Installing LoadRunner* guide.

To configure the MI Listener:

- 1 Open incoming HTTPS service for port 443.
- 2 Stop the LoadRunner 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 > LoadRunner > Advanced Settings, or run `<LR_root_dir>\launch_service\bin\MILsnConfig.exe`.



- 4 Set each option as described in “MI Listener Configuration Settings,” on page 209.
- 5 Click **OK** to save your changes, **Cancel** to cancel them, or **Use Defaults**.
- 6 Restart the LoadRunner agent by double-clicking the shortcut on the desktop, or running it from Start > Programs > LoadRunner.
- 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.

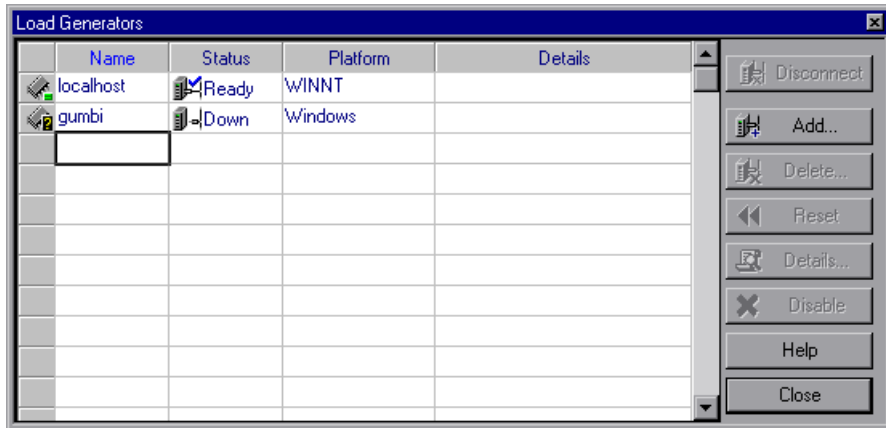
Configuring the Controller to Run or Monitor Vusers over the Firewall

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 Controller 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.

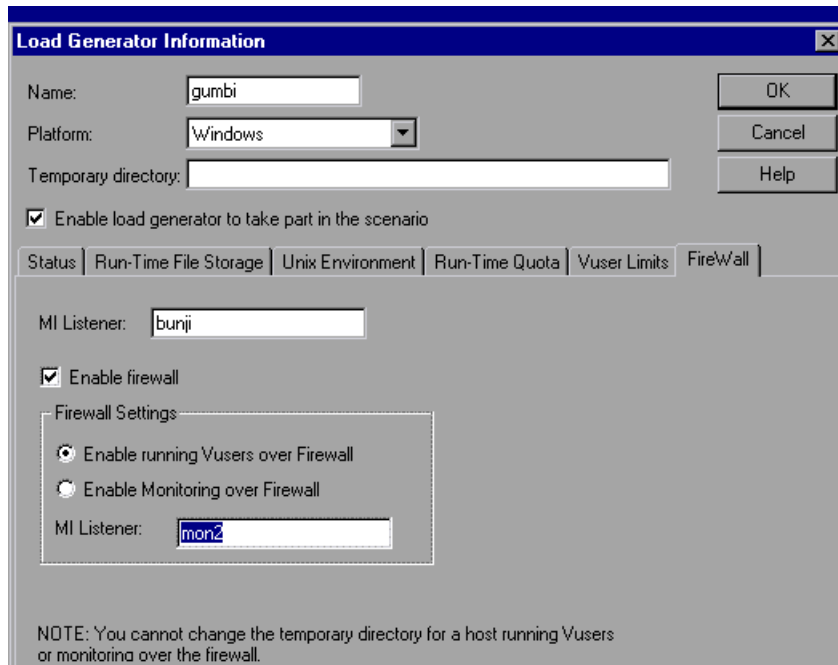
To configure the Controller for running vusers or monitoring over the firewall:

- 1 Run the Controller from Start > Programs > LoadRunner and create a new scenario, 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 IV

Monitoring a Scenario

16

Online Monitoring

You can monitor scenario execution using the LoadRunner 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, and Middleware performance monitors.

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

- Starting 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

LoadRunner provides the following online monitors:

The **Run-Time** monitor displays the number and status of Vusers participating in the scenario, 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 scenario execution. For more information, see Chapter 19, "Run-Time and Transaction Monitoring."

The **Web Resource** monitor measures statistics at the Web server(s) during scenario runs. It provides information about the number of Web connections, throughput volume, HTTP responses, server retries, and downloaded pages during the scenario. For more information on the Web Resource monitor, see Chapter 20, "Web Resource Monitoring."

The **System Resource** monitors gauge the Windows, UNIX, SNMP, Antara FlameThrower, and SiteScope resources used during a scenario. To activate the System Resource monitors, you must set the monitor options before you run your scenario. For information on setting these options, see Chapter 21, "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 scenario. For more information see Chapter 22, "Network Monitoring."

The **Firewall** monitor measures statistics at the firewall servers during the scenario. To activate the Firewall monitor, you must set up a list of resources to monitor before you run your scenario. For more information, see Chapter 23, "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 scenario. To activate the Web Server Resource monitors, you must set up a list of resources to monitor before you run your scenario. For more information, see Chapter 24, "Web Server Resource Monitoring."

The **Web Application Server Resource** monitors measure statistics at the Web application server(s) during the scenario. To activate the Web Application Server Resource monitors, you must set up a list of resources to monitor before you run your scenario. For more information, see Chapter 25, “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 scenario. For more information, see Chapter 26, “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 scenario. For more information, see Chapter 27, “Streaming Media Monitoring.”

The **ERP Server Resource** monitor measures statistics at the ERP servers during the scenario. To activate the ERP Server Resource monitor, you must set up a list of resources to monitor before you run your scenario. For more information, see Chapter 28, “ERP 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 scenario. For more information, see Chapter 29, “Java Performance Monitoring,” and Chapter 30, “J2EE Performance Monitoring.”

The **Application Deployment Solutions** monitor measures statistics of the Citrix MetaFrame XP and 1.8 servers during a scenario run. To activate the Application Deployment Solutions monitor, you must set the monitor options before you run your scenario. For information on setting these options, see Chapter 31, “Application Deployment Solutions.”

The **Middleware Performance** monitors measure statistics of the TUXEDO, and IBM WebSphere MQ servers during a scenario run. To activate the Middleware Performance monitors, you must set the monitor options before you run your scenario. For information on setting these options, see Chapter 32, "Middleware Performance Monitoring."

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


Note: For a detailed list of LoadRunner'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)).

Starting the Monitors

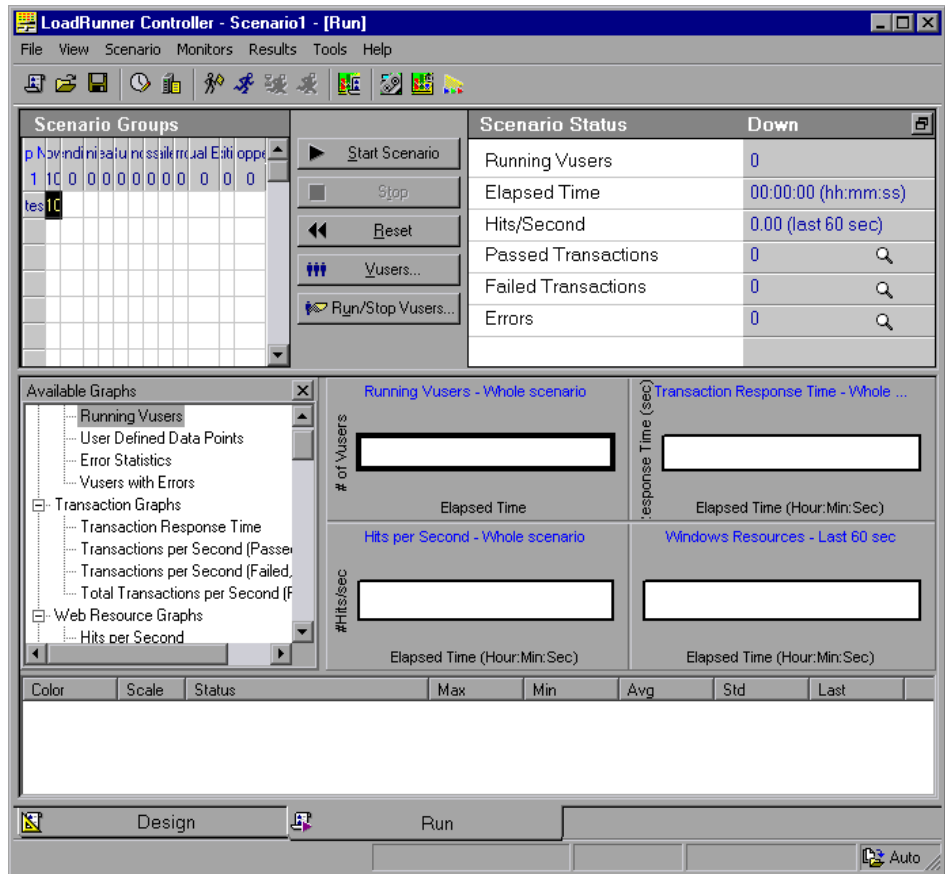
You use the online monitors to monitor Vuser status, errors, transactions, system resources, Web resources, network delay, firewall server resources, Web server resources, Web application server resources, database server resources, streaming media resources, ERP server resources, and Java performance.

To start the online monitors:

- 1 Start the scenario. Select the Vuser groups you want to run and click the **Start Scenario** button, or choose **Scenario > Start**.



- Click the **Run** tab. The default graphs are displayed below the Scenario Groups area.



- Double-click a graph to maximize it. Repeat the operation to restore the tiled view.
- 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.

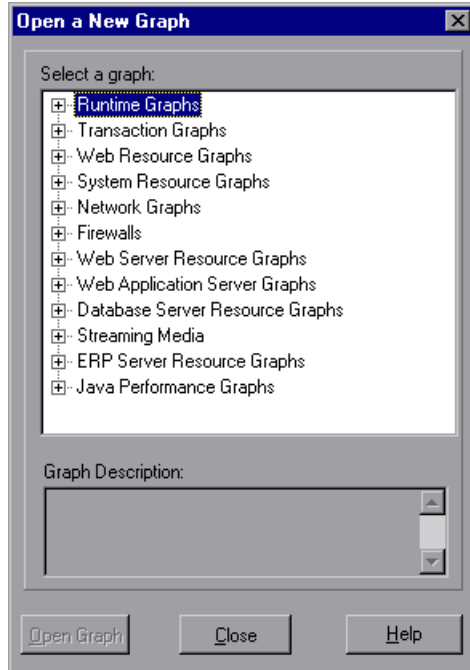
Note: The Transaction Monitor graphs will not contain any data unless transactions are being executed. In addition, the System Resource, Network, Firewall, Web Server, Web Application Server, Database, Streaming Media, ERP Resource, and Java Performance graphs will not contain any data unless you set up a list of resources to monitor before running your scenario.

Opening Online Monitor Graphs

By default, LoadRunner displays four graphs in the Run view: Running Vusers, Transaction Response Time, Hits per Second, and Windows Resources. You can display the other graphs by clicking and dragging them from the graph tree view to the graph view area. Alternatively, you can open a new graph using the Open a New Graph 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.

Customizing the Graph Display View

LoadRunner lets you display up to 16 online monitor graphs simultaneously.

To customize your online graph display:

Click **View > 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

LoadRunner 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 223.

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 225.

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 231.

When you save a scenario, the online monitor configuration settings are saved as well.

Setting Monitor Options

Before running your scenario, 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 225.

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 scenario and reopen it.

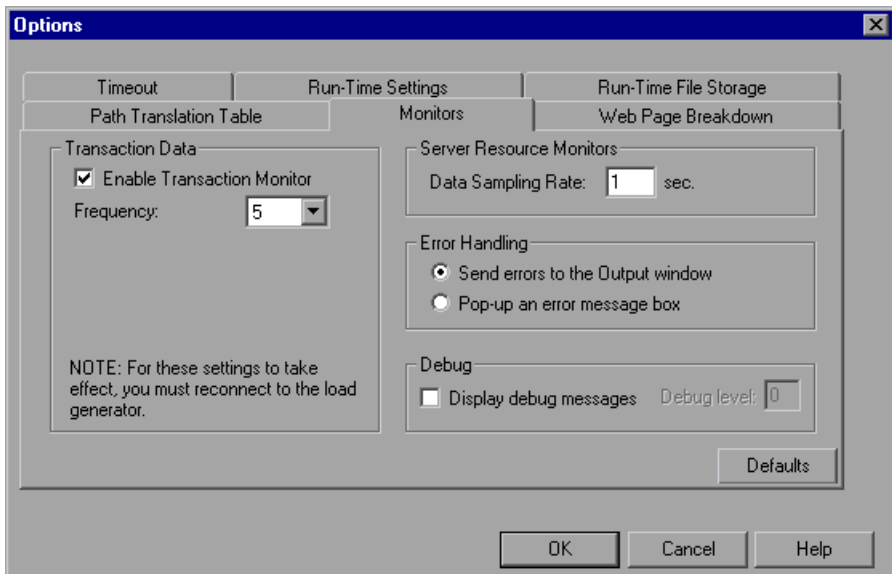
Note: Each monitor has a different minimum sampling rate. If the default sampling rate, or the rate set in the Options Monitors tab is less than a monitor’s minimum sampling rate, the monitor will sample data at its minimum sampling rate. For example, the minimum sampling rate for the Oracle Monitor 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 LoadRunner should behave when a monitor error occurs—issue a popup message box (default) or send error messages to the Output window.
- ▶ **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 Controller 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 Controller.

For information on enabling and disabling the Transaction monitor and Web page breakdown, see Chapter 19, “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 Controller for the Transaction, Data Point, and Web Resource graphs. The default value is 5 seconds. For a small scenario, it is recommended that you use a frequency of 1. For a large scenario, 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 scenario execution; you must stop the scenario 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 C, “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.

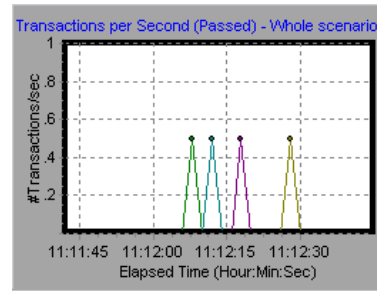
Note: In a large load test, it is recommended to use a refresh rate of three to five seconds. This enables you to avoid problems with CPU resource usage.

X-Axis Style

You can specify how the graph displays the x-axis time: *Don't Show*, *Clock Time*, or *Relative to Scenario Start*. The *Don't Show* setting instructs LoadRunner not to display values for the x-axis. The *Clock Time* setting displays the absolute time, based on the system clock. The *Relative to Scenario Start* setting displays the time relative to the beginning of the scenario. 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 Scenario*, *60*, *180*, *600*, and *3600* seconds.

Display Type

You can specify whether LoadRunner 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 LoadRunner to display graphs using the default y-axis scale, or you can specify a different y-axis scale. Click **Automatic** if you want

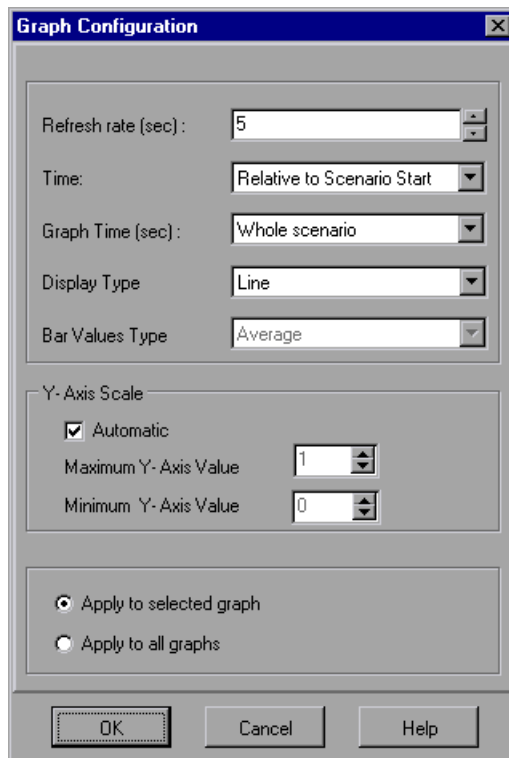
LoadRunner to use the default y-axis values. Specify a maximum or minimum value for the y-axis if you want to modify the y-axis scale.

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.

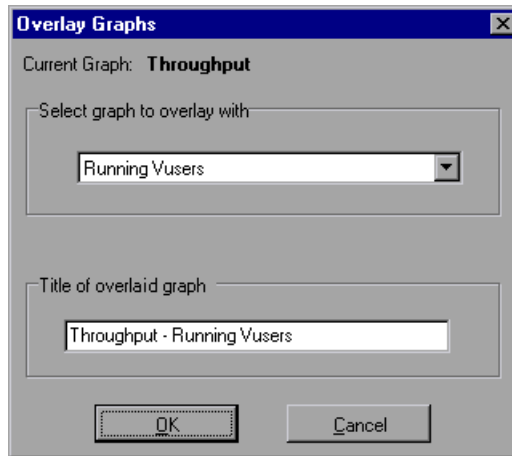
Merging Graphs

LoadRunner lets you merge the results of two graphs from the same scenario 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 performance measurements for those resources being monitored in a scenario. Each measurement is represented on the graph by a colored line, and in the legend which appears beneath the graph (in the same color). The legend displays the measurements for the selected graph only.

The screenshot displays the LoadRunner Controller interface for a scenario named 'Scenario3'. The interface is divided into several sections:

- Scenario Groups:** A table listing various groups and their status. The 'test' group is highlighted.
- Scenario Status:** A table showing the current status of the scenario, which is 'Running'. Key metrics include:

Metric	Value
Running Vusers	10
Elapsed Time	00:00:25 (hh:mm:ss)
Hits/Second	3.00 (last 60 sec)
Passed Transactions	10
Failed Transactions	0
Errors	0
- Available Graphs:** A list of graphs available for monitoring, including 'Running Vusers', 'Transaction Response Time', and various throughput and error metrics.
- Running Vusers - Whole scenario:** A line graph showing the number of running users over time. The y-axis is '# of Vusers' (0 to 20) and the x-axis is 'Elapsed Time' (00:00:00 to 00:00:15). The graph shows a sharp increase from 0 to 10 users within the first few seconds, followed by a steady state at 10 users.
- Transaction Response Time - Whole scenario:** A line graph showing the response time in seconds over time. The y-axis is 'Response Time (sec)' (0 to 1) and the x-axis is 'Elapsed Time (Hour:Min:Sec)' (00:00:00 to 00:00:15). The graph shows a very low and stable response time, consistently near 0.1 seconds.
- Legend:** A table at the bottom of the interface showing the status of different resources.

Color	Scale	Status	Max	Min	Avg	Std	Last
Black	1	Error	0.000	0.000	0.000	0.000	0.000
Blue	1	Finished	0.000	0.000	0.000	0.000	0.000
Purple	1	Ready	0.000	0.000	0.000	0.000	0.000

Note: In a goal-oriented scenario, 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 scenario 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:


- ▶ Line Color
- ▶ Measurement Scale
- ▶ Transaction Display

Line Color

LoadRunner assigns a unique color to each measurement. You can modify the color using the configuration interface.

To change the line color of 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 Click **OK** to accept the settings and close the dialog box.

The specified color changes are reflected in the graph and in the legend beneath the graph. The color is displayed in the first column of the legend.

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	


Measurement Scale

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.


To set the scale of a measurement:

- 1** Select the measurement you want to configure. Right-click and choose **Configure**. The Measurement Configuration dialog box opens.
- 2** Clear the **Autoscale** check box and select the desired ratio from the Scale list.
- 3** Click **OK** to accept the settings and close the dialog box.

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.

Color	Scale	Measurement	Machine	Max	Min	Avg	Std	Last
	10	Processor Queue Length (System)	zeus	3	1	1.823529...	0.705882...	1
	1	File Data Operations/sec (System)	zeus	127.1469...	16.64241...	43.56583...	24.31799...	49.9280

|
scale
factor

By default, LoadRunner uses the *autoscale* option, which automatically scales the measurements by calculating the best ratio for displaying the graph.

Transaction Display

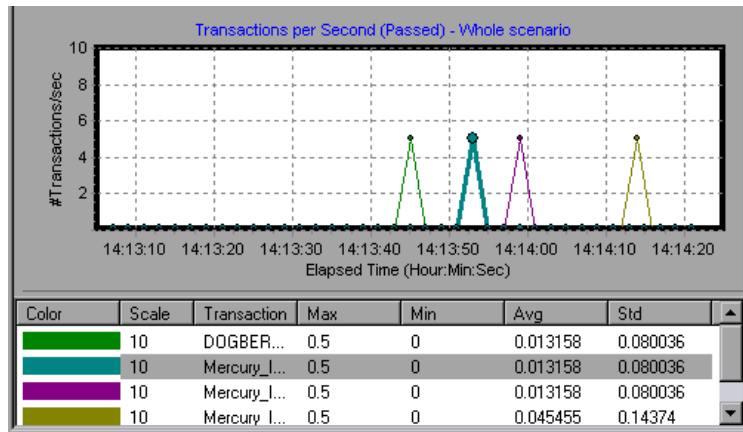
By default, the Transaction Monitor displays a line for each item in the transaction list. You can hide the line for any of the monitored transactions in order to focus on a specific measurement.

To show or hide a transaction:

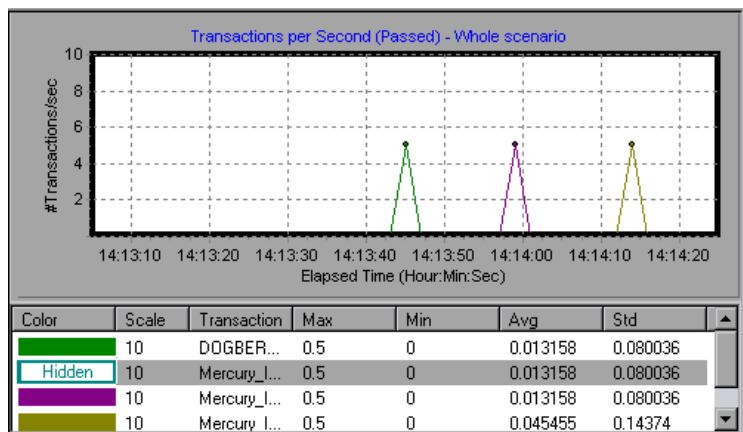
- 1** To hide a measurement, click **Hide**. To show a hidden resource, click **Show**.
- 2** Click **OK** to accept the settings and close the dialog box.

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**.

In the following example, a line is shown for each measurement.



In this example, the second item in the legend is hidden.



Exporting Online Monitor Graphs

LoadRunner 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 scenario run, you can view a graph of the data that was gathered using the LoadRunner 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 LoadRunner Analysis at the conclusion of the scenario, see the *LoadRunner Analysis User's Guide*.

17

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 LoadRunner 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 15, “Working with Firewalls,” and install the Monitoring over Firewall component, as described in “Installing Monitors over Firewall,” on page 238, 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 243.

2 Configure server monitor properties.

Refer to “Configuring Server Monitor Properties,” on page 244.

3 Add and remove measurements.

Add measurements to monitor for each server. If LoadRunner added default measurements, you can edit them as required. Refer to “Adding and Removing Measurements,” on page 246.

4 Configure measurement frequencies.

Set a measurement schedule for each measurement to be reported. Refer to “Configuring Measurement Frequency,” on page 247.

Installing Monitors over Firewall

Monitors over Firewall may have been installed during LoadRunner installation. To check whether it was installed, click **Start > Programs > LoadRunner > Advanced Settings**. If the **Monitor Configuration** option appears on the list of LoadRunner options, then *Monitors over Firewall* was already installed, and you can proceed to “Installing MI_Listener” on page 243.


If Monitors over Firewall was not yet installed, you need to install it using one of the following:

- Perform a custom installation of LoadRunner from the LoadRunner CD, choosing only the Monitors over Firewall option. For instructions on performing a custom installation of LoadRunner, refer to the *Installing LoadRunner* guide.

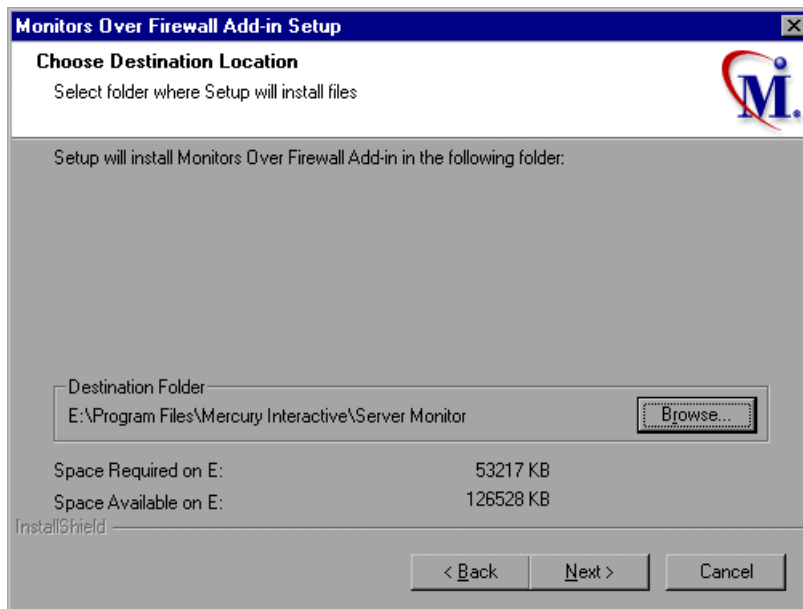
- Obtain the *Monitors over Firewall* file from the Mercury Interactive Customer Support Web site (<http://support.mercuryinteractive.com>). *Monitors over Firewall* is a stand-alone downloadable installation. It comes as a self-extracting installer file.

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.

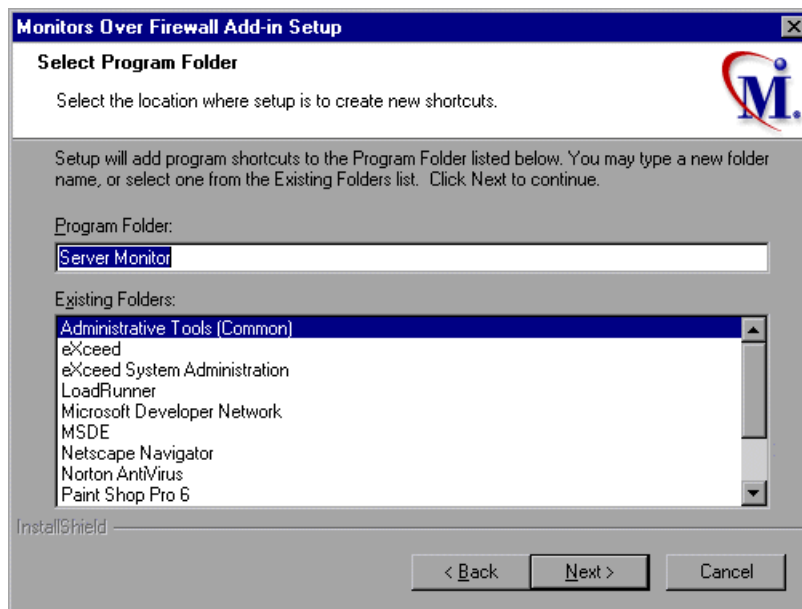


- 3 In the Choose Destination Location screen, specify the folder in which to install the add-in. To select a different location, click **Browse**, choose a folder, and click **OK**.



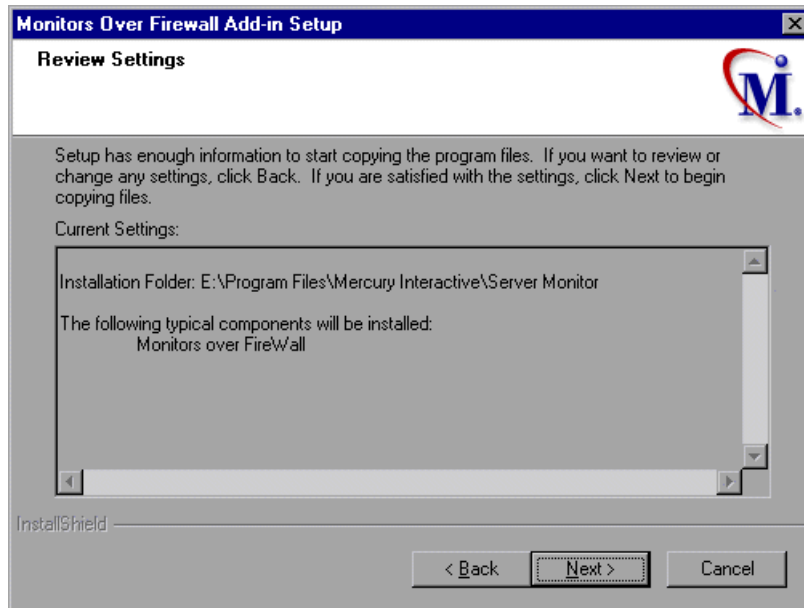
Click **Next**.

- 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 LoadRunner 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 Controller machine. Note that the Controller installation automatically includes the MI Listener, so you can designate the Controller as the MI Listener machine. For instructions, refer to the *Installing LoadRunner* 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 Controller 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 Controller for data collection:

- 1** You should already have configured the LoadRunner agent and the Controller to operate over the firewall, as described in Chapter 15, “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**.

Note: If you are using WAN emulation, you should add the IP address of the MI Listener machine to the Exclude IP list. For more information, refer to “Excluding IP Addresses from WAN Emulation,” on page 76.

- 3** Connect to the load generator. Make sure that you obtain information for the monitors configured inside the firewall.

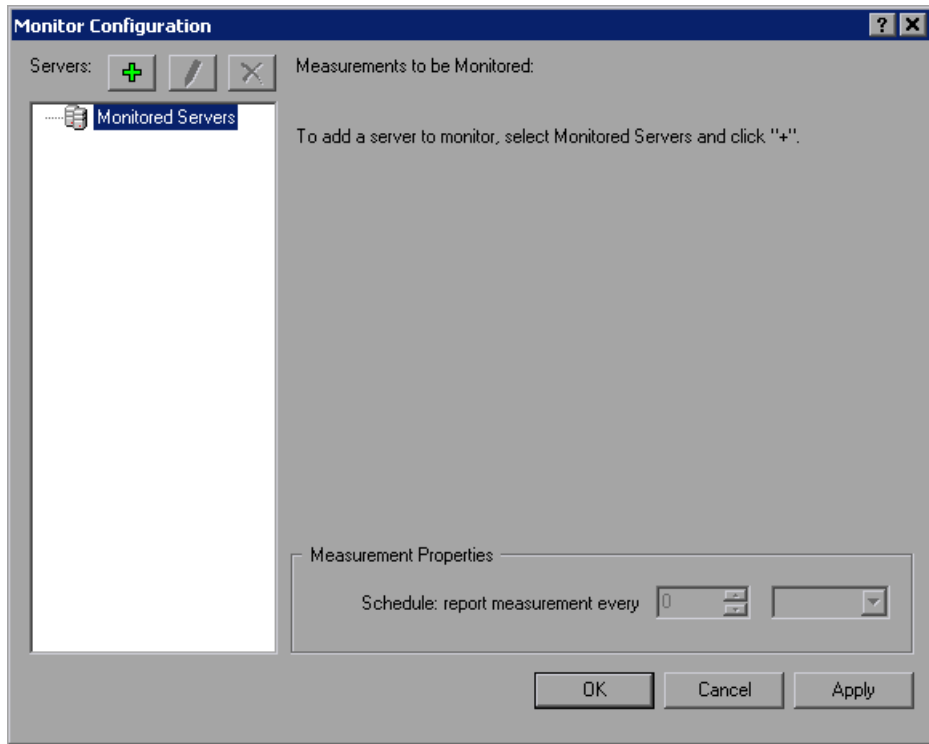
Configuring Server Monitor Properties

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.

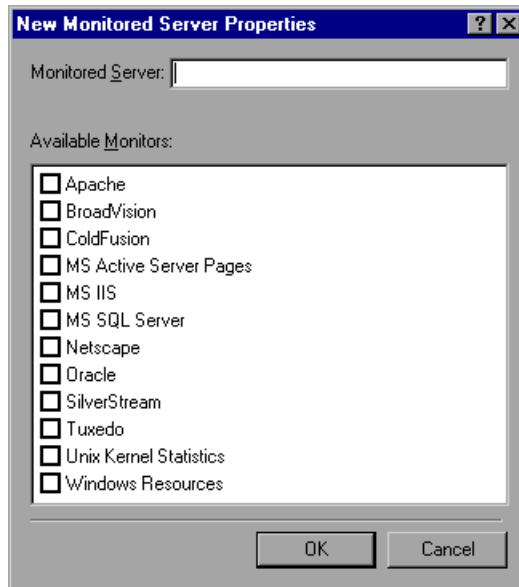
To configure server monitor properties:

- 1 Select **Start > Programs > LoadRunner > Advanced Settings > Monitor Configuration**. For machines without the complete LoadRunner installation, select **Start > Programs > Server Monitor > Monitor Configuration**. The Monitor Configuration dialog box opens.





- Click the **Add Server** button. The New Monitored Server Properties dialog box opens.




- 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.

- From the Available Monitors list, select the monitors appropriate for the server being monitored.

Note: Data can only be viewed for the monitors that are enabled with your LoadRunner license key. To preview your license key information, in the LoadRunner Controller, select **Help > About LoadRunner**.

- 5 Click **OK** to close the New Monitored Server Properties dialog box to display the Monitored Servers list.




Note that, for certain monitors, LoadRunner displays default measurements in the right pane. For details on selecting measurements, see “Adding and Removing Measurements” on page 246.

- 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 LoadRunner 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 LoadRunner 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 Controller machine and the source machine, you must configure the Network Delay Monitor (see “Configuring the Network Monitor,” on page 312), and add the following to step 3 (on page 313):

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 LoadRunner Agent on the source machine.

For example: 12.12.12.3:vds

18

Remote Performance Monitoring

Remote performance monitoring enables additional viewers to monitor a LoadRunner scenario from a remote location using a Web browser. This allows a licensed number of participants to simultaneously view the online test results without requiring access to the Controller machine. Each remote viewer is able to select the graphs that he chooses to monitor, and customize the graph settings to suit his needs.

This chapter describes:

- ▶ Installing the Remote Performance Monitor Server
- ▶ Configuring the Remote Performance Monitor User Settings
- ▶ Connecting to the LoadRunner Remote Performance Monitor
- ▶ Monitoring Load Test Data
- ▶ Viewing Online Graphs
- ▶ Customizing Online Graph Settings

About Remote Performance Monitoring

You can monitor scenario execution and view online results from any browser using LoadRunner's Remote Performance Monitor. The Remote Performance Monitor enables multiple load tests to participate in a scenario run with customized resource monitoring.

During a load test run, the Remote Performance Monitor enables you to view graphs that display information about the load the Vusers generate on your server. Online graphs display data about the run, including Vuser status, errors, 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, and Java performance monitors.

For more information about the available graphs and monitor measurements, see Chapter 16, "Online Monitoring."

Installing the Remote Performance Monitor Server

To monitor the performance of your servers from a remote location, you need to install the Remote Performance Monitor Server from the LoadRunner Controller 7.6 CD.

Installation Requirements

The Remote Performance Monitor Server configures an IIS Web server for the Remote Performance Monitor. This requires a machine with the following components installed:

IIS Server	4.0; 5.0
Operating system	Windows 2000 Server; Windows 2000 Advanced Server; Windows NT Server
Client Browser	Internet Explorer 5.0 and later; Netscape 6.2 and later

The IIS Web server communicates with the Controller and the Remote Performance Monitor to manage the user's requests to produce online graphs and graph legends.

For instructions on installing the Remote Performance Monitor Server, refer to the *LoadRunner Controller Installation Guide*.

Configuring the Remote Performance Monitor User Settings

The Remote Performance Monitor User Configuration tool enables you to change the default or used-defined user name and password that was used during the Remote Performance Monitor installation process on the Web server. This tool is also used to update the Remote Performance Monitor user settings on the Controller machine.

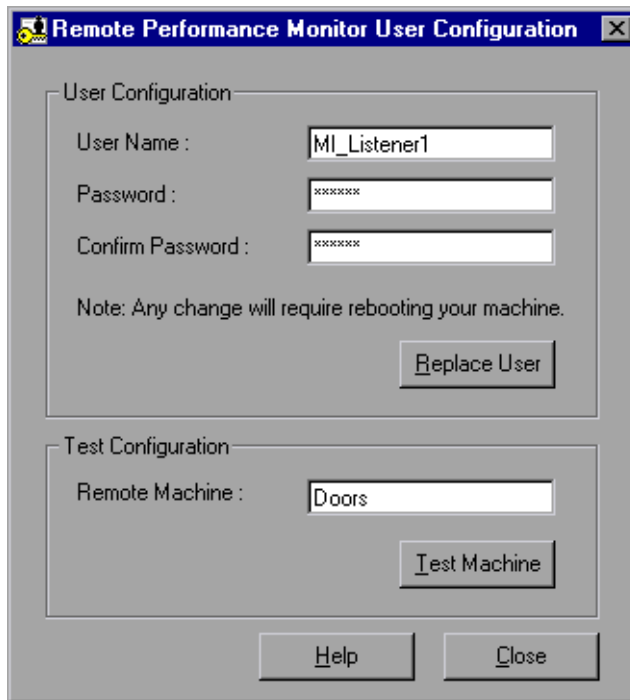
In addition, the Remote Performance Monitor User Configuration tool is used to check that the user name and password on the Web Server and the Controller machine are identical. Since LoadRunner uses the user name and password for authentication between the Web Server and the Controller machine, this information must be the same on both machines.

User Configuration

You must use the Remote Performance Monitor User Configuration tool to configure the user settings on both the Controller and the Web server machines.

To change the user settings on the Controller side:

- 1 Select **Start > Programs > LoadRunner > Tools > RPM User Configuration** to open the Remote Performance Monitor User Configuration tool on the Controller machine.



- 2** In the Remote Performance Monitor User Configuration dialog box, enter your user name and password, and confirm the password.
- 3** Click **Replace User**. The configuration program prompts you to restart the machine. You can delay restarting to a later time.

Note: The changes are performed immediately after you click the **Replace User** button. However, the system will only work properly after rebooting the machine.

To change the user settings on the Web server:

- 1** Select **Start > Programs > RPM Server** to open the Remote Performance Monitor User Configuration tool on the Web server.
- 2** In the Remote Performance Monitor User Configuration dialog box, enter the same user name and password that you entered on the Controller machine. Confirm the password.
- 3** Click **Replace User**. The configuration program prompts you to restart the machine.

Note: The changes are performed immediately after your click the **Replace User** button. However, the system will only work properly after rebooting the machine.

Note: The Remote Performance Monitor user name and password are automatically updated on the IIS Web server.

Configuration Test

After the user settings have been configured on the Controller and the Remote Performance Monitor server machines, and both machines have been restarted, you can run the configuration test.

To run the configuration test:

- 1** On the Controller or the Remote Performance Monitor server machine, open the Remote Performance Monitor User Configuration tool.
- 2** Enter the IP name or address of the other machine (Controller or Web server) in the Remote Machine box of the Test Configuration section.
- 3** Click **Test Machine** to run the configuration test. A message notifies you that the configuration test failed or succeeded.

Connecting to the LoadRunner Remote Performance Monitor

To connect to the LoadRunner Remote Performance Monitor, type the following path in your Web browser:

`http://[name of IIS Web server machine]/remoteview`

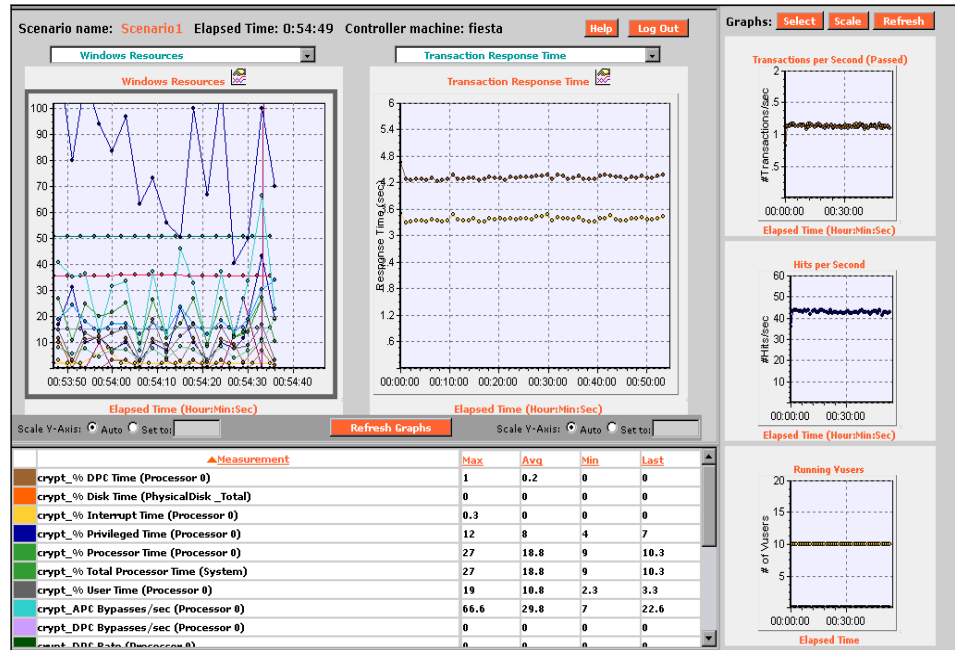
The LoadRunner Remote Performance Monitor log on page opens.



To log on to the LoadRunner Remote Performance Monitor:

- 1** In the **UserID** box, type Admin.
- 2** In the **Password** box, type Admin.
- 3** In the **Controller Machine** box, type the name or IP address of the Controller machine you want to access.

4 Click **Login**. The LoadRunner Remote Performance Monitor page opens.



By default, the left graph is selected, and its measurements are displayed in the measurements legend.

Note: If there is no browser activity for 20 minutes, the Remote Performance monitoring session times out. You need to log in again to continue your session.

Monitoring Load Test Data

You monitor load test data during the load test run to get a quick overview of the test's status and the effects of load on your Web server.

At the top of the Remote Performance Monitor page, you can view the status of the test that is currently running.

Scenario name: **Scenario1.lrs** Elapsed Time: 1:01:41 Controller machine: fiesta

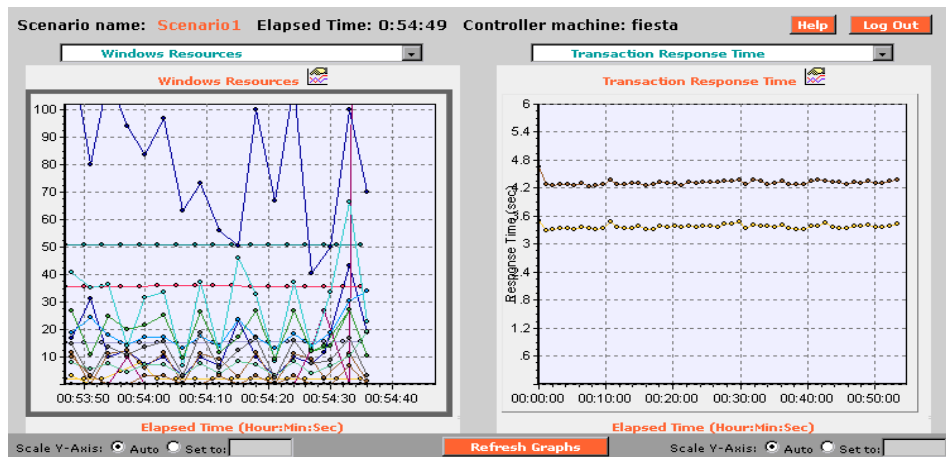
The Remote Performance Monitor page displays the name of the running test, the length of time the test has been running, and the name of the Controller machine.

Viewing Online Graphs

The graphs are tiled to enable you to view five graphs simultaneously: two large graphs and three small graphs. In addition, you can view graph measurements in the legend.

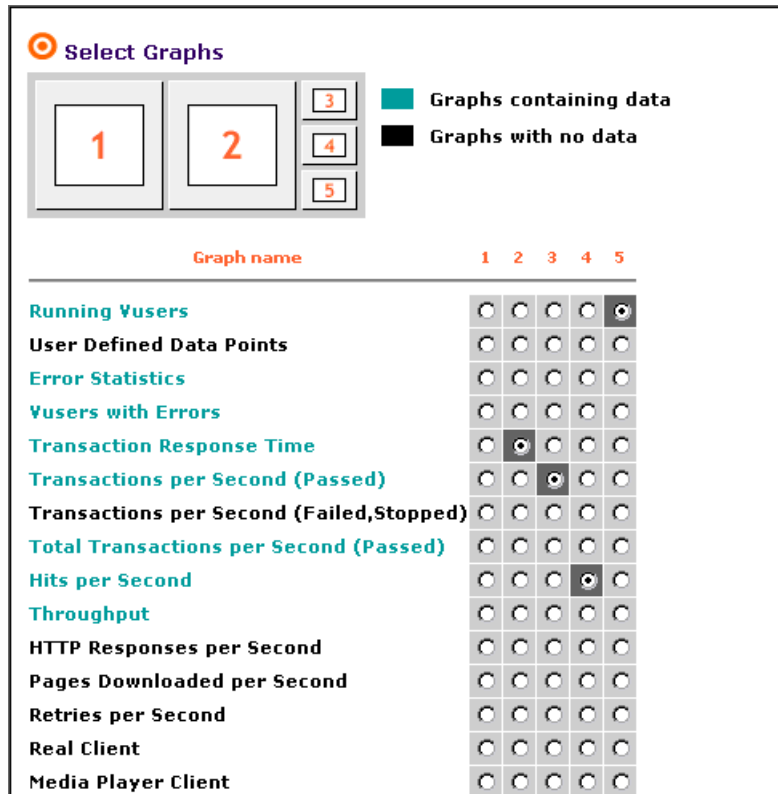
To view graphs during the load test run:

- 1 To display a graph in the large graph pane, select a graph from the drop-down graph list located above the large graph pane. The page reloads with the selected graph.



Note: Available graphs appear in green in the drop-down graph list. If you select an unavailable graph (black), an empty pane will appear.

- To display a graph in the small graph pane, or to change any of the graphs displayed on the screen, click the **Select** button located above the small graphs. The Select Graphs window opens.



- Choose any of the listed graphs and the corresponding position in which they should be displayed. The diagram at the top shows the numbered positions.
- Click **OK** to close the Select Graphs window. The selected graphs appear in the Remote Performance Monitor page.

Graph Legend

You can view the measurements of any graph that is displayed in the large graph window. By default, when the remote viewer opens, the left graph is selected (denoted by a highlighted gray border), and its measurements are displayed in the measurements legend.

Note: The measurements of a small graph cannot be displayed in the legend. Therefore, open the resource you want to measure as a large graph.

To view a graph's legend:

- 1** Click in the graph pane to select the graph. The graph is highlighted with a gray border, and its measurements are displayed in the legend.

To view the measurements of two graphs, select both of the large graphs. The legend splits vertically, displaying the measurements of both graphs.

- 2** The legend displays details about the maximum, average, minimum, and last values for each measurement. To sort the measurements by one of these values, click the column heading (Max, Avg, Min, or Last). An icon is displayed beside the column heading, showing you whether the measurements are sorted in ascending or descending order.
- 3** To close a graph's legend, click anywhere in the graph's pane.

Customizing Online Graph Settings

You can modify the following online graph settings while a load test is running from the Remote Performance Monitor:

- graph scale
- graph refresh rate
- graph configuration measurements

The changes to the default settings only apply to the current run and are not saved for future runs of the load test.

Scaling Graphs

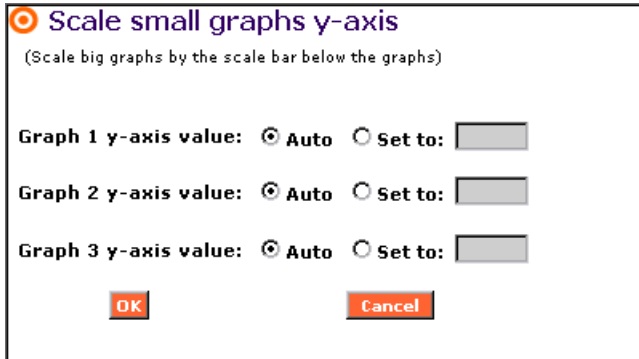
You can modify the scale of a measurement—the relationship between the y-axis and the graph's actual value. The x-axis represents Elapsed Time and cannot be adjusted. By default, LoadRunner uses the **Auto** option, which automatically sets the most suitable measurements for displaying the graph.

To scale the large graphs:

- 1** In the y-axis value section below the large graphs, select **Set to**, type a value in the box, and click **Refresh Graphs**. The graph is redrawn with the specified value as the upper limit of the y-axis.
- 2** To view the graph scaled normally, select **Auto** and click **Refresh Graphs**.

To scale the small graphs:

- 1 Click the **Scale** button located above the small graphs. The Scale Graphs window opens.



You can change the y-axis measurement of the small graphs in positions three, four, and five as shown in the Select Graphs window on page 257.

- 2 To use a different scale, select **Set to**, and, type a value in the box.
- 3 To view the graph scaled normally, select **Auto**.
- 4 Click **OK** to close the Scale graphs window. The graph is redrawn with the specified value as the upper limit of the y-axis.

Note: For more information on graph scale, refer to “Measurement Scale” in Chapter 16, “Online Monitoring.”

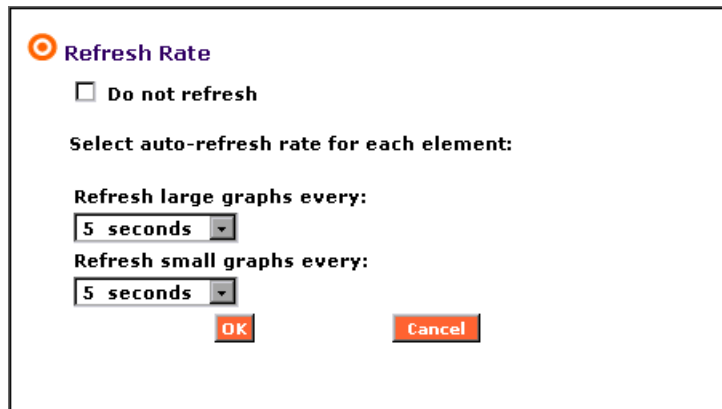
Refresh Rate

By default, graphs on the Remote Performance Monitor page are refreshed every five seconds. You can use the Auto-Refresh option to change the default refresh rate. If you increase the refresh rate, the graphs are refreshed less frequently.

Note: In a large load test, it is recommended to use a slower refresh rate for the small graphs. This enables you to avoid problems with CPU resource usage.

To modify the default refresh rate:

- 1 Click the **Refresh** button located above the small graphs. The Refresh Rate window opens.



By default, the refresh option is enabled to automatically refresh all graphs at 5 second intervals. To disable the auto-refresh option, select the **Do not refresh** check box.

- 2 Select a frequency for refreshing the large graphs and small graphs.
- 3 Click **OK** to modify the auto-refresh rates and return to the Load Test Run page.

Note: To refresh the graphs and the legends immediately, click the **Refresh Graphs** button.

Configure Graph Measurements

You can customize graphs to show, hide, or highlight selected graph measurements.

To configure graph measurements:



- 1 Click the icon at the top of the large graph pane to configure the graph's measurements. The Configure graph's measurements page opens.

Measurement	Show	Bold
crypt_% DPC Time (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_% Disk Time (PhysicalDisk _Total)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_% Interrupt Time (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_% Privileged Time (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_% Processor Time (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_% Total Processor Time (System)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_% User Time (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_APC Bypasses/sec (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_DPC Bypasses/sec (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_DPC Rate (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_DPCs Queued/sec (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_File Data Operations/sec (System)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Interrupts/sec (Processor 0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Page Faults/sec (Memory)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Pages/sec (Memory)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Pool Nonpaged Bytes (Memory)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Private Bytes (Process _Total)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Processor Queue Length (System)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Threads (Objects)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
crypt_Total Interrupts/sec (System)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
(Select/Deselect all)	<input type="checkbox"/>	<input type="checkbox"/>

OK Close

- 2 Select the **Show** check box to show a measurement on the graph. By default, all graph measurements are shown on the graph. To remove a measurement

from being displayed on the graph, clear the **Show** check box. To highlight a measurement in bold on the graph, select the **Bold** check box.

Check the **Select/Deselect All** check box in the Show column to display all measurements on the graph. Clear the **Select/Deselect All** check box to remove all measurements from the graph.

To highlight all measurements in Bold on the graph, check the **Select/Deselect All** check box in the Bold column. To remove the bold highlights, clear the **Select/Deselect All** check box.

- 3 Click **OK** to close the Configure graph measurements page. The new graph configuration settings appear on the refreshed graph.

Log Out

To log out of the LoadRunner Remote Performance Monitor, click the **Log Out** button at the top of the page.

19

Run-Time and Transaction Monitoring

While running a scenario, you can use LoadRunner'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 scenario, 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 scenario execution. For more information about transactions, see "Adding Transactions to a Script" on page 270.

Run-Time Graphs

The monitor's **Running Vusers** graph provides information about the status of the Vusers running in the current scenario 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 scenario 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 scenario 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, see the *LoadRunner Function Reference*.

By default, LoadRunner 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 scenario. For more information, see the *LoadRunner 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 scenario (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 scenario (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 scenario (x-axis).

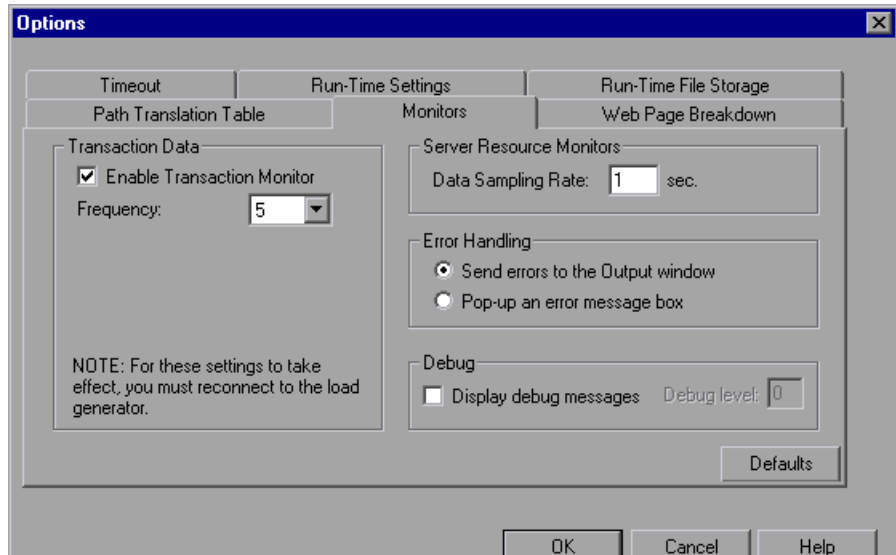
The **Total Transactions per Second (Passed)** graph shows the total number of completed, successful transactions per second (y-axis) as a function of the elapsed time in the scenario (x-axis).

Enabling the Transaction Monitor

The Transaction monitor is enabled by default—it automatically begins monitoring Vuser transactions at the start of a scenario. 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 Scenario 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.

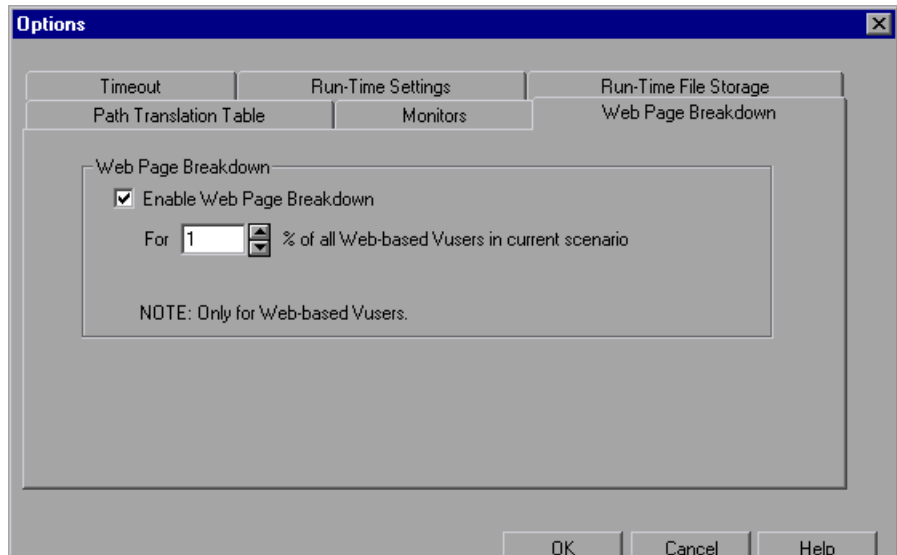
For more information, see the appropriate user's guide as described in the *Welcome* chapter.

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 Controller before running your scenario.

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, see the *LoadRunner Analysis User's Guide*.

20

Web Resource Monitoring

You can obtain information about the performance of your Web server using LoadRunner'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

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 scenario, 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 scenario (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 scenario 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 scenario 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
402	Payment Required

Code	Description
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 scenario run (x-axis). This graph helps you evaluate the amount of load Vusers generate, in terms of the number of pages downloaded.

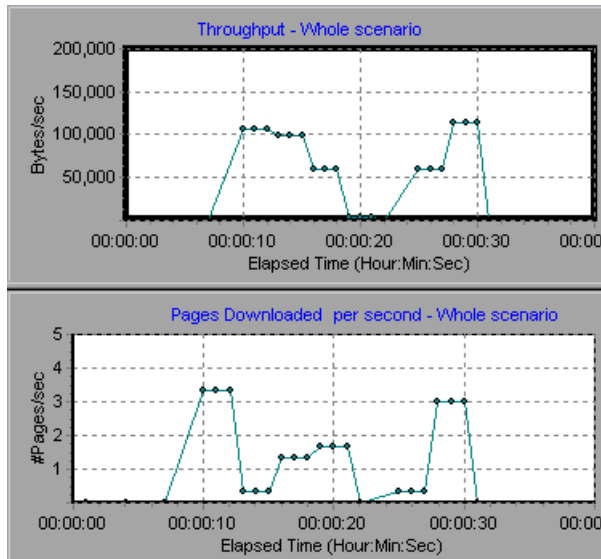
Note: In order to view the Pages Downloaded per Second graph, you must select **Pages per second (HTML Mode only)** from the script's run-time settings Preferences tab before running your scenario.

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 scenario 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 scenario (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.

21

System Resource Monitoring

You can monitor a machine's system resource usage during a scenario run using LoadRunner'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 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 LoadRunner resource monitors, you can monitor the Windows, UNIX, SNMP, Antara FlameThrower, and SiteScope resources on a machine during a scenario 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 289.

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, SMTP, 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, see the relevant SiteScope documentation.

The resource monitors are automatically enabled when you execute a scenario. 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 scenario 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 the Windows Resources graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the Windows Resources dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the Windows Resources dialog box, select the measurements that you want to monitor.

The following default measurements are available for Windows machines:

Object	Measurement	Description
System	% Total Processor Time	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	File Data Operations/sec	The rate at which the computer issues read and write operations to file system devices. This does not include File Control Operations.
Processor	% Processor Time (Windows 2000)	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%.

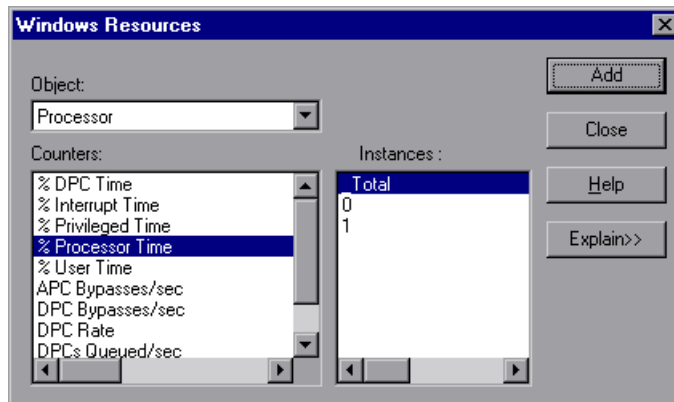
Object	Measurement	Description
System	Processor Queue Length	The instantaneous length of the processor queue in units of threads. This counter is always 0 unless you are also monitoring a thread counter. All processors use a single queue in which threads wait for processor cycles. This length does not 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	Page Faults/sec	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	% Disk Time	The percentage of elapsed time that the selected disk drive is busy servicing read or write requests.
Memory	Pool Nonpaged Bytes	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 623.

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). Proceed to step 5 in order to select counters appropriate for Win2000.

- 5 To select additional measurements, click **Add**. A dialog box displaying the available measurements and server properties opens.



- 6 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.
- 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 Windows Resources 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 Controller machine on the remote Windows machine. To authenticate the Controller machine, create an account, or change the password of the account used to log on to the Controller 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:

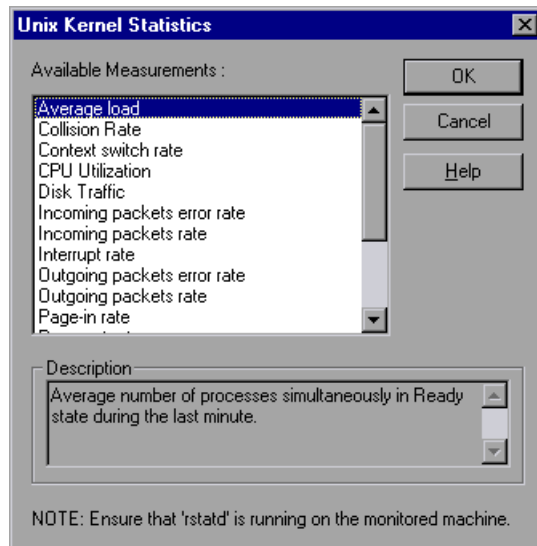
- 1** Click the UNIX Resources graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the UNIX Resources dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select **UNIX** from the list of platforms, and click **OK**.
- 4** In the Resource Measurements section of the UNIX Resources dialog box, select the default measurements you want to monitor.

The following default measurements are available for the UNIX machine:

Measurement	Description
Average load	Average number of processes simultaneously in READY state during the last minute
Collision rate	Collisions per second detected on the Ethernet
Context switches rate	Number of switches between processes or threads, per second
CPU utilization	Percent of time that the CPU is utilized
Disk rate	Rate of disk transfers
Incoming packets error rate	Errors per second while receiving Ethernet packets
Incoming packets rate	Incoming Ethernet packets per second
Interrupt rate	Number of device interrupts per second
Outgoing packets errors rate	Errors per second while sending Ethernet packets
Outgoing packets rate	Outgoing Ethernet packets per second
Page-in rate	Number of pages read to physical memory, per second
Page-out rate	Number of pages written to pagefile(s) and removed from physical memory, per second
Paging rate	Number of pages read to physical memory or written to pagefile(s), per second
Swap-in rate	Number of processes being swapped
Swap-out rate	Number of processes being swapped
System mode CPU utilization	Percent of time that the CPU is utilized in system mode
User mode CPU utilization	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 623.

- 5 To select additional measurements, click **Add**. The UNIX Kernel Statistics dialog box opens, displaying the available measurements and server properties.



- 6 To add UNIX measurements to the monitor list, select the desired measurements, and click **OK**.
- 7 Click **OK** in the UNIX Resources 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 289.

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, LoadRunner connects to default SNMP port 161. You can also specify a machine name in the following format:
`<server name>:<port number>`

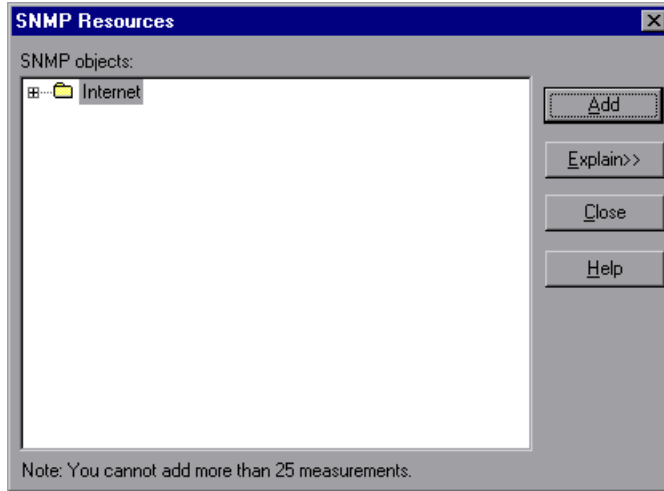
To monitor SNMP resources through a firewall, use ports 161 or 162.

To configure the SNMP Resources monitor:

- 1 Click the SNMP Resources graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the SNMP dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

- 4 In the Resource Measurements section of the SNMP dialog box, click **Add** to select the measurements that you want to monitor.

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 SNMP dialog box to activate the monitor.

You can modify the list of resources that you want to monitor at any point during the scenario. Note that a scenario 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 *<LoadRunner 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 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 the Antara FlameThrower graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.

- 3** In the Monitored Server Machines section of the Antara FlameThrower dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Enter the server name or IP address according to the following format: *<server name>:<port number>*.

For example: merc1:12135

Select the platform on which the machine runs, and click **OK**.

- 4** Click **Add** in the Resource Measurements section of the Antara FlameThrower dialog box to select the measurements that you want to monitor. 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
TxBytes	The total number of Layer 2 data bytes transmitted.
TxByteRate(/sec)	The number of Layer 2 data bytes transmitted per second.
TxFrames	The total number of packets transmitted.
TxFrameRate(/sec)	The number of packets transmitted per second.
RxBytes	The total number of Layer 2 data bytes received.
RxByteRate(/sec)	The number of Layer 2 data bytes received per second.
RxFrames	The total number of packets received.
RxFrameRate(/sec)	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 Rate(/sec)	Number of SYN ACK packets received per second.
TCPConnLatency(milise	Interval between transmitting a SYN packet and receiving a SYN ACK reply packet in msec.
MinTCPConn Latency(milise	Minimum TCPConnectionLatency in msec.
MaxTCPConn Latency(milise	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.

Measurement	Description
AvgHTTPData Latency(miliseconds)	The average HTTP Data Latency over the past second in msec.
HTTPData Latency(miliseconds)	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(miliseconds)	Minimum HTTPDataLatency in msec.
MaxHTTPData Latency(miliseconds)	Maximum HTTPDataLatency in msec.
MinData Throughput (bytes/sec)	Minimum HTTPDataThroughput in seconds.
MaxData Throughput (bytes/sec)	Maximum HTTPDataThroughput in seconds.
SuccHTTPRequests	Total number of successful HTTP Request Replies (200 OK) received.
SuccHTTPRequest Rate(/sec)	Number of successful HTTP Request Replies (200 OK) received per second.
UnSuccHTTP Requests	Number of unsuccessful HTTP Requests.

SSL/HTTPS Performance Counters

Measurement	Description
SSLConnections	Number of ClientHello messages sent by the Client.
SSLConnection Rate(/sec)	Number of ClientHello messages sent per second.
SuccSSL Connections	Number of successful SSL Connections. A successful connection is one in which the Client receives the Server's finished handshake message without any errors.

Measurement	Description
SuccSSLConnection Rate(/sec)	Number of successful SSL connections established per second.
SSLAlertErrors	Number of SSL alert messages received by the client (e.g. bad_record_mac, decryption_failed, handshake_failure, etc..)
SuccSSLResumed Sessions	Number of SSL Sessions that were successfully resumed.
FailedSSLResumed Sessions	Number of SSL Sessions that were unable to be resumed.

Sticky SLB Performance Counters

Measurement	Description
Cookie AuthenticationFail	The number of Cookie's that were not authenticated by the Server.
SuccCookie Authentication	The number of Cookie's authenticated by the server.
SSLClientHellos	The number of Client Hello packets sent to the server.
SSLServerHellos	The number of Server Hello packets sent to back to the client.
SSLSessionsFailed	The number of Session ID's that were not authenticated by the server.
SSLSessions Resumed	The number of Session ID's authenticated by the server.
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.
SuccFTPPassRate (/sec)	Number of successful FTP PASS replies received per second.
FTPControl Connections	Total number of SYN packets transmitted by the FTP client.
FTPControl ConnectionRate (/sec)	Number of SYN packets transmitted by the FTP client per second.

Measurement	Description
SuccFTPControl Connections	Total number of SYN ACK packets received by the FTP client.
SuccFTPControl ConnectionRate (/sec)	Number of SYN ACK packets received by the FTP Client per second.
FTPData Connections	Number of SYN ACK packets received by the FTP client per second.
FTPDataConnection Rate(/sec)	Number of SYN ACK packets transmitted by the FTP Client or received by the FTP Server per second.
SuccFTPData Connections	Total number of SYN ACK packets transmitted by the FTP Client or received by the FTP Server.
SuccFTPData ConnectionRate (/sec)	Number of SYN ACK packets received by the FTP server per second.
FtpAuthFailed	Total number of error replies received by the FTP client.
FTPGets	Total number of client Get requests.
FTPPuts	Total number of client Put requests.
SuccFTPGets	Total number of successful Get requests (data has been successfully transferred from server to client).
SuccFTPPuts	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 (milisechs)	Interval between transmitting a HELO packet and receiving a response in msec.
MinSMTPHelo Latency(milisechs)	Minimum SMTPHeloLatency in msec.

Measurement	Description
MaxSMTPHeloLatency(miliseconds)	Maximum SMTPHeloLatency in msec.
SuccSMTPHelos	Total number of successful HELO replies received.
SuccSMTPHeloRate(/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.
SMTPMailFromLatency(miliseconds)	Interval between transmitting a Mail From packet and receiving a response in msec.
MinSMTPMailFromLatency(miliseconds)	Minimum SMTPMailFromLatency in msec.
MaxSMTPMailFromLatency(miliseconds)	Maximum SMTPMailFromLatency in msec.
SuccSMTPMailFroms	Total number of successful Mail From replies received.
SuccSMTPMailFromRate(/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.
SMTPRcptToLatency(miliseconds)	Interval between transmitting a RcptTo packet and receiving a response in msec.
MinSMTPRcptToLatency(miliseconds)	Minimum SMTPRcptToLatency in msec.
MaxSMTPRcptToLatency(miliseconds)	Maximum SMTPRcptToLatency in msec.
SuccSMTPRcptTos	Total number of successful RcptTo replies received.
SuccSMTPRcptToRate(/sec)	Number of successful RcptTo replies received per second.

Measurement	Description
SMTPDdatas	Total number of Data packets transmitted.
SMTPDataRate(/sec)	Number of Data packets transmitted per second.
SMTPDataLatency (milisechs)	Interval between transmitting a Data packet and receiving a response in msec.
MinSMTPData Latency(milisechs)	Minimum SMTPDataLatency in msec.
MaxSMTPData Latency(milisechs)	Maximum SMTPDataLatency in msec.
SuccSMTPDdatas	Total number of successful Data replies received.
SuccSMTPDataRate (/sec)	Number of successful Data replies received per second.

POP3 Performance Counters

Measurement	Description
POP3Users	Total number of Pop3 User command packets transmitted.
POP3UserRate(/sec)	Number of Pop3 User command packets transmitted per second.
POP3UserLatency (milisechs)	Interval between transmitting a Pop3 User command packet and receiving a response in msec.
MinPOP3User Latency(milisechs)	Minimum POP3UserLatency in msec.
MaxPOP3User Latency(milisechs)	Maximum POP3UserLatency in msec.
SuccPOP3Users	Total number of successful Pop3 User replies received.
SuccPOP3UserRate (/sec)	Number of successful Pop3 User replies received per second.
POP3Passes	Total number of Pop3 Pass command packets transmitted.

Measurement	Description
POP3PassRate(/sec)	Number of Pop3 Pass command packets transmitted per second.
POP3PassLatency (milisecs)	Interval between transmitting a Pop3 Pass packet and receiving a response in msec.
MinPOP3Pass Latency(milisecs)	Minimum POP3PassLatency in msec.
MaxPOP3Pass Latency(milisecs)	Maximum POP3PassLatency in msec.
SuccPOP3Passes	Total number of successful Pop3 Pass replies received.
SuccPOP3PassRate (/sec)	Number of successful Pop3 Pass replies received per second.
POP3Stats	Total number of Pop3 Stat command packets sent.
POP3StatRate(/sec)	Number of Pop3 Stat command packets transmitted per second.
POP3StatLatency (milisecs)	Interval between transmitting a Pop3 Stat packet and receiving a response in msec.
MinPOP3Stat Latency(milisecs)	Minimum POP3StatLatency in msec.
MaxPOP3Stat Latency(milisecs)	Maximum POP3StatLatency in msec.
SuccPOP3Stats	Total number of successful Pop3 Stat replies received.
SuccPOP3StatRate (/sec)	Number of successful Pop3 Stat replies received per second.
POP3Lists	Total number of Pop3 List command packets transmitted.
POP3ListRate(/sec)	Number of Pop3 List command packets transmitted per second.
POP3ListLatency (milisecs)	Interval between transmitting a Pop3 List packet and receiving a response in msec.
MinPOP3List Latency(milisecs)	Minimum POP3ListLatency in msec.

Measurement	Description
MaxPOP3List Latency(miliseCs)	Maximum POP3ListLatency in msec.
SuccPOP3Lists	Total number of successful Pop3Lists received.
SuccPOP3ListRate (/sec)	Number of successful Pop3Lists received per second.
POP3Retrs	Total number of Pop3 Retr packets transmitted.
POP3RetrRate(/sec)	Number of Pop3 Retr packets transmitted per second.
POP3RetrLatency (miliseCs)	Interval between transmitting a Pop3 Retr packet and receiving a response in msec.
MinPOP3Retr Latency(miliseCs)	Minimum POP3RetrLatency in msec.
MaxPOP3Retr Latency(miliseCs)	Maximum POP3RetrLatency in msec.
SuccPOP3Retrs	Total number of successful Pop3Retrs received.
SuccPOP3RetrRate (/sec)	Number of successful Pop3Retrs received per second.

DNS Performance Counters

Measurement	Description
SuccPrimaryDNS Request	Total number of Successful DNS requests made to the Primary DNS server.
SuccSecondaryDNS Request	Total number of Successful DNS requests made to the Secondary DNS server.
SuccDNSData RequestRate(/sec)	Number of Successful DNS Request packets transmitted per second.
PrimaryDNSFailure	Total number of DNS requests failures received from the Primary DNS server.
PrimaryDNSRequest	Total number of DNS requests made to the Primary DNS server.

Measurement	Description
SecondaryDNS Failure	Total number of DNS requests failures received from the Secondary DNS server.
SecondaryDNS Request	Total number of DNS requests made to the Secondary DNS server.
MinDNSData Latency	Minimum DNS Data Latency in msec.
MaxDNSData Latency	Maximum DNS Data Latency in msec.
CurDNSData Latency	Interval between sending a DNS request packet and receiving a response in msec.
DNSDataRequest Rate(/sec)	Number of DNS Request packets transmitted per second.
NoOf ReTransmission	Total number of DNS Request packets re
NoOfAnswers	Total number of Answers to the DNS Request packets.

Attacks Performance Counters

Measurement	Description
Attacks	Total number of attack packets transmitted (All Attacks)
AttackRate(/sec)	Number of attack packets transmitted per second (ARP, Land, Ping, SYN, and Smurf)
Havoc Flood	Number of Havoc packets generated (Stacheldraht only)
Icmp Flood	Number of ICMP attack packets generated (TFN, TFN2K, & Stacheldraht)
Mix Flood	Number of Mix packets generated (TFN2K only)
Mstream Flood	Number of Mstream packets generated (Stacheldraht only)
Null Flood	Number of Null packets generated (Stacheldraht only)

Measurement	Description
Smurf Flood	Number of Smurf packets generated (TFN, TFN2K, & Stacheldraht)
Syn Flood	Number of SYN packets generated (TFN, TFN2K, & Stacheldraht)
Targa Flood	Number of Targa packets generated (TFN2K only)
Udp Flood	Number of UDP packets generated (All DDoS Attacks only)

- Click **OK** in the Antara FlameThrower Monitor Configuration dialog box, and in the Antara FlameThrower 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.

Before setting up the SiteScope monitor:

- Make sure that SiteScope has been installed on a server. Although you can install SiteScope on the Controller machine, we recommend installing it on a dedicated server.
- On the machine where SiteScope is installed, configure SiteScope to monitor the required servers. 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.

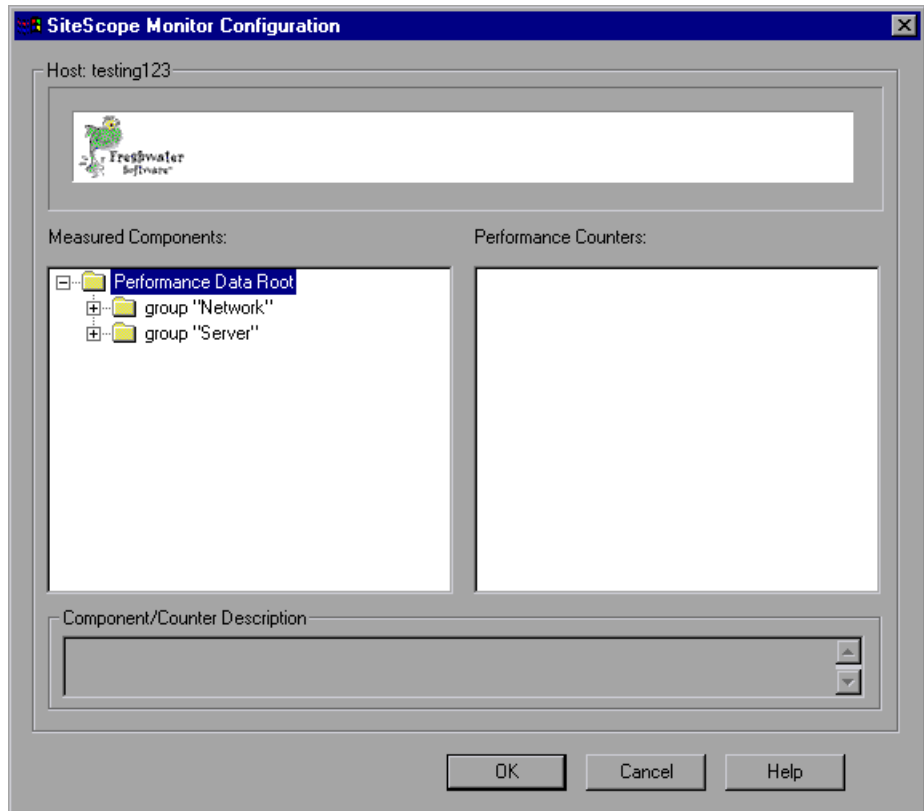
Note: SiteScope's default sampling rate is 10 minutes, and its minimum rate is 15 seconds.

- Verify that SiteScope is collecting the required data from the servers it is monitoring.

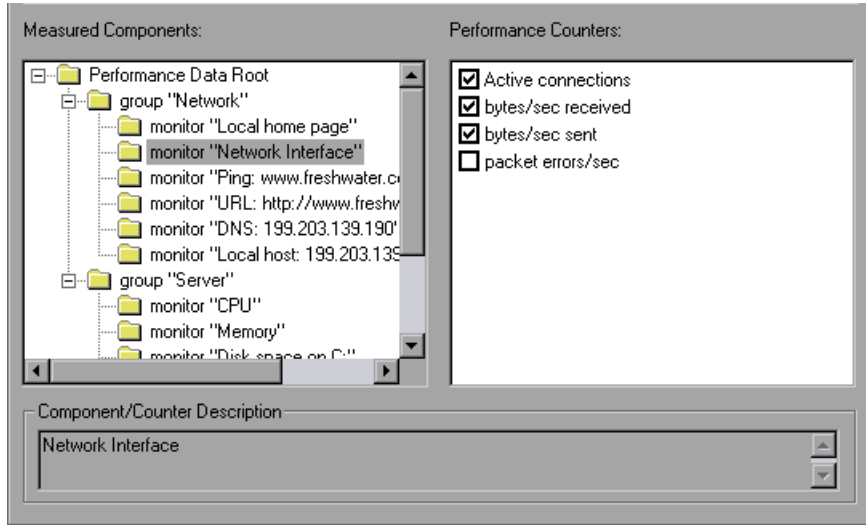
To configure the SiteScope monitor:

- 1 Click the SiteScope graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the SiteScope dialog box, click **Add** to enter the machine where SiteScope is installed. Select the platform on which the machine runs, and click **OK**.
- 4 In the Resource Measurements section of the SiteScope dialog box, click **Add** to select the measurements that you want to monitor.

The SiteScope Monitor Configuration dialog box opens.



- 5 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.



- 6 Check the required performance counters in the SiteScope Monitor Configuration window's right pane. For details of the performance counters that SiteScope can monitor, see the relevant SiteScope documentation.
- 7 Click **OK** in the SiteScope Monitor Configuration dialog box, and in the SiteScope dialog box, to activate the SiteScope monitor.

Note: SiteScope can be monitored by only one Controller at a time.

22

Network Monitoring

You can use Network monitoring to determine whether your network is causing a delay in the scenario. 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. 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 LoadRunner 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 LR_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 LoadRunner is installed on the network:

In a LoadRunner 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 *LR_installation/bin* to anywhere on the local disk of the source machine. For example, to copy the file to the */local/LR* directory, type: `cp /net/tools/LR_installation/bin/merc_webtrace /local/LR`

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/LR*), since all of them use the same *mdrv.dat*.

- 2** Add the following line to the *LR_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 LR_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 Controller 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 Controller machine, type: `cd LR_installation/bin` to change to the *bin* directory.

Note: In a network or workstation installation, *LR_installation/bin* is the location on the network in which you installed the LoadRunner 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 *LR_installation/dat/monitors/ndm.cfg* file, in the [hosts] section:

Host = <encrypted string copied from clipboard>
- 6** Close and open the current scenario. LoadRunner will read the updated configuration file and recognize the source machine for monitoring.

Connecting to the Source Machine Through the Agent

If the Controller 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 D, “Troubleshooting the Controller.”

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 *LR_installation/bin* directory.
- 2** Make sure that the agent daemon is running whenever you activate the Network monitor.
- 3** To stop the Network Delay Monitor agent daemon, type `m_daemon_setup -remove`.

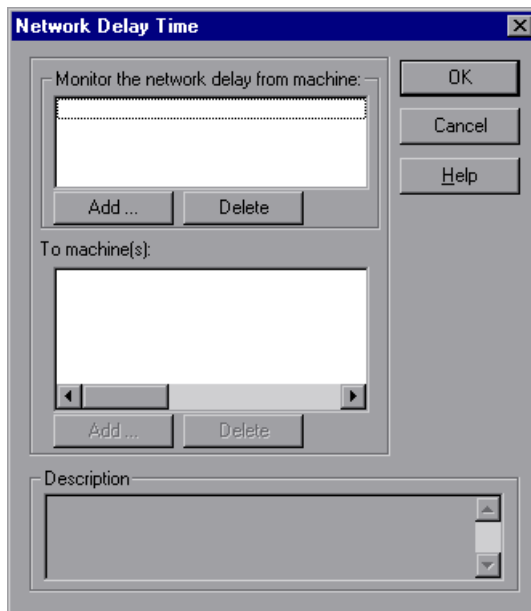
Configuring the Network Monitor

You configure the Network monitor from the Run view of the Controller before you begin running a scenario. Using the Network Delay Time and Add Destination Machines for Network Delay Monitoring dialog boxes, you select the network path you want to monitor.

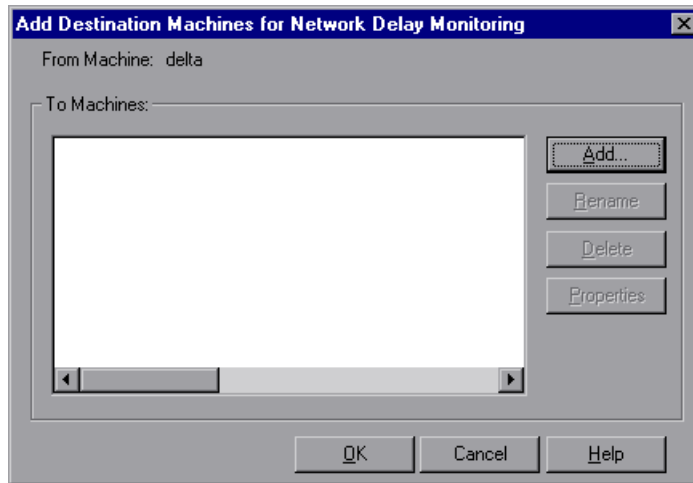
Note: To enable network monitoring, you must install the LoadRunner agent on the source machine. You do not have to install the LoadRunner agent on the destination machine.

To configure the Network monitor:

- 1** In the graph tree view, select the **Network Delay Time** graph and drag it into the right pane.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**. The Network Delay Time dialog box opens.



- 3 In the **Monitor the network delay from machine** section, click **Add** to enter the server name or IP address of the source machine, from which you want the network path monitoring to begin. Select the platform on which the machine runs, and click **OK**.
- 4 In the **To machine(s)** section of the Network Delay Time dialog box, click **Add** to enter the name of the machine at the final destination of the path you want to monitor. The Add Destination Machines for Network Delay Monitoring dialog box opens.



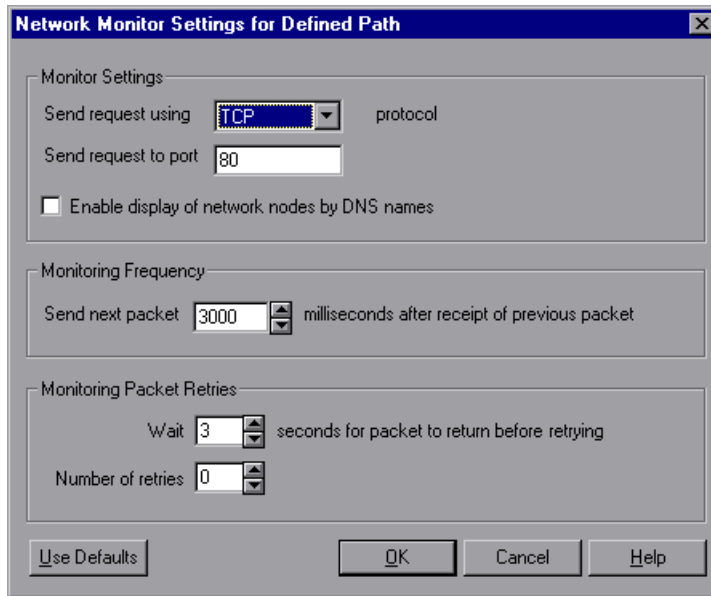
- 5 Click **Add**, enter the name of the destination machine, and click **OK**. The name of the machine appears in the Add Destination Machines for Network Delay Monitoring dialog box. Repeat this procedure for each path you want to monitor.

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**.

- 6 Click **Properties** to configure additional network monitor settings. The Network Monitor Settings for Defined Path dialog box opens.



- 7 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.
- 8 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.
- 9 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 scenario, you can increase the interval by several seconds.

- 10** 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 Controller machine and the source machine, you must configure the LoadRunner agent, MI Listener, and Network monitor for monitoring over a firewall. For more information see “Configuring the LoadRunner Agents in LAN1,” on page 199, “Installing and Configuring the MI_Listener in LAN2,” on page 208, and “Configuring the Network Delay Monitor over a Firewall,” on page 247.

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 scenario 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.

23

Firewall Server Performance Monitoring

During a scenario run, you can monitor the firewall server in order to isolate server performance bottlenecks.

This chapter describes:

- ▶ Configuring the Check Point FireWall-1 Server Monitor

About the Firewall Server Monitor

The Firewall server online monitor measures the performance of a Firewall server during scenario execution. In order to obtain performance data, you need to activate the Firewall server monitor before executing the scenario, and indicate which statistics and measurements you want to monitor.

Configuring the Check Point FireWall-1 Server Monitor

To monitor the Check Point FireWall-1 server, you must select the counters you want the Check Point FireWall-1 server monitor to measure. You select these counters using the Check Point FireWall-1 SNMP Resources dialog box.

To configure the Check Point FireWall-1 server monitor:

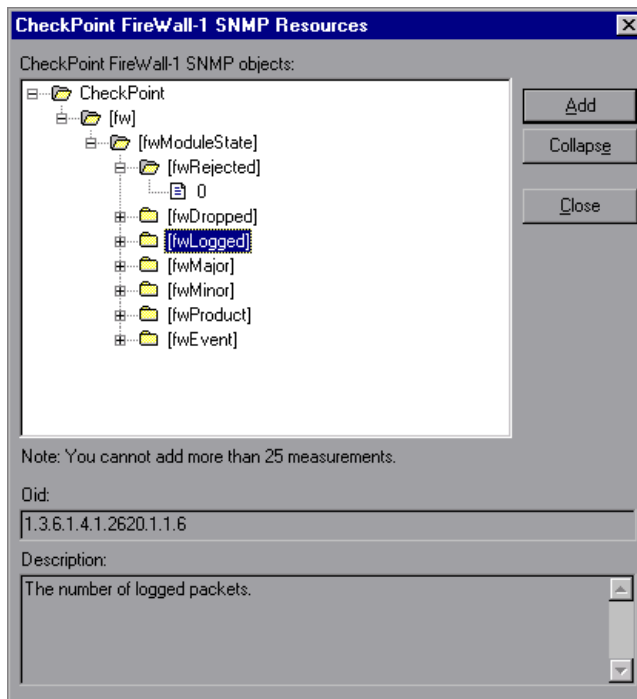
- 1** Click the Check Point FireWall-1 graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.

- 3 In the Monitored Server Machines section of the Check Point FireWall-1 dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

Note: You can specify a port number in the *snmp.cfg* file. If you do not specify a port number, LoadRunner connects to port 260, the default port for the Check Point FireWall-1 SNMP agent. You can also specify a machine name and port number in the Add Machine dialog box using the following format:

<machine name>:<port number>

- 4 Click **Add** in the Resource Measurements section of the Check Point FireWall-1 dialog box. The Check Point 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 Check Point FireWall-1 monitor can only monitor up to 25 measurements.

- 7** Click **OK** in the Check Point FireWall-1 dialog box to activate the monitor.

Note: You can improve the level of measurement information for the Check Point 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 *<LoadRunner 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 Check Point FireWall-1 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.

24

Web Server Resource Monitoring

Using LoadRunner's Web Server Resource monitors, you can monitor the Apache, Microsoft IIS, iPlanet (SNMP), and iPlanet/Netscape servers during a scenario 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 scenario 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 scenario.

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 624.

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 the Apache graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the Apache dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the Apache dialog box, click **Add** to select the measurements that you want to monitor.

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 `<LR root folder>\dat\monitors\` directory.
- 2 Edit the following parameters after the `Delimiter=:` statement:

InfoURL	server statistics information URL
ServerPort	server port number
SamplingRate	rate (milliseconds) at which the LoadRunner monitor will poll the server for the statistics information. If this value is greater than 1000, LoadRunner 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:

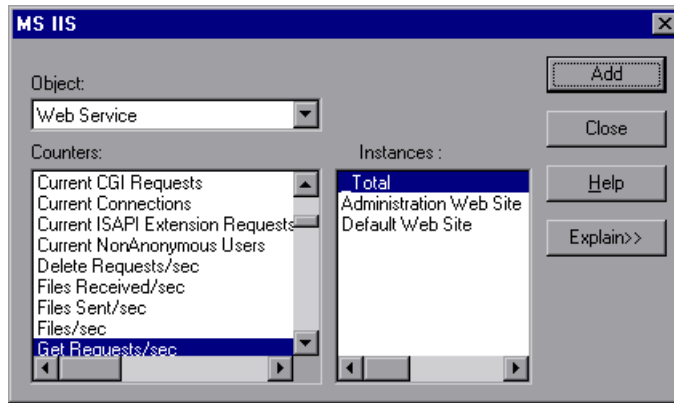
- 1** Click the MS IIS graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the MS IIS dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the MS IIS dialog box, select the measurements you want to monitor. The following table describes the default measurements that can be monitored:

Object	Measurement	Description
Web Service	Bytes Sent/sec	The rate at which the data bytes are sent by the Web service
Web Service	Bytes Received/sec	The rate at which the data bytes are received by the Web service
Web Service	Get Requests/sec	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.
Web Service	Post Requests/sec	The rate at which HTTP requests using the POST method are made. Post requests are generally used for forms or gateway requests.

Object	Measurement	Description
Web Service	Maximum Connections	The maximum number of simultaneous connections established with the Web service
Web Service	Current Connections	The current number of connections established with the Web service
Web Service	Current NonAnonymous Users	The number of users that currently have a non-anonymous connection using the Web service
Web Service	Not Found Errors/sec	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.
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 Microsoft IIS Server monitor, see “Changing a Monitor’s Default Counters” on page 623.

- 5 To select additional measurements, click **Add**. A dialog box displaying the Web Service object, its counters, and instances opens.



- 6 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.
- 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 MS IIS 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
```

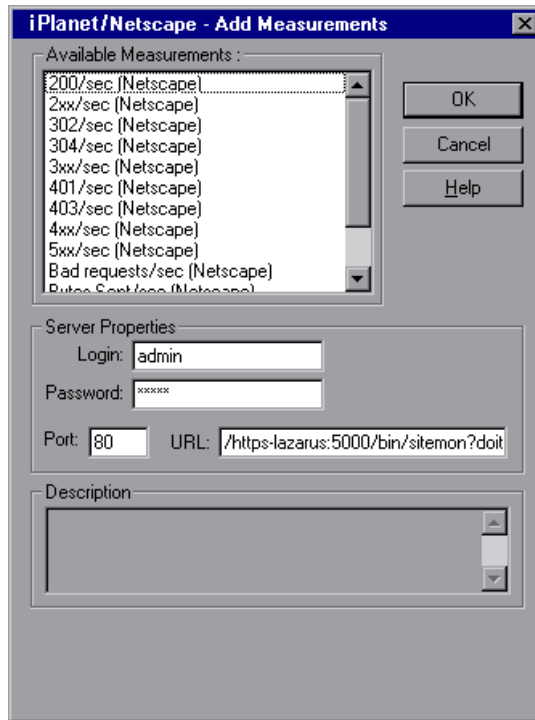
Note: In some server configurations, the URL must contain the administration server name and not the IP address.

In addition, the administration server name may differ from the iPlanet/Netscape server name.

To activate the iPlanet/Netscape monitor from the Controller:

- 1** Click the iPlanet/Netscape graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the iPlanet/Netscape dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the iPlanet/Netscape dialog box, click **Add** to select the measurements that you want to monitor.

Another 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

Measurement	Description
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 *<LR root folder>\dat\monitors* directory.
- 2 Edit the following parameters in the [Netscape] section:

Counters	number of counters that the LoadRunner iPlanet/Netscape monitor will show you. This value should match the number of counters defined in the file.
InfoURL	server statistics information URL
ServerPort	server port number
ServerLogin	login name to the server
ServerPassword	login password for the login name
SamplingRate	rate (milliseconds) at which the LoadRunner monitor will poll the server for the statistics information. If this value is greater than 1000, LoadRunner 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 the iPlanet (SNMP) graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the iPlanet (SNMP) dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

Note: You need to define the port number if the iPlanet SNMP agent is running on a different port than the default SNMP port. Enter the following information in the Add Machine dialog box:

<server name:port number>

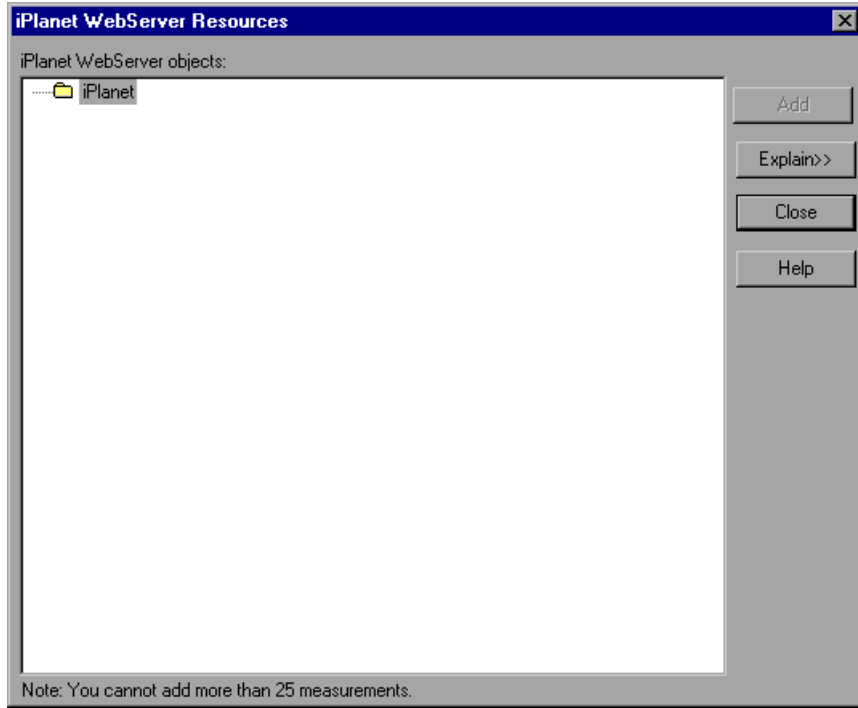
For example: digi:8888

In addition, you can define the default port for your iPlanet server in the configuration file, *snmp.cfg*, located in *<LoadRunner 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:

```
; iPlanet (WebServer)
[cm_snmp_mon_iws60]
port=8888
```

- 4 In the Resource Measurements section of the iPlanet (SNMP) dialog box, click **Add** to select the measurements that you want to monitor.

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
iwsInstanceTable	iPlanet Web Server instances
iwsInstanceEntry	iPlanet Web Server instances
iwsInstanceIndex	Server instance index
iwsInstanceId	Server instance identifier
iwsInstanceVersion	Server instance software version

Measurement	Description
iwsInstanceDescription	Description of server instance
iwsInstanceOrganization	Organization responsible for server instance
iwsInstanceContact	Contact information for person(s) responsible for server instance
iwsInstanceLocation	Location of server instance
iwsInstanceStatus	Server instance status
iwsInstanceUptime	Server instance uptime
iwsInstanceDeathCount	Number of times server instance processes have died
iwsInstanceRequests	Number of requests processed
iwsInstanceInOctets	Number of octets received
iwsInstanceOutOctets	Number of octets transmitted
iwsInstanceCount2xx	Number of 200-level (Successful) responses issued
iwsInstanceCount3xx	Number of 300-level (Redirection) responses issued
iwsInstanceCount4xx	Number of 400-level (Client Error) responses issued
iwsInstanceCount5xx	Number of 500-level (Server Error) responses issued
iwsInstanceCountOther	Number of other (neither 2xx, 3xx, 4xx, nor 5xx) responses issued
iwsInstanceCount200	Number of 200 (OK) responses issued
iwsInstanceCount302	Number of 302 (Moved Temporarily) responses issued
iwsInstanceCount304	Number of 304 (Not Modified) responses issued
iwsInstanceCount400	Number of 400 (Bad Request) responses issued
iwsInstanceCount401	Number of 401 (Unauthorized) responses issued

Measurement	Description
iwsInstanceCount403	Number of 403 (Forbidden) responses issued
iwsInstanceCount404	Number of 404 (Not Found) responses issued
iwsInstanceCount503	Number of 503 (Unavailable) responses issued
iwsInstanceLoad 1MinuteAverage	System load average for 1 minute
iwsInstanceLoad 5MinuteAverage	System load average for 5 minutes
iwsInstanceLoad 15MinuteAverage	System load average for 15 minutes
iwsInstanceNetwork InOctets	Number of octets transmitted on the network per second
iwsInstanceNetwork OutOctets	Number of octets received on the network per second
iwsVsTable	iPlanet Web Server virtual servers
iwsVsEntry	iPlanet Web Server virtual server
iwsVsIndex	Virtual server index
iwsVsId	Virtual server identifier
iwsVsRequests	Number of requests processed
iwsVsInOctets	Number of octets received
iwsVsOutOctets	Number of octets transmitted
iwsVsCount2xx	Number of 200-level (Successful) responses issued
iwsVsCount3xx	Number of 300-level (Redirection) responses issued
iwsVsCount4xx	Number of 400-level (Client Error) responses issued
iwsVsCount5xx	Number of 500-level (Server Error) responses issued

Measurement	Description
iwsVsCountOther	Number of other (neither 2xx, 3xx, 4xx, nor 5xx) responses issued
iwsVsCount200	Number of 200 (OK) responses issued
iwsVsCount302	Number of 302 (Moved Temporarily) responses issued
iwsVsCount304	Number of 304 (Not Modified) responses issued
iwsVsCount400	Number of 400 (Bad Request) responses issued
iwsVsCount401	Number of 401 (Unauthorized) responses issued
iwsVsCount403	Number of 403 (Forbidden) responses issued
iwsVsCount404	Number of 404 (Not Found) responses issued
iwsVsCount503	Number of 503 (Unavailable) responses issued
iwsProcessTable	iPlanet Web Server processes
iwsProcessEntry	iPlanet Web Server process
iwsProcessIndex	Process index
iwsProcessId	Operating system process identifier
iwsProcessThreadCount	Number of request processing threads
iwsProcessThreadIdle	Number of request processing threads currently idle
iwsProcessConnectionQueueCount	Number of connections currently in connection queue
iwsProcessConnectionQueuePeak	Largest number of connections that have been queued simultaneously
iwsProcessConnectionQueueMax	Maximum number of connections allowed in connection queue
iwsProcessConnectionQueueTotal	Number of connections that have been accepted
iwsProcessConnectionQueueOverflows	Number of connections rejected due to connection queue overflow

Measurement	Description
iwsProcessKeepalive Count	Number of connections currently in keepalive queue
iwsProcessKeepaliveMax	Maximum number of connections allowed in keepalive queue
iwsProcessSizeVirtual	Process size in kbytes
iwsProcessSizeResident	Process resident size in kbytes
iwsProcessFraction SystemMemoryUsage	Fraction of process memory in system memory
iwsListenTable	iPlanet Web Server listen sockets
iwsListenEntry	iPlanet Web Server listen socket
iwsListenIndex	Listen socket index
iwsListenId	Listen socket identifier
iwsListenAddress	Address socket is listening on
iwsListenPort	Port socket is listening on
iwsListenSecurity	Encryption support
iwsThreadPoolTable	iPlanet Web Server thread pools
iwsThreadPoolEntry	iPlanet Web Server thread pool
iwsThreadPoolIndex	Thread pool index
iwsThreadPoolId	Thread pool identifier
iwsThreadPoolCount	Number of requests queued
iwsThreadPoolPeak	Largest number of requests that have been queued simultaneously
iwsThreadPoolMax	Maximum number of requests allowed in queue
iwsCpuTable	iPlanet Web Server CPUs
iwsCpuEntry	iPlanet Web Server CPU
iwsCpuIndex	CPU index
iwsCpuId	CPU identifier

Measurement	Description
iwsCpuIdleTime	CPU Idle Time
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 iPlanet (SNMP) 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 *<LoadRunner 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

LoadRunner allows you to monitor using the Apache and Netscape monitors when there is a proxy server between the Controller 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 LoadRunner to obtain proxy settings from your Internet Explorer connection configuration, or from the proxy settings in the configuration file.

To have LoadRunner 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 LoadRunner 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 LoadRunner 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.

25

Web Application Server Resource Monitoring

You can monitor a Web application server during a scenario run and isolate application server performance bottlenecks using LoadRunner'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 scenario 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 scenario.

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.

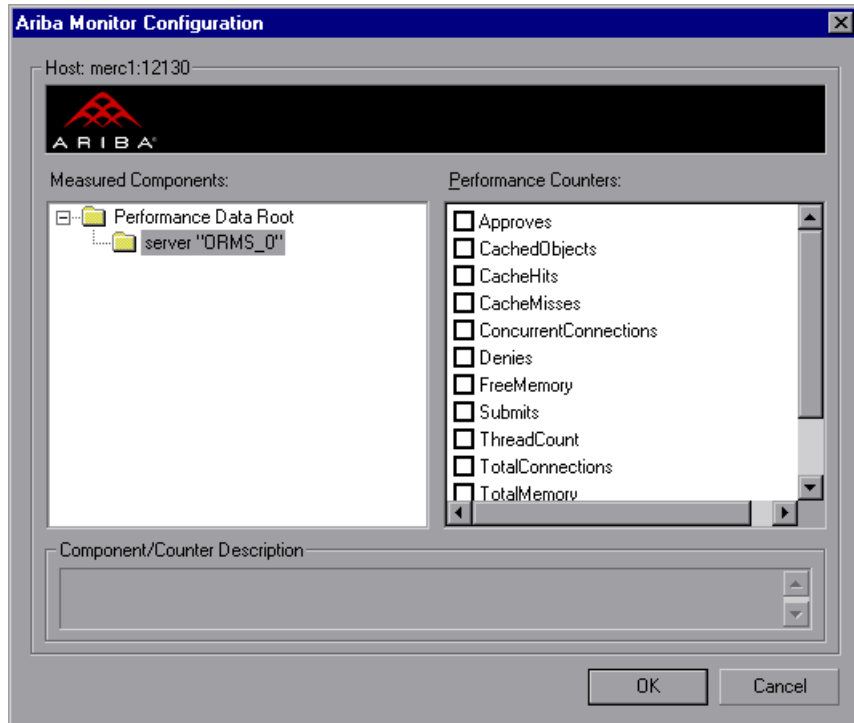
To configure the Ariba monitor:

- 1 Click the Ariba graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the Ariba dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Enter the server name or IP address according to the following format:
<server name>:<port number>.

For example: merc1:12130

Select the platform on which the machine runs, and click **OK**.

- 4 Click **Add** in the Resource Measurements section of the Ariba dialog box to select the measurements that you want to monitor. 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
Requisitions Finished	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.
Worker Queue Length	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.
Concurrent Connections	The instantaneous reading of the number of concurrent user connections at the moment this metric is obtained
Total Connections	The cumulative number of concurrent user connections since Ariba Buyer was started.
Total Memory	The instantaneous reading of the memory (in KB) being used by Ariba Buyer at the moment this metric is obtained
Free Memory	The instantaneous reading of the reserved memory (in KB) that is not currently in use at the moment this metric is obtained
Up Time	The amount of time (in hours and minutes) that Ariba Buyer has been running since the previous time it was started
Number of Threads	The instantaneous reading of the number of server threads in existence at the moment this metric is obtained
Number of Cached Objects	The instantaneous reading of the number of Ariba Buyer objects being held in memory at the moment this metric is obtained
Average Session Length	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

- 7 Click **OK** in the Ariba Monitor Configuration dialog box, and in the Ariba 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 the ATG Dynamo graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the ATG Dynamo dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

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 *<LoadRunner 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
```

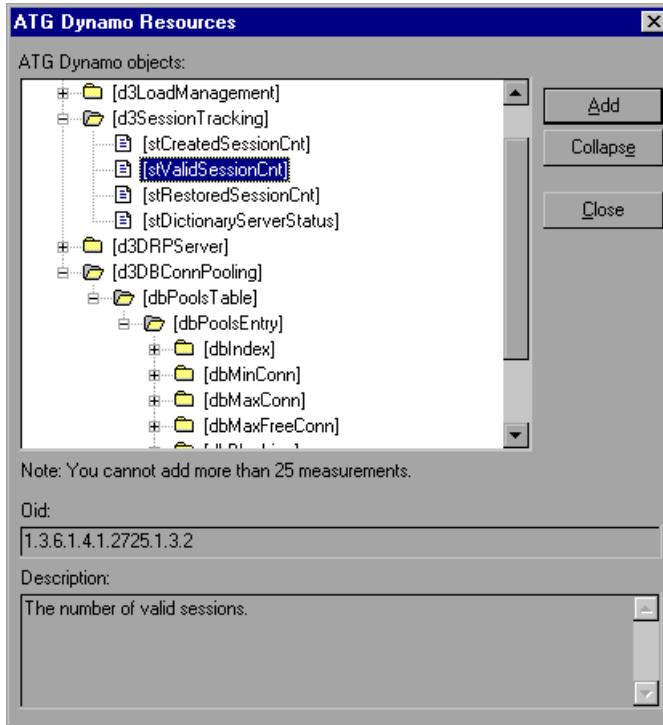
You can also specify a machine name and port number in the Add Machine dialog box using the following format:

<server name:port number>

For example: digi:8888

- 4 Click **Add** in the Resource Measurements section of the ATG Dynamo dialog box. 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
sysTotalMem	The total amount of memory currently available for allocating objects, measured in bytes
sysFreeMem	An approximation of the total amount of memory currently available for future allocated objects, measured in bytes
sysNumInfoMsgs	The number of system global info messages written

Measurement	Description
sysNumWarningMsgs	The number of system global warning messages written
sysNumErrorMsgs	The number of system global error messages written

d3LoadManagement

Measurement	Description
ImIsManager	True if the Dynamo is running a load manager
ImManagerIndex	Returns the Dynamo's offset into the list of load managing entities
ImIsPrimaryManager	True if the load manager is an acting primary manager
ImServicingCMs	True if the load manager has serviced any connection module requests in the amount of time set as the connection module polling interval
ImCMLDRPPort	The port of the connection module agent
ImIndex	A unique value for each managed entity
ImSNMPPort	The port for the entry's SNMP agent
ImProbability	The probability that the entry will be given a new session
ImNewSessions	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.
ImNewSessionOverride	The override set for whether or not a server is accepting new sessions

d3SessionTracking

Measurement	Description
stCreatedSessionCnt	The number of created sessions
stValidSessionCnt	The number of valid sessions

Measurement	Description
stRestoredSessionCnt	The number of sessions migrated to the server
StDictionaryServerStatus	d3Session Tracking

d3DRPServer

Measurement	Description
drpPort	The port of the DRP server
drpTotalReqsServed	Total number of DRP requests serviced
drpTotalReqTime	Total service time in msec for all DRP requests
drpAvgReqTime	Average service time in msec for each DRP request
drpNewessions	True if the Dynamo is accepting new sessions

d3DBConnPooling

Measurement	Description
dbPoolsEntry	A pooling service entry containing information about the pool configuration and current status
dbIndex	A unique value for each pooling service
dbPoolID	The name of the DB connection pool service
dbMinConn	The minimum number of connections pooled
dbMaxConn	The maximum number of connections pooled
dbMaxFreeConn	The maximum number of free pooled connections at a time
dbBlocking	Indicates whether or not the pool is to block out check outs
dbConnOut	Returns the number of connections checked out

Measurement	Description
dbFreeResources	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.
dbTotalResources	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.


- 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 ATG Dynamo monitor can only monitor up to 25 measurements.

- 7** Click **OK** in the ATG Dynamo 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 *<LoadRunner 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 Controller 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

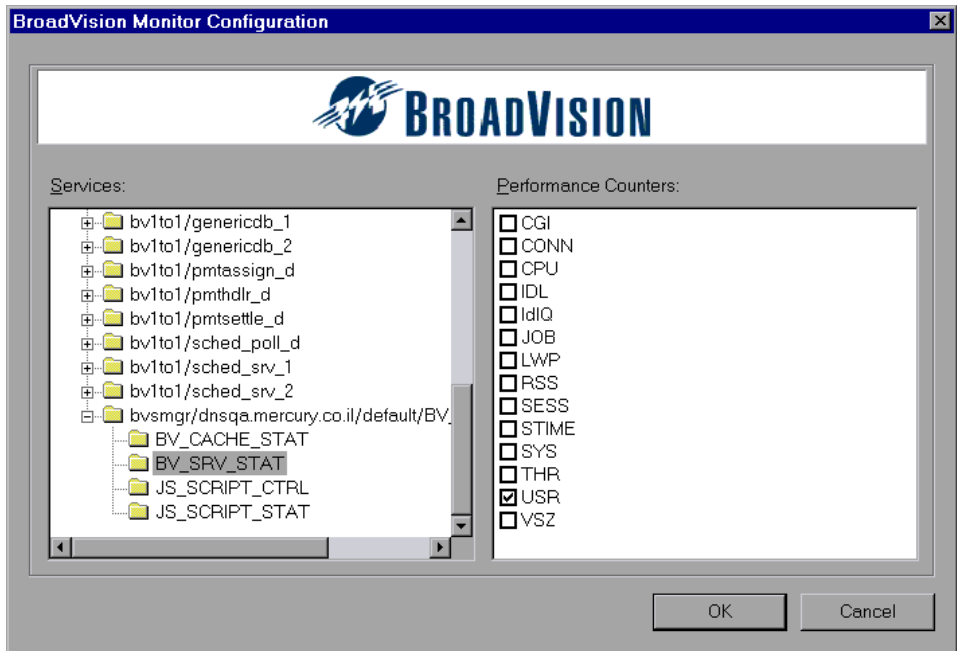
- 3** The registry should contain the correct path to the Java run-time environment (JRE) under the JRE 1.2 installation directory. Verify the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Java Runtime Environment\1.2\JavaHome

To configure the BroadVision online monitor:

- 1** Right-click the graph in the graph view area and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 2** In the Monitored Server Machines section of the BroadVision dialog box, click **Add** to enter the BroadVision server name or IP address with the port number according to the following format: *<server name>:<port number>*. For example: dnsqa:1221. Select the machine platform, and click **OK**.
- 3** Click **Add** in the Resource Measurements section of the BroadVision dialog box.

The BroadVision Monitor Configuration dialog box opens, displaying the available measurements:



- 4 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.
- 5 Click **OK** in the BroadVision Monitor Configuration dialog box, and in the BroadVision 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
bvconf_srv	No	One-To-One configuration management server. There must be one.
cmsdb	Yes	Visitor management database server.
cntdb	Yes	Content database server.
deliv_smtp_d	Yes	Notification delivery server for e-mail type messages. Each instance of this server must have its own ID, numbered sequentially starting with "1".
deliv_comp_d	No	Notification delivery completion processor.
extdbacc	Yes	External database accessor. You need at least one for each external data source.
genericdb	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.
hostmgr	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.
g1_ofbe_srv	No	Order fulfillment back-end server.
g1_ofdb	Yes	Order fulfillment database server.
g1_om_srv	No	Order management server.
pmtassign_d	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>pmtldlr_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
- Rebind
- Rsvl
- Unbind

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:

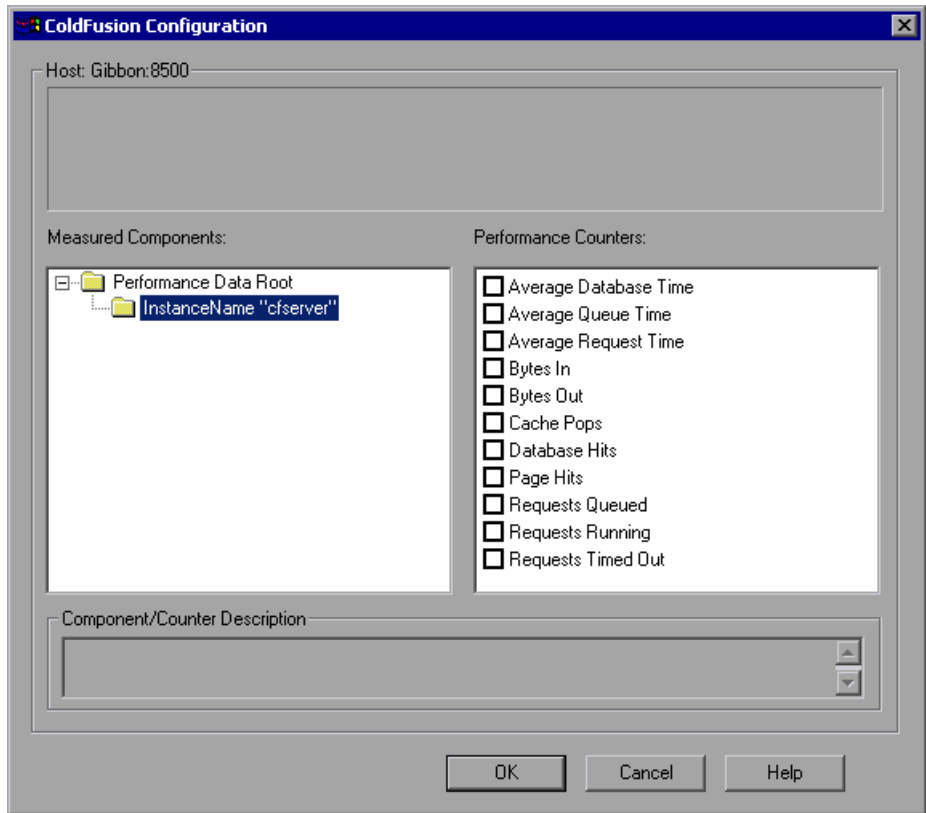
Copy the `<LR installation>\dat\monitors\perfmon.cfm` file into the `<ColdFusion Home>\cfide\administrator` directory. By default, the ColdFusion monitor checks for the `<ColdFusion Home>\cfide\administrator\perfmon.cfm` file.

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 the ColdFusion graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the ColdFusion dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4 In the Resource Measurements section of the ColdFusion dialog box, click **Add** to select the measurements that you want to monitor. The ColdFusion Monitor Configuration dialog box displays the available measurements.

5 Browse the Measured Components tree.



6 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
Avg. Database Time (msec)	The running average of the amount of time, in milliseconds, that it takes ColdFusion to process database requests.
Avg. Queue Time (msec)	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.
Avg Req Time (msec)	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.
Bytes In/sec	The number of bytes per second sent to the ColdFusion server.
Bytes Out/sec	The number of bytes per second returned by the ColdFusion server.
Cache Pops	Cache pops.
Database Hits/sec	This is the number of database hits generated per second by the ColdFusion server.
Page Hits/sec	This is the number of Web pages processed per second by the ColdFusion server.
Queued Requests	The number of requests currently waiting to be processed by the ColdFusion server.
Running Requests	The number of requests currently being actively processed by the ColdFusion server.
Timed Out Requests	The number of requests that timed out due to inactivity timeouts.

- 7 Click **OK** in the ColdFusion Monitor Configuration dialog box, and in the ColdFusion 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 the Fujitsu INTERSTAGE graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the Fujitsu INTERSTAGE dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

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. Enter the following information in the Add Machine dialog box:

<server name:port number>

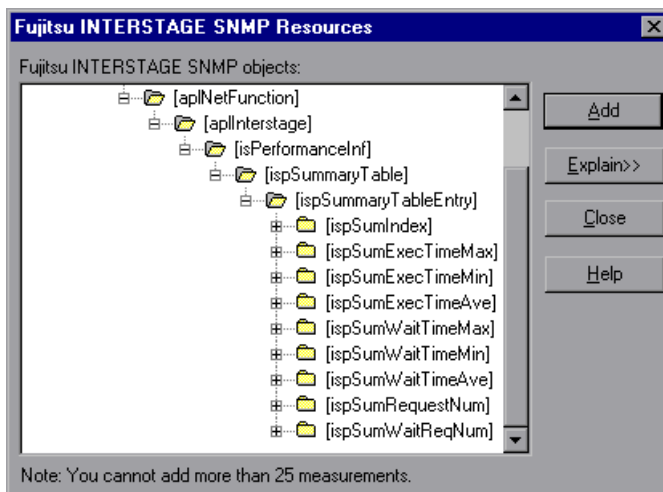
For example: digi:8888

In addition, you can define the default port for your Fujitsu INTERSTAGE server in the configuration file, *snmp.cfg*, located in *<LoadRunner 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
```

-
- 4** Click **Add** in the Resource Measurements section of the Fujitsu INTERSTAGE dialog box. The Fujitsu INTERSTAGE SNMP Resources dialog box opens.

- 5 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
IspSumObjectName	The object name of the application for which performance information is measured
IspSumExecTimeMax	The maximum processing time of the application within a certain period of time
IspSumExecTimeMin	The minimum processing time of the application within a certain period of time
IspSumExecTimeAve	The average processing time of the application within a certain period of time
IspSumWaitTimeMax	The maximum time required for INTERSTAGE to start an application after a start request is issued
IspSumWaitTimeMin	The minimum time required for INTERSTAGE to start an application after a start request is issued
IspSumWaitTimeAve	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 Fujitsu INTERSTAGE 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 Controller

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 the iPlanet (NAS) graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the iPlanet (NAS) dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

Note: You need to define the port number if the iPlanet SNMP agent is running on a different port than the default SNMP port. Enter the following information in the Add Machine dialog box:

<server name:port number>

For example: digi:8888

In addition, you can define the default port for your iPlanet server in the configuration file, *snmp.cfg*, located in *<LoadRunner 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:

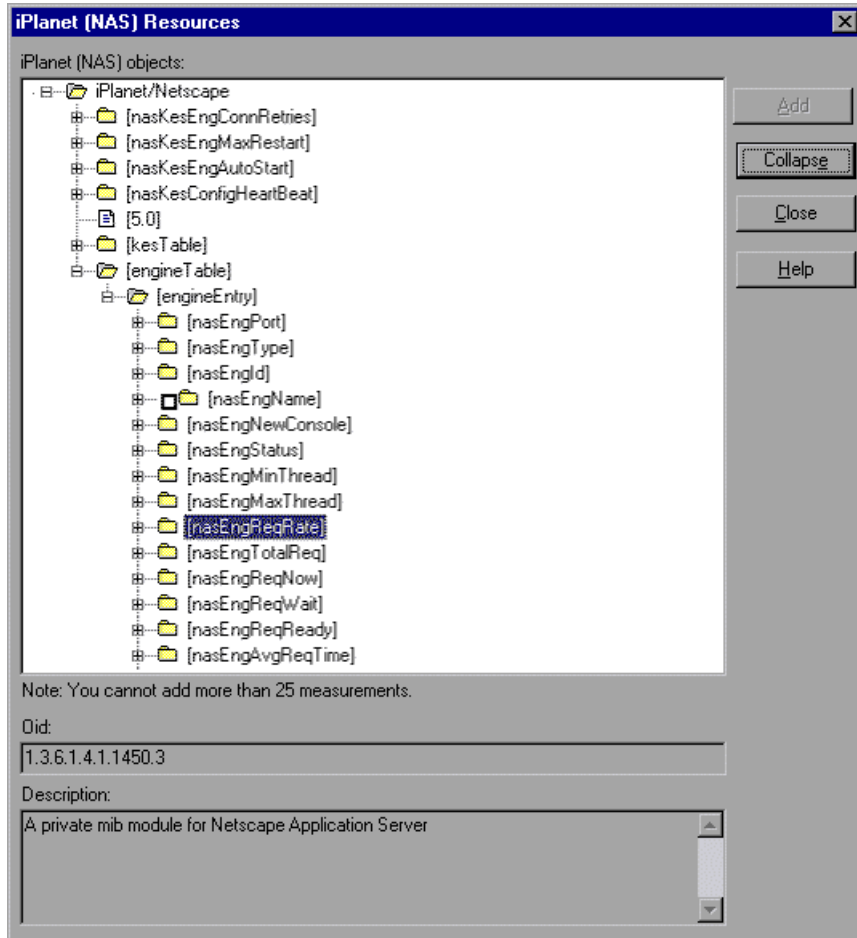
```
; iPlanet (NAS)
```

```
[cm_snmp_mon_nas]
```

```
port=8888
```

- 4** In the Resource Measurements section of the iPlanet (NAS) dialog box, click **Add** to select the measurements that you want to monitor.

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
nasKesEngConnRetries	The maximum number of times the administration server will try to connect to an engine.
nasKesEngMaxRestart	The maximum number of times the administration server will restart an engine after a failure.
nasKesEngAutoStart	Start all the engines at startup of the administration server.
nasKesConfigHeartBeat	Heart Beat.

KES Performance Counters

Measurement	Description
nasKesId	The ID of the KES this engine belongs to.
nasKesMinThread	The default minimum number of threads per engine.
nasKesMaxThread	The default maximum number of threads per engine.
nasKesLoadBalancerDisable	Enable or Disable the load balancer service.
nasKesCpuLoad	The total CPU usage on this host.
nasKesDiskLoad	The total disk usage on this host.
nasKesMemLoad	The total memory usage on this host.
nasKesRequestLoad	The number of requests on this NAS.
nasKesCpuLoadFactor	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
nasKesDiskLoadFactor	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%.
nasKesMemLoadFactor	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%.
nasKesAppLogicsRunningFactor	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%.
nasKesResultsCacheFactor	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%.
nasKesAvgExecTimeFactor	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%.
nasKesLastExecTimeFactor	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
nasKesHitsFactor	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%.
nasKesServerLoadFactor	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%.
nasKesBroadcastInterval	The length of time in seconds, between each broadcast attempt from the load balancer daemon.
nasKesApplogicBroadcastInterval	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.
nasKesServerBroadcastInterval	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.
nasKesServerLoadUpdateInterval	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.
nasKesCpuLoadUpdateInterval	The length of time, in seconds, between each sampling of CPU usage.
nasKesDiskLoadUpdateInterval	The length of time, in seconds, between each sampling of disk usage.
nasKesMemLoadUpdateInterval	The length of time, in seconds, between each sampling of memory thrashes.
nasKesTotalReqsUpdateInterval	The length of time, in seconds, between each sampling of the number of requests.
nasKesMaxHops	The maximum number of times a request can be load-balanced to another server.

Measurement	Description
nasKesODBCReqMinThread	The minimum number of threads reserved to process asynchronous requests.
nasKesODBCReqMaxThread	The maximum number of threads reserved to process asynchronous requests.
nasKesODBCCacheMaxConns	The maximum number of connections opened between NAS and the database.
nasKesODBCCacheFreeSlots	The minimum number of cached connections established between NAS and the database.
nasKesODBCCacheTimeout	The time after which an idle connection is dropped.
nasKesODBCCacheInterval	The interval in seconds at which the cache cleaner will try to disconnect connections already idle for longer than the specified timeout.
nasKesODBCConnGiveupTime	Maximum time the driver will try to connect to the database.
nasKesODBCCacheDebug	Turns on the connection cache debug information.
nasKesODBCResultSetInitRows	The number of rows fetched at once from the database.
nasKesODBCResultSetMaxRows	The maximum number of rows the cached result set can contain.
nasKesODBCResultSetMaxSize	The maximum size of result set the driver will cache
nasKesODBCSqlDebug	Turns on SQL debug information.
nasKesODBCEnableParser	Turns on SQL parsing.
nasKesORCLReqMinThread	The minimum number of threads reserved to process asynchronous requests.
nasKesORCLReqMaxThread	The maximum number of threads reserved to process asynchronous requests.

Measurement	Description
nasKesORCLCacheMaxConns	The maximum number of connections opened between NAS and the database.
nasKesORCLCacheFreeSlots	The minimum number of cached connections established between NAS and the database.
nasKesORCLCacheTimeout	The time after which an idle connection is dropped.
nasKesORCLCacheInterval	The interval in seconds at which the cache cleaner will try to disconnect connections already idle for longer than the specified timeout.
nasKesORCLConnGiveupTime	The maximum time the driver will spend trying to obtain a connection to Oracle.
nasKesORCLCacheDebug	Turns on the connection cache debug information.
nasKesORCLResultSetInitRows	The number of rows fetched at once from the database.
nasKesORCLResultSetMaxRows	The maximum number of rows the cached result set can contain.
nasKesORCLResultSetMaxSize	The maximum size of result set the driver will cache.
nasKesORCLSqlDebug	Turns on SQL debug information.
nasKesSYBReqMinThread	The minimum number of threads reserved to process asynchronous requests.
nasKesSYBReqMaxThread	The maximum number of threads reserved to process asynchronous request.
nasKesSYBCacheMaxConns	The maximum number of connections opened between NAS and the database.
nasKesSYBCacheFreeSlots	The minimum number of cached connections established between NAS and the database.
nasKesSYBCacheTimeout	The time after which an idle connection is dropped.

Measurement	Description
nasKesSYBCacheInterval	The interval time between cached connections.
nasKesSYBConnGiveupTime	The maximum time the driver will spend trying to obtain a connection to Sybase before giving up.
nasKesSYBCacheDebug	Turns on the connection cache debug information.
nasKesSYBResultSetInitRows	The number of rows fetched at once from the database.
nasKesSYBResultSetMaxRows	The maximum number of rows the cached result set can contain.
nasKesSYBResultSetMaxSize	The maximum size of result set the driver will cache.

Engine Performance Counters

Measurement	Description
nasEngKesPort	The port of the KXS this engine serves. This is supplied as part of the object ID and cannot be modified after creation.
nasEngPort	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.
nasEngType	Type of the engine: executive(0), Java(1000), C++(3000).
nasEngId	The ID is an incremental number starting at 0. The ID cannot be modified.
nasEngName	The name of this engine. This is an informational string that contains kcs, kxs or kjs.
nasEngNewConsole	Starts each engine in a new console window.
nasEngStatus	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
nasEngMinThread	The default minimum number of threads per engine.
nasEngMaxThread	The default maximum number of threads per engine.
nasEngReqRate	The rate at which requests arrive.
nasEngTotalReq	The total number of requests processed since engine startup.
nasEngReqNow	The number of requests being processed.
nasEngReqWait	The requests waiting to be serviced.
nasEngReqReady	The requests that are ready to be serviced.
nasEngAvgReqTime	The average request processing time.
nasEngThreadNow	Number of threads in use by the request manager.
nasEngThreadWait	The number of idle threads.
nasEngWebReqQueue	The number of web requests that are queued.
nasEngFailedReq	The number of requests that failed.
nasEngTotalConn	The total number of connections opened.
nasEngTotalConnNow	The total number of connections in use.
nasEngTotalAccept	The total number of connections listening to incoming requests.
nasEngTotalAcceptNow	The total number of connections listening to incoming connections in use.
nasEngTotalSent	The total number of packets sent.
nasEngTotalSentBytes	The total number of bytes sent.
nasEngTotalRecv	The total number of packets received.
nasEngTotalRecvBytes	The total number of bytes received.
nasEngBindTotal	The number of AppLogic bound since startup.

Measurement	Description
nasEngBindTotalCached	The number of AppLogic cached since startup.
nasEngTotalThreads	Total number of threads created in this process.
nasEngCurrentThreads	Total number of threads in use in this process.
nasEngSleepingThreads	Number of threads sleeping in this process.
nasEngDAETotalQuery	Total number of queries executed since startup.
nasEngDAEQueryNow	The number of queries being processed.
nasEngDAETotalConn	The number of logical connections created since startup.
nasEngDAEConnNow	The number of logical connections in use.
nasEngDAECacheCount	The number of caches.
nasEngODBCQueryTotal	Total number of queries executed since startup.
nasEngODBCPreparedQueryTotal	Total number of odbc prepared queries executed since startup.
nasEngODBCConnTotal	Total number of connections opened since startup.
nasEngODBCConnNow	Number of connections currently opened.
nasEngORCLQueryTotal	Total number of queries executed since startup.
nasEngORCLPreparedQueryTotal	Total number of prepared queries executed since startup.
nasEngORCLConnTotal	Total number of connections established with Oracle since startup.

Measurement	Description
nasEngORCLConnNow	Number of connections opened with Oracle now.
nasEngSYBQueryTotal	Total number of queries the driver processed since startup.
nasEngSYBPreparedQueryTotal	Total number of prepared queries processed since startup.
nasEngSYBConnTotal	Total number of connections opened since startup.
nasEngSYBConnNow	Number of SYB connections opened now.
nasStatusTrapEntry	The KES definition.
nasTrapKesIpAddress	The IP Address of KES host.
nasTrapKesPort	The port of the main engine of this NAS.
nasTrapEngPort	The port of the engine generating this event.
nasTrapEngState	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 iPlanet (NAS) 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 *<LoadRunner 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:

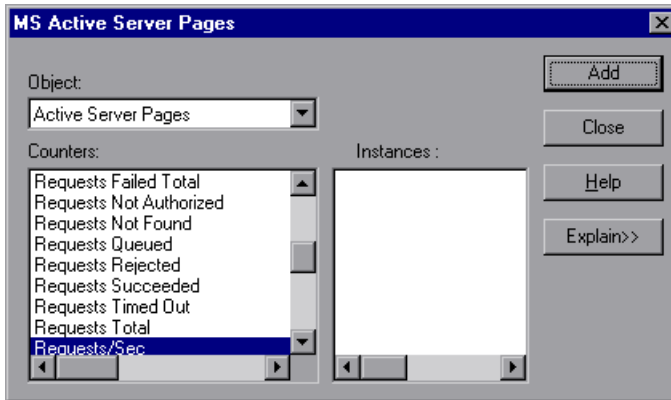
- 1** Click the MS Active Server Pages graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the MS Active Server Pages dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the MS Active Server Pages dialog box, select the measurements you want to monitor. 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.

Measurement	Description
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 623.

- To select additional measurements, click **Add**. A dialog box displaying the Active Server Pages object, its counters, and instances opens.



- 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 MS Active Server Pages dialog box to activate the monitor.

Configuring the Oracle9iAS HTTP Monitor

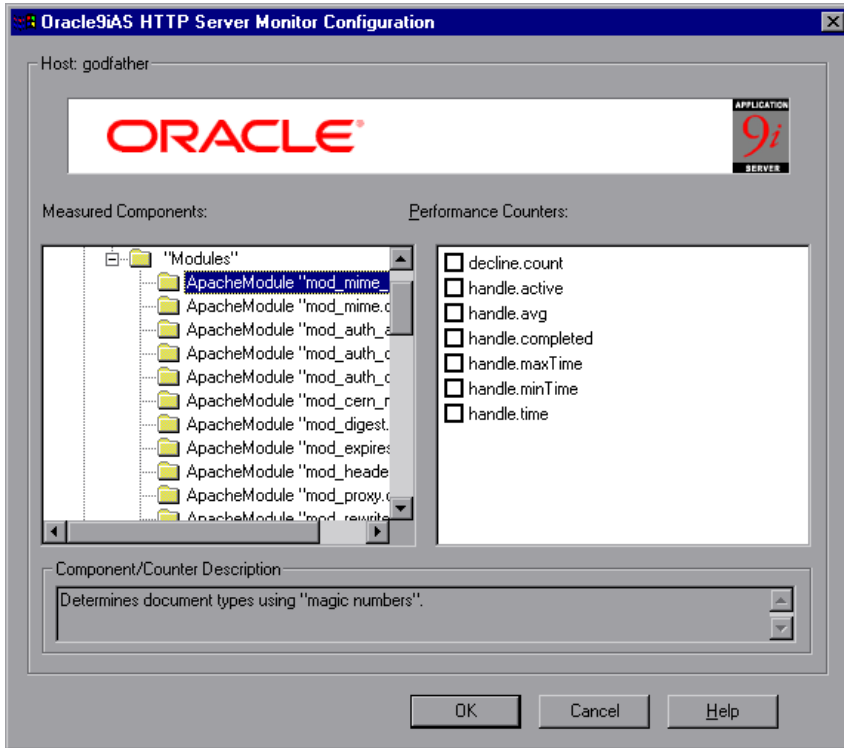
You select measurements to monitor the Oracle9iAS HTTP server using the Oracle HTTP Server Monitor Configuration dialog box. Note that you must start running the Oracle9iAS HTTP server before you begin selecting the measurements you want to monitor.

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 the Oracle9iAS HTTP graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the Oracle9iAS HTTP Server dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select any platform, and click **OK**.
- 4 Click **Add** in the Resource Measurements section of the Oracle9iAS HTTP Server dialog box to select the measurements that you want to monitor. The Oracle 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
mod_mime.c	Determines document types using file extensions
mod_mime_magic.c	Determines document types using "magic numbers"
mod_auth_anon.c	Provides anonymous user access to authenticated areas
mod_auth_dbm.c	Provides user authentication using DBM files
mod_auth_digest.c	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
mod_alias.c	Maps different parts of the host file system in the document tree, and redirects URLs
mod_userdir.c	Handles user home directories
mod_cgi.c	Invokes CGI scripts
mod_dir.c	Handles the basic directory
mod_autoindex.c	Provides automatic directory listings
mod_include.c	Provides server-parsed documents
mod_negotiation.c	Handles content negotiation
mod_auth.c	Provides user authentication using text files
mod_access.c	Provides access control based on the client hostname or IP address
mod_so.c	Supports loading modules (.so on UNIX, .dll on Win32) at run-time
mod_oprocmgr.c	Monitors JServ processes and restarts them if they fail
mod_jserv.c	Routes HTTP requests to JServ server processes. Balances load across multiple JServs by distributing new requests in round-robin order
mod_ose.c	Routes requests to the JVM embedded in Oracle's database server
http_core.c	Handles requests for static Web pages

The following table describes the counters that can be monitored:

Measurement	Description
handle.minTime	The minimum time spent in the module handler
handle.avg	The average time spent in the module handler
handle.active	The number of threads currently in the handle processing phase

Measurement	Description
handle.time	The total amount of time spent in the module handler
handle.completed	The number of times the handle processing phase was completed
request.maxTime	The maximum amount of time required to service an HTTP request
request.minTime	The minimum amount of time required to service an HTTP request
request.avg	The average amount of time required to service an HTTP request
request.active	The number of threads currently in the request processing phase
request.time	The total amount of time required to service an HTTP request
request.completed	The number of times the request processing phase was completed
connection.maxTime	The maximum amount of time spent servicing any HTTP connection
connection.minTime	The minimum amount of time spent servicing any HTTP connection
connection.avg	The average amount of time spent servicing HTTP connections
connection.active	The number of connections with currently open threads
connection.time	The total amount of time spent servicing HTTP connections
connection.completed	The number of times the connection processing phase was completed
numMods.value	The number of loaded modules
childFinish.count	The number of times the Apache parent server started a child server, for any reason

Measurement	Description
childStart.count	The number of times “children” finished “gracefully.” There are some ungraceful error/crash cases that are not counted in childFinish.count
Decline.count	The number of times each module declined HTTP requests
internalRedirect.count	The number of times that any module passed control to another module using an “internal redirect”
cpuTime.value	The total CPU time utilized by all processes on the Apache server (measured in CPU milliseconds)
heapSize.value	The total heap memory utilized by all processes on the Apache server (measured in kilobytes)
pid.value	The process identifier of the parent Apache process
upTime.value	The amount of time the server been running (measured in milliseconds)

- 7 Click **OK** in the Oracle HTTP Server Monitor Configuration dialog box, and in the Oracle9iAS HTTP Server 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`

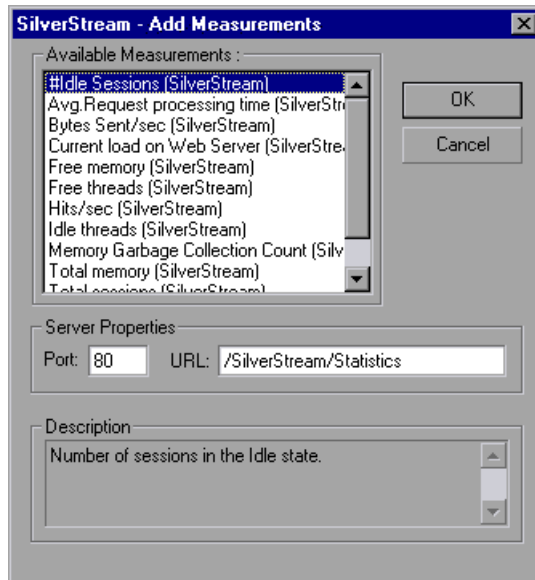
for example:

`http://199.203.78.57:80/SilverStream/Statistics`

To configure the SilverStream monitor:

- 1** Click the SilverStream graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the SilverStream dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the SilverStream dialog box, click **Add** to select the measurements that you want to monitor.

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 SilverStream 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 `<LR root folder>\dat\monitors\` directory.
- 2 Edit the following parameters at the end of the file:

InfoURL	server statistics information URL
ServerPort	server port number
SamplingRate	rate (milliseconds) at which the LoadRunner monitor will poll the server for the statistics information. If this value is greater than 1000, LoadRunner 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 the WebLogic (SNMP) graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the WebLogic (SNMP) dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

Note: You need to define the port number if the WebLogic SNMP agent is running on a different port than the default SNMP port. Enter the following information in the Add Machine dialog box:

<server name:port number>

For example: digi:8888

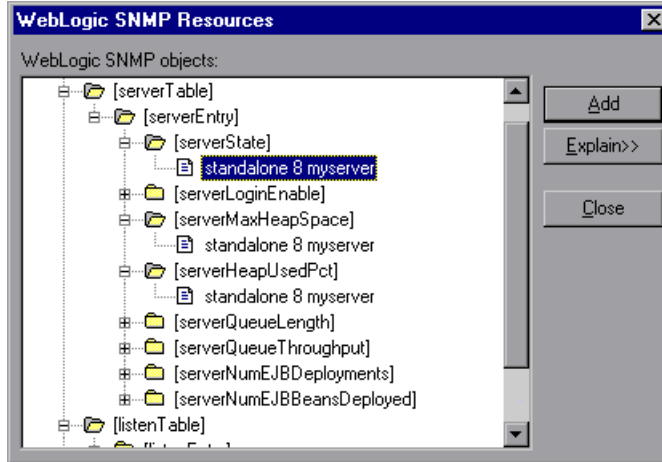
In addition, you can define the default port for your WebLogic server in the configuration file, *snmp.cfg*, located in *<LoadRunner root folder>\dat\monitors*. For example, if the port used by the SNMP agent on your WebLogic server is 8888, you should edit the *snmp.cfg* file as follows:

```
; WebLogic  
[cm_snmp_mon_isp]  
port=8888
```

-
- 4** In the Resource Measurements section of the WebLogic (SNMP) dialog box, click **Add** to select the measurements that you want to monitor. The WebLogic SNMP Resources dialog box displays the available measurements.

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
ServerState	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.
ServerLoginEnable	This value is true if client logins are enabled on the server.
ServerMaxHeapSpace	The maximum heap size for this server, in KB
ServerHeapUsedPct	The percentage of heap space currently in use on the server

Measurement	Description
ServerQueueLength	The current length of the server execute queue
ServerQueueThroughput	The current throughput of execute queue, expressed as the number of requests processed per second
ServerNumEJBDeployment	The total number of EJB deployment units known to the server
ServerNumEJBBeansDeployed	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
ListenPort	Port number.
ListenAdminOK	True if admin requests are allowed on this (protocol, ipAddr, port); otherwise false
ListenState	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
CPType	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.
CPIndex	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 WebLogic (SNMP) 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 Controller 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 *<LR 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 *<LR root folder>\classes*.
- 2** If the classes file is not located in the default *<LR root folder>* folder, you need to specify the full path to it in the *<LR 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:

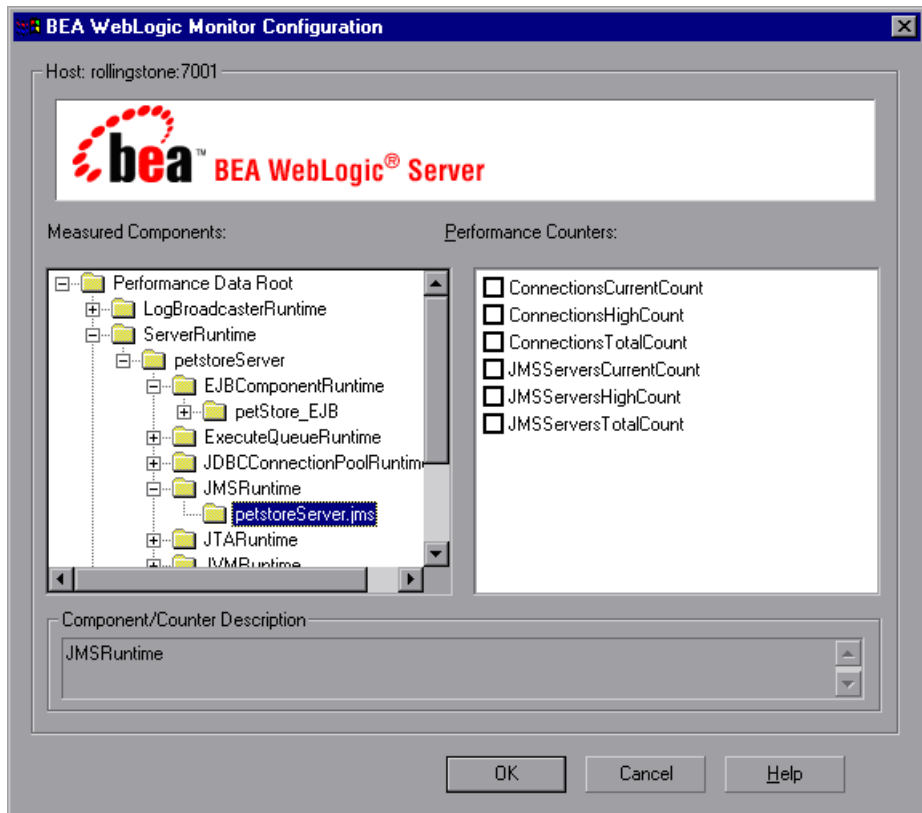
- 1** Click the WebLogic (JMX) graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the WebLogic (JMX) dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Enter the server name or IP address according to the following format: *<server name>:<port number>*.

For example: mercury:8111

Select the platform on which the machine runs, and click **OK**.

- 4** Click **Add** in the Resource Measurements section of the WebLogic (JMX) dialog box. 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 402.

5 Browse the Measured Components tree.



6 Check the required performance counters in the BEA WebLogic Monitor Configuration window's right pane.

7 Click **OK** in the BEA WebLogic Monitor Configuration dialog box, and in the WebLogic (JMX) 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).

- ServletRuntime
- WebAppComponentRuntime
- EJBStatefulHomeRuntime
- JTARuntime
- JVMRuntime
- EJBEntityHomeRuntime.
- DomainRuntime
- EJBComponentRuntime
- DomainLogHandlerRuntime
- JDBCConnectionPoolRuntime
- ExecuteQueueRuntime
- ClusterRuntime
- JMSRuntime
- TimeServiceRuntime

- EJBStatelessHomeRuntime
- WLECConnectionServiceRuntime

ServerSecurityRuntime

Measurement	Description
UnlockedUsersTotalCount	Returns the number of times a user has been unlocked on the server
InvalidLoginUsersHighCount	Returns the high-water number of users with outstanding invalid login attempts for the server
LoginAttemptsWhileLockedTotalCount	Returns the cumulative number of invalid logins attempted on the server while the user was locked
Registered	Returns "false" if the MBean represented by this object has been unregistered.
LockedUsersCurrentCount	Returns the number of currently locked users on the server
CachingDisabled	Private property that disables caching in proxies.
InvalidLoginAttemptsTotalCount	Returns the cumulative number of invalid logins attempted on the server
UserLockoutTotalCount	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 LoadRunner 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 `<LR root folder>\dat\monitors\xmlmonitorshared.ini` file and add the following line to the [WebSphereMonitor] section:

`QueryLoginInfo=1`

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:

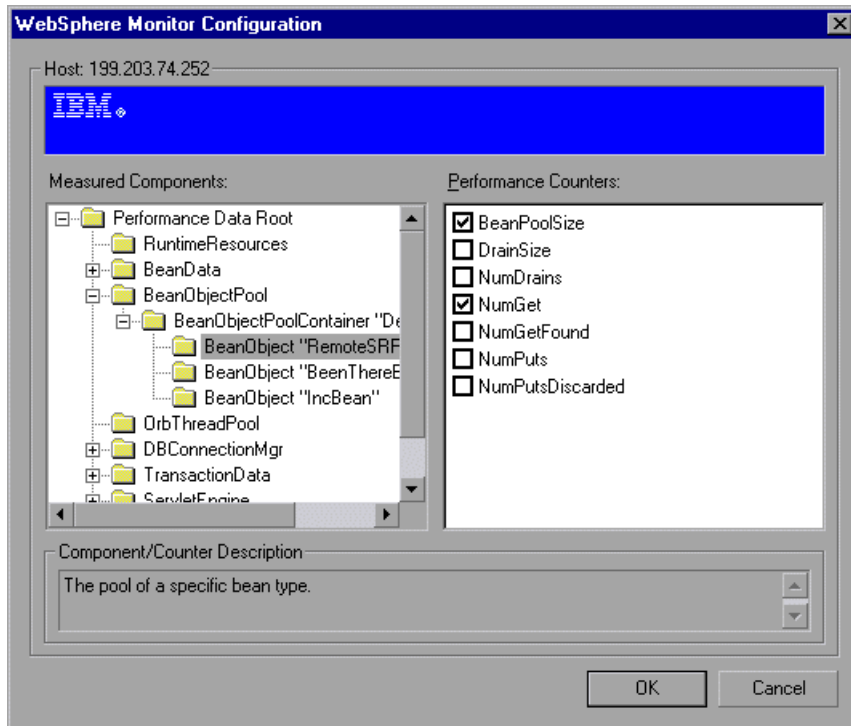
`http://<server name:port number>/<servlet_folder>/com.ibm.ivb.epm.servlet.
PerformanceServlet`

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 the WebSphere or WebSphere 4.x - 5.x graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the WebSphere dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4 In the Resource Measurements section of the WebSphere dialog box, click **Add** to select the measurements that you want to monitor. 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 413.
- 7 Click **OK** in the WebSphere Monitor Configuration dialog box, and in the WebSphere 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, LoadRunner 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`, LoadRunner 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.EPMAApplicationListener
com.ibm.servlet.debug.OLTServletManager
listeners.application=
```

Should be:

```
listeners.application=com.ibm.servlet.engine.EPMAApplicationListener
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
BeanPassivates	The number of bean passivates pertaining to a specific bean. Applies to an individual bean that is either 'stateful' or 'entity'
EntityBeanPassivates	The number of entity bean passivates
StatefulBeanPassivates	The number of stateful bean passivates
BeanActivates	The number of bean activates pertaining to a specific bean. Applies to an individual bean that is either 'stateful' or 'entity'
EntityBeanActivates	The number of entity bean activates
StatefulBeanActivates	The number of stateful bean activates
BeanLoads	The number of times the bean data was loaded. Applies to entity
BeanStores	The number of times the bean data was stored in the database. Applies to entity

Measurement	Description
BeanInstantiates	The number of times a bean object was created. This applies to an individual bean, regardless of its type.
StatelessBeanInstantiates	The number of times a stateless session bean object was created
StatefulBeanInstantiates	The number of times a stateful session bean object was created
EntityBeanInstantiates	The number of times an entity bean object was created
BeanDestroys	The number of times an individual bean object was destroyed. This applies to any bean, regardless of its type
StatelessBeanDestroys	The number of times a stateless session bean object was destroyed
StatefulBeanDestroys	The number of times a stateful session bean object was destroyed
EntityBeanDestroys	The number of times an entity bean object was destroyed
BeansActive	The average number of instances of active beans pertaining to a specific bean. Applies to an individual bean that is either 'stateful' or 'entity'
EntityBeansActive	The average number of active entity beans
StatefulBeansActive	The average number of active session beans
BeansLive	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.
StatelessBeansLive	The average number of stateless session bean objects that are instantiated but not yet destroyed
StatefulBeansLive	The average number of stateful session bean objects that are instantiated but not yet destroyed

Measurement	Description
EntityBeansLive	The average number of entity bean objects that are instantiated but not yet destroyed
BeanMethodRT	The average method response time for all methods defined in the remote interface to this bean. Applies to all beans
BeanMethodActive	The average number of methods being processed concurrently. Applies to all beans
BeanMethodCalls	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
BeanObjectPoolContainer	The pool of a specific bean type
BeanObject	The pool specific to a home
NumGet	The number of calls retrieving an object from the pool
NumGetFound	The number of calls to the pool that resulted in finding an available bean
NumPuts	The number of beans that were released to the pool
NumPutsDiscarded	The number of times releasing a bean to the pool resulted in the bean being discarded because the pool was full
NumDrains	The number of times the daemon found the pool was idle and attempted to clean it

Measurement	Description
DrainSize	The average number of beans discarded by the daemon during a clean
BeanPoolSize	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
ActiveThreads	The average number of active threads in the pool
TotalThreads	The average number of threads in the pool
PercentTimeMaxed	The average percent of the time that the number of threads in the pool reached or exceeded the desired maximum number
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

Measurement	Description
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
NumTransactions	The number of transactions processed
ActiveTransactions	The average number of active transactions
TransactionRT	The average duration of each transaction
BeanObjectCount	The average number of bean object pools involved in a transaction
RolledBack	The number of transactions rolled back
Committed	The number of transactions committed
LocalTransactions	The number of transactions that were local
TransactionMethodCount	The average number of methods invoked as part of each transaction
Timeouts	The number of transactions that timed out due to inactivity timeouts
TransactionSuspended	The average number of times that a transaction was suspended

► ServletEngine

These are resources that are related to servlets and JSPs.

Measurement	Description
ServletsLoaded	The number of servlets currently loaded
ServletRequests	The number of requests serviced
CurrentRequests	The number of requests currently being serviced
ServletRT	The average response time for each request
ServletsActive	The average number of servlets actively processing requests
ServletIdle	The amount of time that the server has been idle (i.e., time since last request)
ServletErrors	The number of requests that resulted in an error or an exception
ServletBeanCalls	The number of bean method invocations that were made by the servlet
ServletBeanCreates	The number of bean references that were made by the servlet
ServletDBCalls	The number of database calls made by the servlet
ServletDBConAlloc	The number of database connections allocated by the servlet
SessionLoads	The number of times the servlet session data was read from the database
SessionStores	The number of times the servlet session data was stored in the database
SessionSize	The average size, in bytes, of a session data
LoadedSince	The time that has passed since the server was loaded (UNC time)

► Sessions

These are general metrics regarding the HTTP session pool.


Measurement	Description
SessionsCreated	The number of sessions created on the server
SessionsActive	The number of currently active sessions
SessionsInvalidated	The number of invalidated sessions. May not be valid when using sessions in the database mode
SessionLifetime	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 Controller machine. You may also need to copy the security keyring.

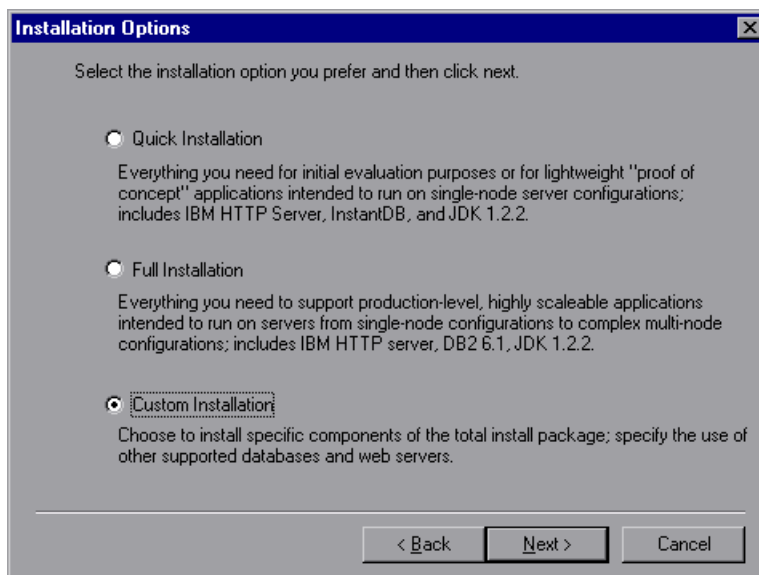
To install the IBM WebSphere Administrator's Console:

- 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. Follow the remaining instructions.

3 Click **Next** to proceed. The Installation Options dialog box opens.

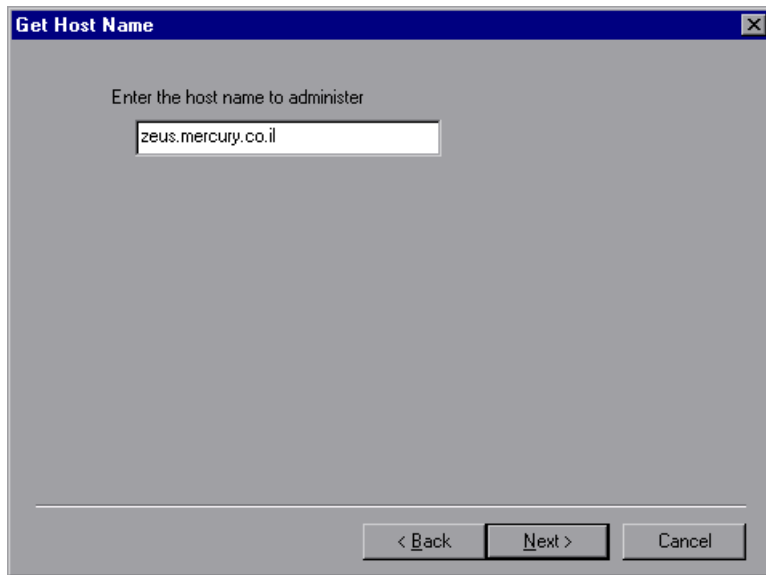


- 4 Select **Custom Installation**, and click **Next**. The Choose Application Server Components dialog box opens.



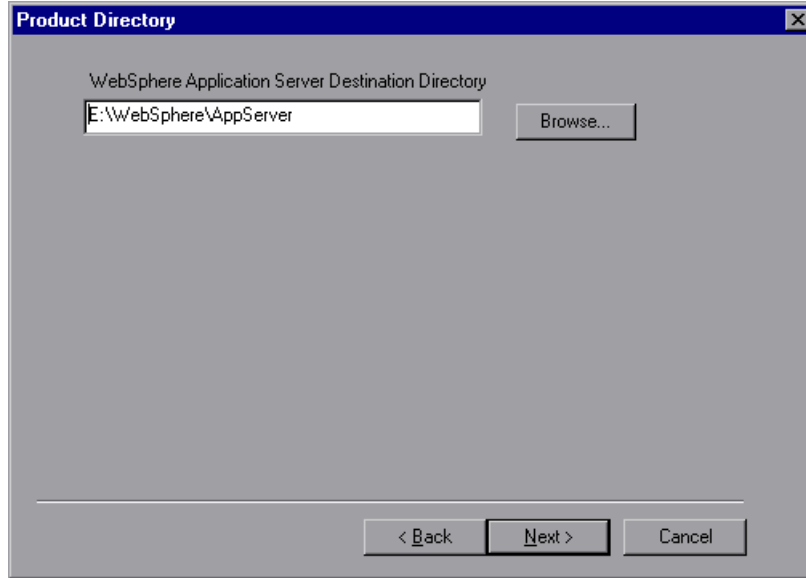
- 5 Select **Administrator's Console** and **IBM JDK 1.2.2**. Clear all the other options.

- 6 Click **Next**. The Get Host Name dialog box opens.



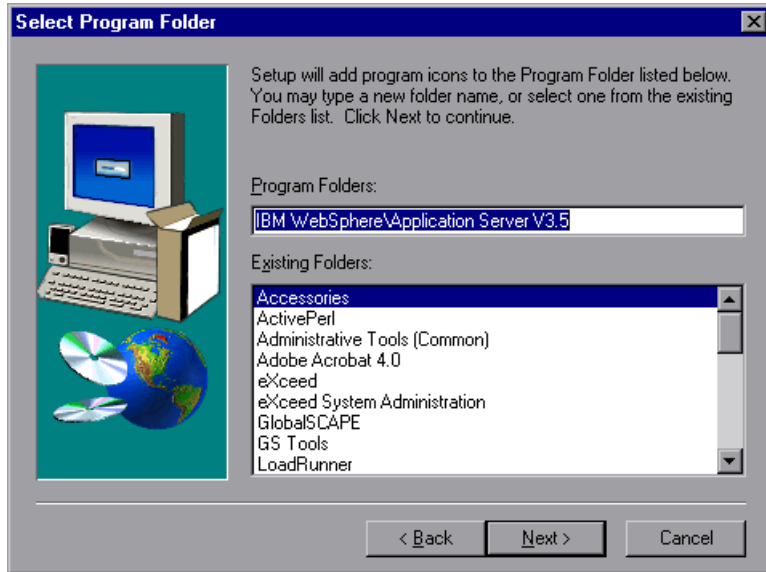
- 7 Type the name of the machine that you want to monitor.

- 8 Click **Next**. The Product Directory dialog box opens.




- 9 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**.

- 10 Click **Next**. The Select Program Folder dialog box opens.




- 11 Specify a program folder, or accept the default folder, IBM WebSphere\Application Server V3.5.
- 12 Click **Next**. The installation process begins. To pause or quit the installation, click **Cancel**.



When the installation is complete, the Setup Complete dialog box opens.



- 13** In the Setup Complete dialog box, 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**.
- 14** Click **Finish** to complete the installation program. The Restarting Windows dialog box opens.



- 15 Select either to restart your computer and complete the installation now (recommended) or to wait and complete the installation later.
- 16 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 <LR 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

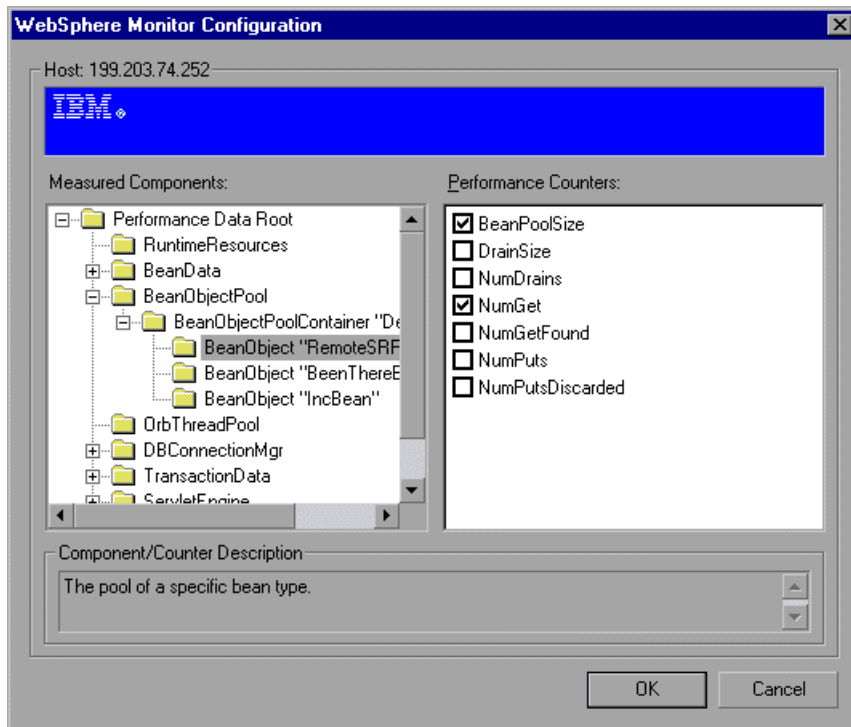
Activating the WebSphere (EPM) Monitor

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 the WebSphere (EPM) graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.

- 3 In the Monitored Server Machines section of the WebSphere (EPM) dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4 In the Resource Measurements section of the WebSphere (EPM) dialog box, click **Add** to select the measurements that you want to monitor. 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 413.
- 7 Click **OK** in the WebSphere Monitor Configuration dialog box, and in the WebSphere (EPM) dialog box, to activate the WebSphere (EPM) monitor.

26

Database Resource Monitoring

You can monitor DB2, Oracle, SQL Server, or Sybase database resource usage during a scenario run using LoadRunner'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 scenario 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 scenario. 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 scenario 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 Controller 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 Controller

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:

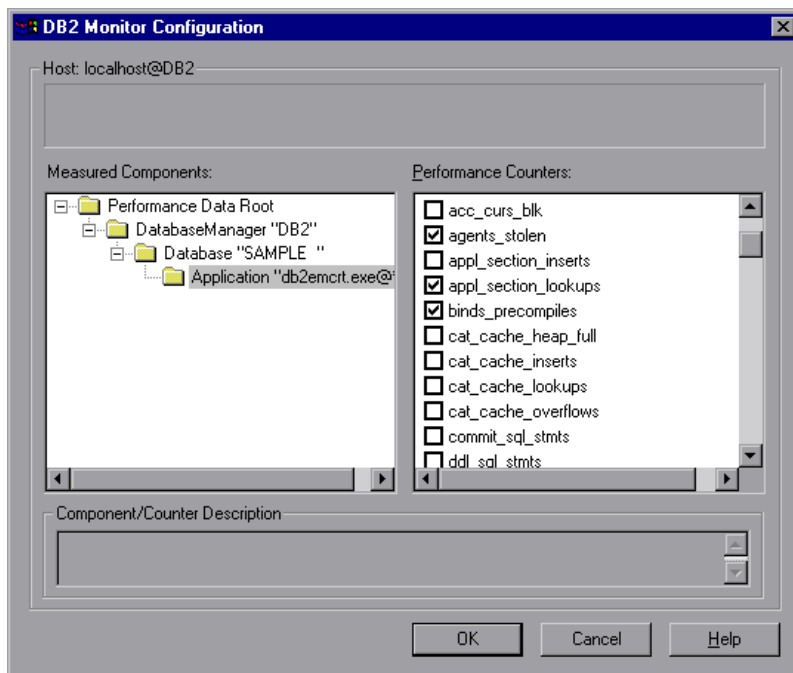
- 1 Click the DB2 graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**. The DB2 dialog box opens.
- 3 Click the **Add** button in the Monitored Server Machines section of the dialog box. The Add Machine dialog box opens.
- 4 In the Name box, enter the DB2 server machine name followed by the @ sign and the database instance you specified in the DB2 Control Center. In the Platform box, select **N/A**.



Click **OK** to save the information you entered and close the dialog box.

- 5 Click **Add** in the Resource Measurements section of the DB2 dialog box. 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
rem_cons_in	The current number of connections initiated from remote clients to the instance of the database manager that is being monitored.
rem_cons_in_exec	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.
local_cons	The number of local applications that are currently connected to a database within the database manager instance being monitored.
local_cons_in_exec	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.
con_local_databases	The number of local databases that have applications connected.
agents_registered	The number of agents registered in the database manager instance that is being monitored (coordinator agents and subagents).
agents_waiting_on_token	The number of agents waiting for a token so they can execute a transaction in the database manager.
idle_agents	The number of agents in the agent pool that are currently unassigned to an application and are therefore "idle".
agents_from_pool	The number of agents assigned from the agent pool
agents_created_empty_pool	The number of agents created because the agent pool was empty.

Measurement	Description
agents_stolen	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.
comm_private_mem	The amount of private memory that the instance of the database manager has currently committed at the time of the snapshot.
inactive_gw_agents	The number of DRDA agents in the DRDA connections pool that are primed with a connection to a DRDA database, but are inactive.
num_gw_conn_switches	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.
sort_heap_allocated	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.
post_threshold_sorts	The number of sorts that have requested heaps after the sort heap threshold has been reached.
piped_sorts_requested	The number of piped sorts that have been requested.
piped_sorts_accepted	The number of piped sorts that have been accepted.

Database

Measurement	Description
appls_cur_cons	Indicates the number of applications that are currently connected to the database.
appls_in_db2	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
total_sec_cons	The number of connections made by a sub-agent to the database at the node.
num_assoc_agents	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.
sort_heap_allocated	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.
total_sorts	The total number of sorts that have been executed.
total_sort_time	The total elapsed time (in milliseconds) for all sorts that have been executed.
sort_overflows	The total number of sorts that ran out of sort heap and may have required disk space for temporary storage.
active_sorts	The number of sorts in the database that currently have a sort heap allocated.
total_hash_joins	The total number of hash joins executed.
total_hash_loops	The total number of times that a single partition of a hash join was larger than the available sort heap space.
hash_join_overflows	The number of times that hash join data exceeded the available sort heap space
hash_join_small_overflows	The number of times that hash join data exceeded the available sort heap space by less than 10%.
pool_data_l_reads	Indicates the number of logical read requests for data pages that have gone through the buffer pool.
pool_data_p_reads	The number of read requests that required I/O to get data pages into the buffer pool.

Measurement	Description
pool_data_writes	Indicates the number of times a buffer pool data page was physically written to disk.
pool_index_l_reads	Indicates the number of logical read requests for index pages that have gone through the buffer pool.
pool_index_p_reads	Indicates the number of physical read requests to get index pages into the buffer pool.
pool_index_writes	Indicates the number of times a buffer pool index page was physically written to disk.
pool_read_time	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.
pool_write_time	Provides the total amount of time spent physically writing data or index pages from the buffer pool to disk.
files_closed	The total number of database files closed.
pool_async_data_reads	The number of pages read asynchronously into the buffer pool.
pool_async_data_writes	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.
pool_async_index_writes	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.
pool_async_index_reads	The number of index pages read asynchronously into the buffer pool by a pre-fetcher.
pool_async_read_time	The total elapsed time spent reading by database manager pre-fetchers.

Measurement	Description
pool_async_write_time	The total elapsed time spent writing data or index pages from the buffer pool to disk by database manager page cleaners.
pool_async_data_read_reqs	The number of asynchronous read requests.
pool_lsn_gap_clns	The number of times a page cleaner was invoked because the logging space used had reached a pre-defined criterion for the database.
pool_drty_pg_steal_clns	The number of times a page cleaner was invoked because a synchronous write was needed during the victim buffer replacement for the database.
pool_drty_pg_thrsh_clns	The number of times a page cleaner was invoked because a buffer pool had reached the dirty page threshold criterion for the database.
prefetch_wait_time	The time an application spent waiting for an I/O server (pre-fetcher) to finish loading pages into the buffer pool.
pool_data_to_estore	The number of buffer pool data pages copied to extended storage.
pool_index_to_estore	The number of buffer pool index pages copied to extended storage.
pool_data_from_estore	The number of buffer pool data pages copied from extended storage.
pool_index_from_estore	The number of buffer pool index pages copied from extended storage.
direct_reads	The number of read operations that do not use the buffer pool.
direct_writes	The number of write operations that do not use the buffer pool.
direct_read_reqs	The number of requests to perform a direct read of one or more sectors of data.
direct_write_reqs	The number of requests to perform a direct write of one or more sectors of data.

Measurement	Description
direct_read_time	The elapsed time (in milliseconds) required to perform the direct reads.
direct_write_time	The elapsed time (in milliseconds) required to perform the direct writes.
cat_cache_lookups	The number of times that the catalog cache was referenced to obtain table descriptor information.
cat_cache_inserts	The number of times that the system tried to insert table descriptor information into the catalog cache.
cat_cache_overflows	The number of times that an insert into the catalog cache failed due the catalog cache being full.
cat_cache_heap_full	The number of times that an insert into the catalog cache failed due to a heap-full condition in the database heap.
pkg_cache_lookups	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.
pkg_cache_inserts	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.
pkg_cache_num_overflows	The number of times that the package cache overflowed the bounds of its allocated memory.
appl_section_lookups	Lookups of SQL sections by an application from its SQL work area.
appl_section_inserts	Inserts of SQL sections by an application from its SQL work area.
sec_logs_allocated	The total number of secondary log files that are currently being used for the database.

Measurement	Description
log_reads	The number of log pages read from disk by the logger.
log_writes	The number of log pages written to disk by the logger.
total_log_used	The total amount of active log space currently used (in bytes) in the database.
locks_held	The number of locks currently held.
lock_list_in_use	The total amount of lock list memory (in bytes) that is in use.
deadlocks	The total number of deadlocks that have occurred.
lock_escalations	The number of times that locks have been escalated from several row locks to a table lock.
x_lock_escalations	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.
lock_timeouts	The number of times that a request to lock an object timed-out instead of being granted.
lock_waits	The total number of times that applications or connections waited for locks.
lock_wait_time	The total elapsed time waited for a lock.
locks_waiting	Indicates the number of agents waiting on a lock.
rows_deleted	The number of row deletions attempted.
rows_inserted	The number of row insertions attempted.
rows_updated	The number of row updates attempted.
rows_selected	The number of rows that have been selected and returned to the application.

Measurement	Description
int_rows_deleted	The number of rows deleted from the database as a result of internal activity.
int_rows_updated	The number of rows updated from the database as a result of internal activity.
int_rows_inserted	The number of rows inserted into the database as a result of internal activity caused by triggers.
static_sql_stmts	The number of static SQL statements that were attempted.
dynamic_sql_stmts	The number of dynamic SQL statements that were attempted.
failed_sql_stmts	The number of SQL statements that were attempted, but failed.
commit_sql_stmts	The total number of SQL COMMIT statements that have been attempted.
rollback_sql_stmts	The total number of SQL ROLLBACK statements that have been attempted.
select_sql_stmts	The number of SQL SELECT statements that were executed.
uid_sql_stmts	The number of SQL UPDATE, INSERT, and DELETE statements that were executed.
ddl_sql_stmts	This element indicates the number of SQL Data Definition Language (DDL) statements that were executed.
int_auto_rebinds	The number of automatic rebinds (or recompiles) that have been attempted.
int_commits	The total number of commits initiated internally by the database manager.
int_rollbacks	The total number of rollbacks initiated internally by the database manager.

Measurement	Description
int_deadlock_rollbacks	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.
binds_precompiles	The number of binds and pre-compiles attempted.

Application

Measurement	Description
agents_stolen	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.
num_assoc_agents	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.
total_sorts	The total number of sorts that have been executed.
total_sort_time	The total elapsed time (in milliseconds) for all sorts that have been executed.
sort_overflows	The total number of sorts that ran out of sort heap and may have required disk space for temporary storage.
total_hash_joins	The total number of hash joins executed.
total_hash_loops	The total number of times that a single partition of a hash join was larger than the available sort heap space.
hash_join_overflows	The number of times that hash join data exceeded the available sort heap space

Measurement	Description
hash_join_small_overflows	The number of times that hash join data exceeded the available sort heap space by less than 10%.
pool_data_l_reads	Indicates the number of logical read requests for data pages that have gone through the buffer pool.
pool_data_p_reads	The number of read requests that required I/O to get data pages into the buffer pool.
pool_data_writes	Indicates the number of times a buffer pool data page was physically written to disk.
pool_index_l_reads	Indicates the number of logical read requests for index pages that have gone through the buffer pool.
pool_index_p_reads	Indicates the number of physical read requests to get index pages into the buffer pool.
pool_index_writes	Indicates the number of times a buffer pool index page was physically written to disk.
pool_read_time	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.
prefetch_wait_time	The time an application spent waiting for an I/O server (pre-fetcher) to finish loading pages into the buffer pool.
pool_data_to_estore	The number of buffer pool data pages copied to extended storage.
pool_index_to_estore	The number of buffer pool index pages copied to extended storage.
pool_data_from_estore	The number of buffer pool data pages copied from extended storage.
pool_index_from_estore	The number of buffer pool index pages copied from extended storage.

Measurement	Description
direct_reads	The number of read operations that do not use the buffer pool.
direct_writes	The number of write operations that do not use the buffer pool.
direct_read_reqs	The number of requests to perform a direct read of one or more sectors of data.
direct_write_reqs	The number of requests to perform a direct write of one or more sectors of data.
direct_read_time	The elapsed time (in milliseconds) required to perform the direct reads.
direct_write_time	The elapsed time (in milliseconds) required to perform the direct writes.
cat_cache_lookups	The number of times that the catalog cache was referenced to obtain table descriptor information.
cat_cache_inserts	The number of times that the system tried to insert table descriptor information into the catalog cache.
cat_cache_overflows	The number of times that an insert into the catalog cache failed due the catalog cache being full.
cat_cache_heap_full	The number of times that an insert into the catalog cache failed due to a heap-full condition in the database heap.
pkg_cache_lookups	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.
pkg_cache_inserts	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.

Measurement	Description
appl_section_lookups	Lookups of SQL sections by an application from its SQL work area.
appl_section_inserts	Inserts of SQL sections by an application from its SQL work area.
uow_log_space_used	The amount of log space (in bytes) used in the current unit of work of the monitored application.
locks_held	The number of locks currently held.
deadlocks	The total number of deadlocks that have occurred.
lock_escals	The number of times that locks have been escalated from several row locks to a table lock.
x_lock_escals	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.
lock_timeouts	The number of times that a request to lock an object timed-out instead of being granted.
lock_waits	The total number of times that applications or connections waited for locks.
lock_wait_time	The total elapsed time waited for a lock.
locks_waiting	Indicates the number of agents waiting on a lock.
uow_lock_wait_time	The total amount of elapsed time this unit of work has spent waiting for locks.
rows_deleted	The number of row deletions attempted.
rows_inserted	The number of row insertions attempted.
rows_updated	The number of row updates attempted.
rows_selected	The number of rows that have been selected and returned to the application.

Measurement	Description
rows_written	The number of rows changed (inserted, deleted or updated) in the table.
rows_read	The number of rows read from the table.
int_rows_deleted	The number of rows deleted from the database as a result of internal activity.
int_rows_updated	The number of rows updated from the database as a result of internal activity.
int_rows_inserted	The number of rows inserted into the database as a result of internal activity caused by triggers.
open_rem_curs	The number of remote cursors currently open for this application, including those cursors counted by 'open_rem_curs_blk'.
open_rem_curs_blk	The number of remote blocking cursors currently open for this application.
rej_curs_blk	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.
acc_curs_blk	The number of times that a request for an I/O block was accepted.
open_loc_curs	The number of local cursors currently open for this application, including those cursors counted by 'open_loc_curs_blk'.
open_loc_curs_blk	The number of local blocking cursors currently open for this application.
static_sql_stmts	The number of static SQL statements that were attempted.
dynamic_sql_stmts	The number of dynamic SQL statements that were attempted.
failed_sql_stmts	The number of SQL statements that were attempted, but failed.

Measurement	Description
commit_sql_stmts	The total number of SQL COMMIT statements that have been attempted.
rollback_sql_stmts	The total number of SQL ROLLBACK statements that have been attempted.
select_sql_stmts	The number of SQL SELECT statements that were executed.
uid_sql_stmts	The number of SQL UPDATE, INSERT, and DELETE statements that were executed.
ddl_sql_stmts	This element indicates the number of SQL Data Definition Language (DDL) statements that were executed.
int_auto_rebinds	The number of automatic rebinds (or recompiles) that have been attempted.
int_commits	The total number of commits initiated internally by the database manager.
int_rollbacks	The total number of rollbacks initiated internally by the database manager.
int_deadlock_rollbacks	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.
binds_precompiles	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 DB2 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 Controller 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 Controller 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 Controller machine running the Oracle monitor. If you have a 16-bit and a 32-bit Oracle client installation on the controller 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 Controller 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 Controller 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 Controller 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 LoadRunner 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.

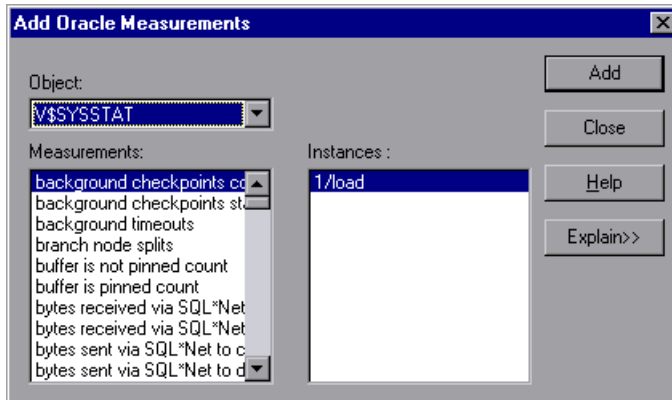
To configure the Oracle monitor:

- 1** Click the Oracle graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the Oracle dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select any platform, and click **OK**.
- 4** In the Resource Measurements section of the Oracle dialog box, click **Add** to select the measurements that you want to monitor.

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 454.

The following measurements are most commonly used when monitoring the Oracle server (from the V\$SYSSTAT table):

Measurement	Description
CPU used by this session	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.
Bytes received via SQL*Net from client	The total number of bytes received from the client over Net8
Logons current	The total number of current logons
Opens of replaced files	The total number of files that needed to be reopened because they were no longer in the process file cache
User calls	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.
SQL*Net roundtrips to/from client	The total number of Net8 messages sent to, and received from, the client
Bytes sent via SQL*Net to client	The total number of bytes sent to the client from the foreground process(es)
Opened cursors current	The total number of current open cursors

Measurement	Description
DB block changes	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.
Total file opens	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 Oracle 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 the SQL Server graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the SQL Server dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

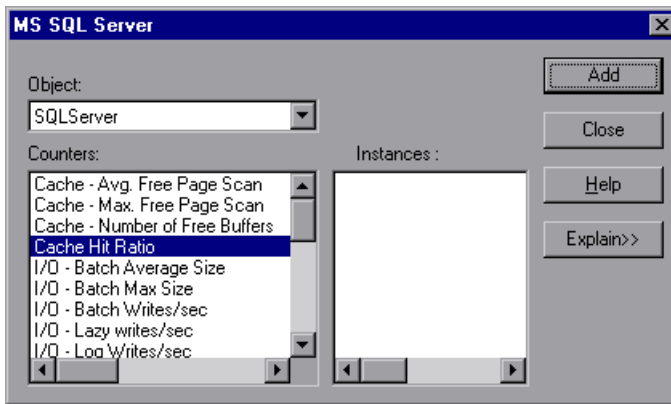
- 4 In the Resource Measurements section of the SQL Server dialog box, select the measurements you want to monitor. The following table describes the default counters that can be monitored on version 6.5 of the SQL Server:

Measurement	Description
% Total Processor Time (NT)	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.
% Processor Time (Win 2000)	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%.
Cache Hit Ratio	The percentage of time that a requested data page was found in the data cache (instead of being read from disk)
I/O - Batch Writes/sec	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.
I/O - Lazy Writes/sec	The number of 2K pages flushed to disk per second by the Lazy Writer
I/O - Outstanding Reads	The number of physical reads pending

Measurement	Description
I/O - Outstanding Writes	The number of physical writes pending
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 623.

- To select additional measurements, click **Add**. A dialog box displaying the SQL Server object, its counters, and instances opens.



- 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.

- 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 SQL Server 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 624.

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 Controller machine.

- 2 Verify a connection between the client and server on the Controller machine. To do so, use the Sybase client's *dsedit* tool to ping the Adaptive Server Enterprise Monitor Server.

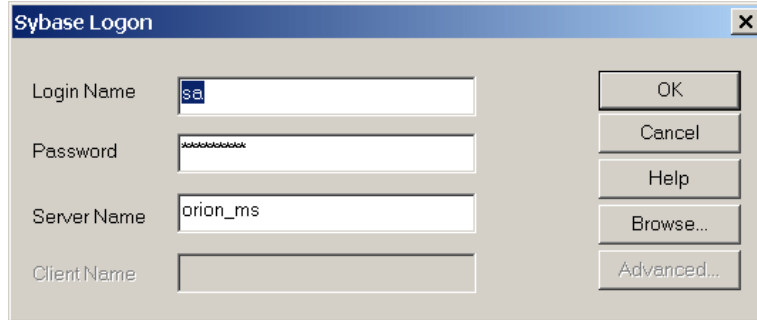


Note: The port you use to monitor a Sybase server through a firewall depends on the configuration of the Sybase server. Configuration information for the connection between the client and server is located in the Sybase client *sql.ini* file.

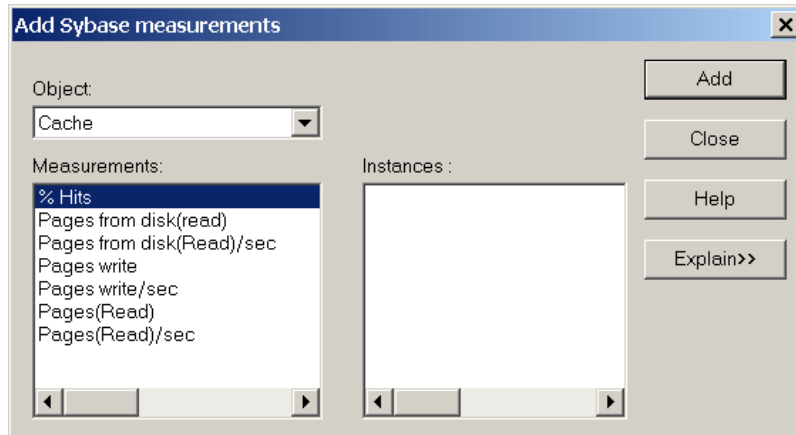
To configure the Sybase ASE monitor:

- 1 Click the Sybase graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the Sybase dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select any platform, and click **OK**.
- 4 In the Resource Measurements section of the Sybase dialog box, click **Add** to select the measurements that you want to monitor.

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 Sybase dialog box to activate the monitor.

27

Streaming Media Monitoring

During a scenario 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, see the *Creating Vuser Scripts* 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 scenario, and indicate which statistics and measurements you want to monitor. The RealPlayer Client and Media Player Client do not require pre-scenario 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.

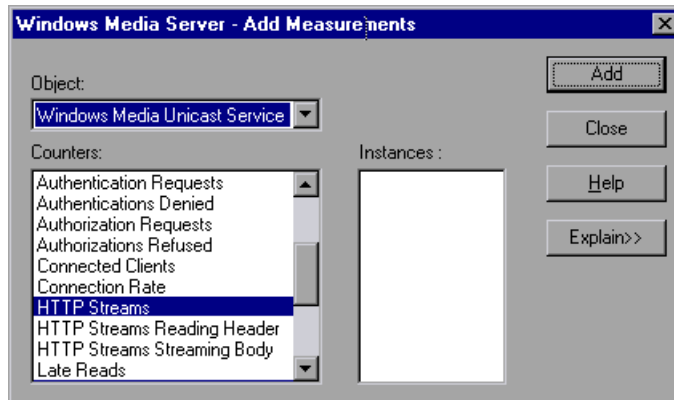
To configure the Windows Media Server monitor:

- 1** Click the Windows Media Server graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the Windows Media Server dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the Windows Media Server dialog box, select the measurements you want to monitor. The following table describes the default counters that can be monitored:

Measurement	Description
Active Live Unicast Streams (Windows)	The number of live unicast streams that are being streamed
Active Streams	The number of streams that are being streamed
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

Measurement	Description
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 To select additional measurements, click **Add**. The Windows Media Server - Add Measurements dialog box opens, displaying the Windows Media Unicast Service object, its counters, and instances.



- 6 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.

- 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 Windows Media Server dialog box to activate the monitor.

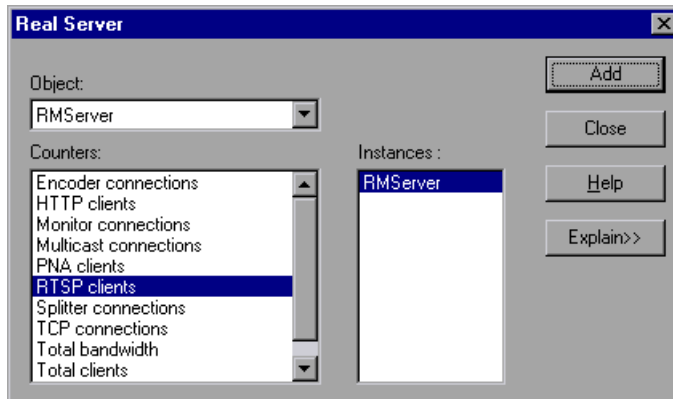
Configuring the RealPlayer Server Monitor

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 the Real Server graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3 In the Monitored Server Machines section of the Real Server dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4 Click **Add** in the Resource Measurements section of the Real Server dialog box to select the measurements that you want to monitor.

Another 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 Real Server 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 graph tree into the right pane of the Run view.

The following table describes the RealPlayer Client measurements that are monitored:

Measurement	Description
Current Bandwidth (Kbits/sec)	The number of kilobytes in the last second
Buffering Event Time (sec)	The average time spent on buffering
Network Performance	The ratio (percentage) between the current bandwidth and the actual bandwidth of the clip
Percentage of Recovered Packets	The percentage of error packets that were recovered
Percentage of Lost Packets	The percentage of packets that were lost
Percentage of Late Packets	The percentage of late packets
Time to First Frame Appearance (sec)	The time for first frame appearance (measured from the start of the replay)
Number of Buffering Events	The average number of all buffering events
Number of Buffering Seek Events	The average number of buffering events resulting from a seek operation
Buffering Seek Time	The average time spent on buffering events resulting from a seek operation
Number of Buffering Congestion Events	The average number of buffering events resulting from network congestion
Buffering Congestion Time	The average time spent on buffering events resulting from network congestion

Measurement	Description
Number of Buffering Live Pause Events	The average number of buffering events resulting from live pause
Buffering Live Pause Time	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 graph tree into the right pane of the Run view.

The following table describes the Media Player Client measurements that are monitored:

Measurement	Description
Stream Quality (Packet-level)	The percentage ratio of packets received to total packets
Current bandwidth (Kbits/sec)	The number of kbits per second received
Stream Packet Rate	The number of packets received
Total number of recovered packets	The number of lost packets that were recovered. This value is only relevant during network playback.
Total number of lost packets	The number of lost packets that were not recovered. This value is only relevant during network playback.
Stream Quality (Sampling-level)	The percentage of stream samples received on time (no delays in reception)

28

ERP Server Resource Monitoring

During a scenario run, you can monitor ERP server resources in order to isolate server performance bottlenecks.

This chapter describes:

- ▶ Setting Up the SAP Monitor
- ▶ Configuring the SAP Monitor

About ERP Server Resource Monitoring

The ERP server resource monitor provides you with performance information for the SAP R/3 system server. You can use the SAP monitor to view:

- ▶ the number of configured instances for each SAP system.
- ▶ data for all application instances (not just the one you logged on to)
- ▶ transactions used and the users that call them
- ▶ number of users working on the different instances
- ▶ performance history for recent periods of all instances
- ▶ response time distribution and resource consumption for any application server
- ▶ application server workload for today or for a recent period

In order to obtain this data, you need to activate the ERP server resource monitor before executing the scenario, and indicate which statistics and measurements you want to monitor.

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 Controller 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 Controller machine, using the username and password defined in the Controller.

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:

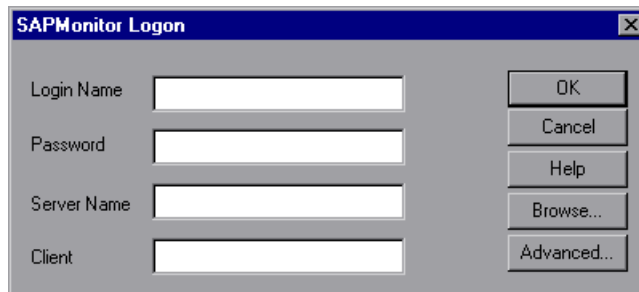
- 1** Click the SAP graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the SAP dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

Note: You can also specify a system number and IP address in the Add Machine dialog box using the following format:

<system number:IP address>

For example: 199.35.106.162:00

- 4** Click **Add** in the Resource Measurements section of the SAP dialog box. 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:

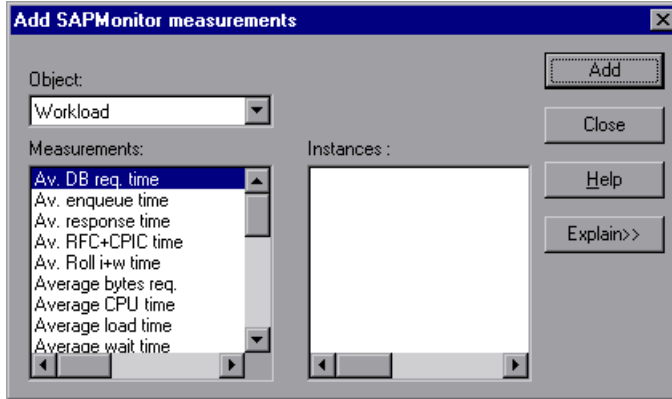
<RouterString/ServerIP/S/sapdpxx>

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 SAP dialog box to activate the monitor.

29

Java Performance Monitoring

During a scenario run, you can monitor the resource usage of Java 2 Platform, Enterprise Edition (J2EE) objects, 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

Note: The J2EE performance monitor is described separately in Chapter 30, “J2EE Performance Monitoring.”

About Java Performance Monitoring

The Java performance monitors provide you with performance information for Java 2 Platform, Enterprise Edition (J2EE) objects, Enterprise Java Bean (EJB) objects, and Java-based applications, using the J2EE, EJB, JProbe, and Sitraka JMonitor during scenario execution. In order to obtain this data, you need to activate the Java performance monitors before executing the scenario, 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 scenario 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, see the *Creating Vuser Scripts* 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 `<LoadRunner 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\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.
- 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 LoadRunner 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 LoadRunner 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 LoadRunner 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 Controller's EJB Monitor Configuration dialog box.

To configure the EJB monitor:

- 1 Click the EJB graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**. The EJB dialog box opens.



- 3 Click **Add** in the Monitored Server Machines box to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4 Click **Add** in the Resource Measurements section of the EJB dialog box. The EJB Monitor Configuration dialog box opens, displaying the available EJBs.



- 5 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.

- 6 Click **OK** in the EJB Monitor Configuration dialog box, and in the EJB dialog box, to activate the EJB monitor.

JProbe Performance Monitoring

You can monitor Java-based applications during a scenario 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 Controller to select the counters you want the JProbe Java profiling tool to measure.

Note: You must run the JPlauncher before opening the Controller.

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 Controller). In the section relating to the port to be used, select **Use Default Port**.

Note: The Controller 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 Controller'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="<Controller machine>:<Port>"
```


The JPlauncher establishes communication with the Controller machine. The Controller receives data from the JProbe tool through the above default port.

Configuring the JProbe Tool in the Controller

To monitor a Java-based application, you must select the counters you want the JProbe tool to measure. You select these counters using the Controller's JProbe dialog box.

To configure the JProbe tool:

- 1 Click the JProbe graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph in the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**. The JProbe dialog box opens.



- 3 Click **Add** in the Monitored Server Machines box to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

- 4 Click **Add** in the Resource Measurements section of the JProbe dialog box. 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 510.

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 `<LR_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\wlserver6.0\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\jrutils.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_ia32
```

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 482.

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 `<LR_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\wlserver6.0/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 ${CLASSPATH}  
-Dweblogic.Domain=mydomain -Dweblogic.Name=myserver "  
-Dbea.home=usr/bea"  
"-Djava.security.policy==usr/bea/wlserver6.0/lib/weblogic.policy"  
-Dweblogic.management.password=${WLS_PW} weblogic.Server
```


- 5 Run the `<WebLogic Home>\config\<domain name>\startWeblogicMercury.cmd` file to start the server with JMonitor and LoadRunner 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/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, 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
```

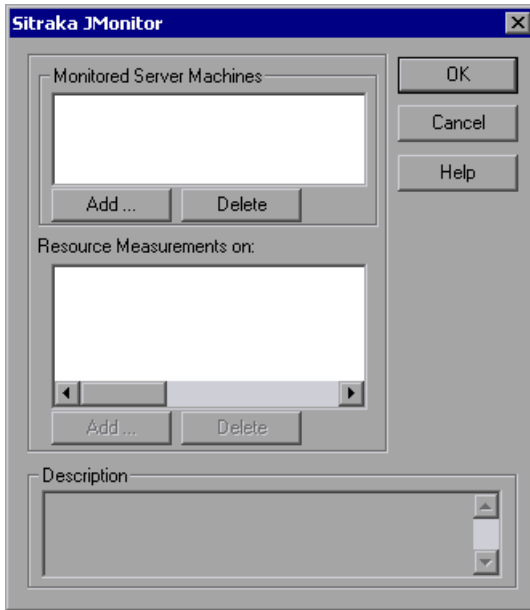
- 5 Click the **Advanced JVM Settings** button. In the Command line arguments field, add the following value for Windows 2000/NT, Solaris, and AIX:
-Xrunjdkhook:jmonitor:load_mws,proxy
- 6 Click the **OK** and **Apply** buttons to save the changes for the Application server. You can now start and stop your WebSphere server using the Sitraka JMonitor.

Configuring the Sitraka JMonitor in the Controller

To monitor a Java-based application, you must select the counters you want the Sitraka JMonitor to measure. You select these counters using the Controller's Sitraka JMonitor dialog box.

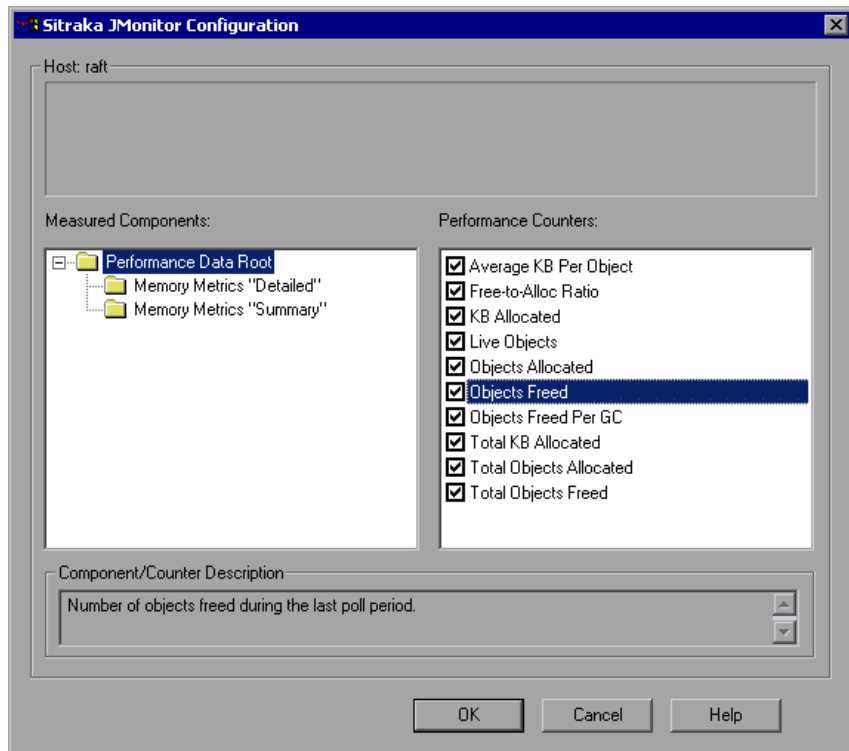
To configure the Sitraka JMonitor monitor:

- 1 Click the Sitraka JMonitor graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph in the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**. The Sitraka JMonitor dialog box opens.



- 3 Click **Add** in the Monitored Server Machines box to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

- 4 Click **Add** in the Resource Measurements section of the Sitraka JMonitor dialog box. The Sitraka JMonitor 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
% Free Heap Space	The percentage of free heap space since the last report.
% GC-To-Elapsed Time	The percentage of garbage collection to elapsed time.
% GC-To-Poll Time	The percentage of garbage collection to poll time.
% Used Heap	The percentage of used heap space since the last report.
Average GC Time (ms)	The average time, in milliseconds, spent performing garbage collections since the metric was enabled. (Disabling metric resets value to zero).
GC Time (ms)	The time, in milliseconds, spent performing garbage collections during the last poll period.
Heap Size (KB)	Total heap size, in kilobytes.
KB Freed	The number of kilobytes freed in the last poll period.
KB Freed Per GC	Average number of kilobytes freed per garbage collection since the metric was enabled (Disabling the metric resets value to zero).
Number of GCs	The number of garbage collections during the last poll period.
Total GC Time (ms)	The total time, in milliseconds, spent performing garbage collections since the metric was enabled. (Disabling the metric resets the value to zero).
Total GCs	The total number of garbage collections since the metric was enabled. (Disabling the metric resets value to zero).

Measurement	Description
Total KB Freed	The total number of kilobytes freed since the metric was enabled. (Disabling the metric resets value to zero).
Used Heap (KB)	Used heap size, in kilobytes.

DetailedMemoryMetrics

Measurement	Description
Average KB Per Object	The average number of kilobytes per object since the metric was enabled. (Disabling the metric resets value to zero).
Free-to-Alloc Ratio	The objects freed to objects allocated ratio since the metric was enabled (Disabling the metric resets value to zero).
KB Allocated	The number of kilobytes allocated since the metric was enabled. (Disabling the metric resets value to zero).
Live Objects	The change in number of live objects during the last poll period.
Objects Allocated	The number of objects allocated in the last poll period.
Objects Freed	The number of objects freed during the last poll period.
Objects Freed Per GC	The average number of objects freed per garbage collection since the metric was enabled. (Disabling the metric resets value to zero).
Total KB Allocated	The kilobytes allocated since metric was enabled. (Disabling the metric resets value to zero).

Measurement	Description
Total Objects Allocated	The number of objects allocated since the metric was enabled. (Disabling the metric resets value to zero).
Total Objects Freed	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 Sitraka JMonitor dialog box, to activate the Sitraka JMonitor monitor.

30

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 LoadRunner users to analyze J2EE component metrics during a scenario 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 LoadRunner Controller 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 523.

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 scenario 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 `<LoadRunner CD>\Add-ins\J2EE\jmonitor_<platform>.jar` file into that directory.

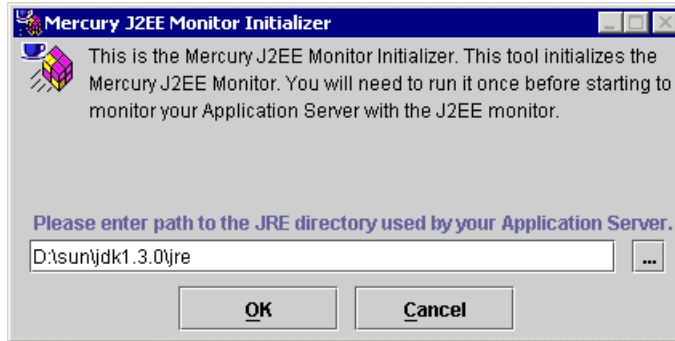
If you do not have WinZip to unzip the installation file, use the following command line to extract the installation file:

```
<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 530 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 525.

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 526.

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

To configure EJB information retrieval to include container methods:

- 1 Open `<J2EEMonitor Home Directory>\dat\java_monitor.ini`.
- 2 In the EJB_CONFIG section of the file, change the `hook_files=auto_detect` setting to the following:

```
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 Controller'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 Controller 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.

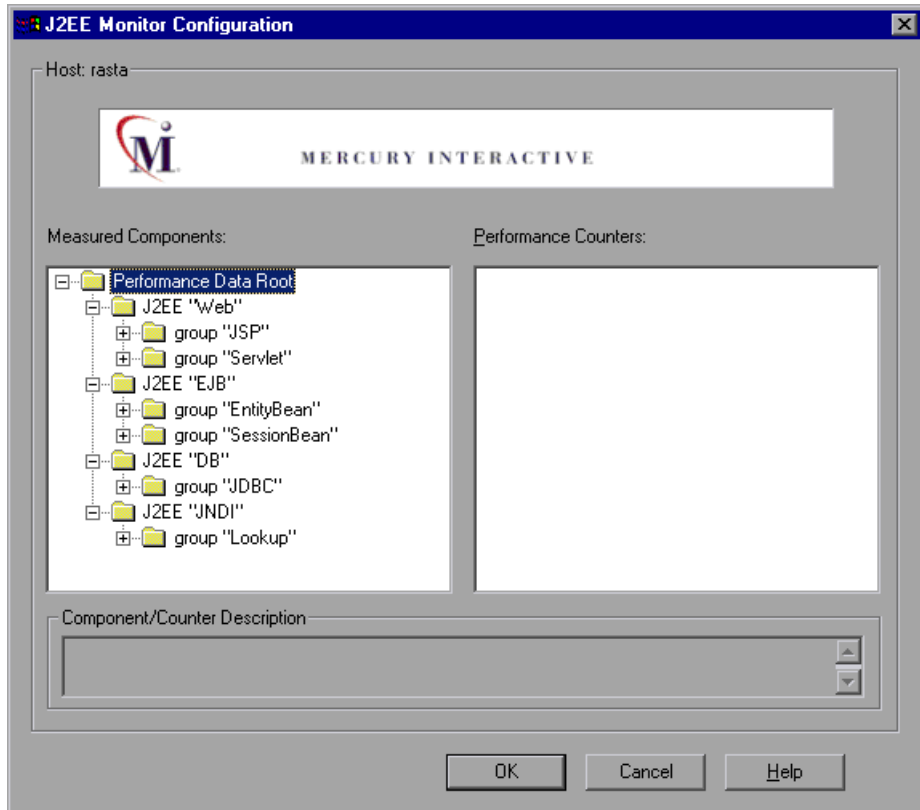
To configure the J2EE monitor:

- 1 Click the J2EE graph in the graph tree, and drag it into the right pane of the Run view.
- 2 Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**. The J2EE dialog box opens.



- 3 Click **Add** in the Monitored Server Machines box to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.

- 4 Click **Add** in the Resource Measurements section of the J2EE dialog box. The J2EE Monitor Configuration dialog box opens, displaying the available J2EE counters.



- 5 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 J2EE 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:

```
"%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 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 LoadRunner 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 LoadRunner 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 LoadRunner 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 LoadRunner, 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 LoadRunner machine, open *<LoadRunner 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.

Application Deployment Solutions

Using LoadRunner's Application Deployment Solution monitor, you can monitor the Citrix MetaFrame XP or 1.8 server during a scenario run and isolate server performance bottlenecks.

This chapter describes:

- ▶ Configuring the Citrix MetaFrame Server Monitor

About Application Deployment Solutions Monitoring

LoadRunner's Citrix MetaFrame XP monitor provides you with information about the application deployment usage of the Citrix MetaFrame XP and 1.8 servers during a scenario 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 scenario.

Configuring the Citrix MetaFrame Server Monitor

To monitor the Citrix server performance, you must first activate the Citrix MetaFrame XP monitor on the application server machine and enable the counters you want to monitor on the Citrix server. Then select the counters you want the Citrix MetaFrame XP monitor to measure. You select these counters using the Controller's 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 configuring the monitor:

- 1** From the Controller machine, map a network drive to the Citrix server machine. This ensures that the required authentication is provided to the Controller to access the resource counters.
- 2** Launch PerfMon from the Controller 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.
- 3** To provide the Controller with a list of measurements that will be available for monitoring, you must first initialize Vusers before running the scenario. After you have initialized the Vusers, you can then configure the Citrix Monitor and add the ICA Session counters.

Note: Measurements that monitor instances are valid for the currently running Citrix session only. If you run this scenario 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 XP graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the Citrix MetaFrame XP dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** In the Resource Measurements section of the dialog box, click **Add** to select the measurements that you want to monitor.

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>).

The following table describes some of the counters that can be measured.

Non-Virtual Counters

Measurement	Description
% Disk Time	The percentage of elapsed time that the selected disk drive is busy servicing read or write requests.
% Processor Time	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%.
File data Operations/sec	The rate that the computer is issuing Read and Write operations to file system devices. It does not include File Control Operations.

Measurement	Description
Interrupts/sec	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.
Output Session Line Speed	This value represents the line speed from server to client for a session in bps.
Input Session Line Speed	This value represents the line speed from client to server for a session in bps.
Page Faults/sec	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
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.
Pool Nonpaged Bytes	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.
Private Bytes	The current number of bytes this process has allocated that cannot be shared with other processes.
Processor Queue Length	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.
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.

Measurement	Description
Latency – Session Average	This value represents the average client latency over the life of a session.
Latency – Last Recorded	This value represents the last recorded latency measurement for this session.
Latency – Session Deviation	This value represents the difference between the minimum and maximum measured values for a session.
Input Session Bandwidth	This value represents the bandwidth from client to server traffic for a session in bps.
Input Session Compression	This value represents the compression ratio for client to server traffic for a session.
Output Session Bandwidth	This value represents the bandwidth from server to client traffic for a session in bps.
Output Session Compression	This value represents the compression ratio for server to client traffic for a session.
Output Session Linespeed	This value represents the line speed from server to client for a session in bps.

Virtual Channel Counters

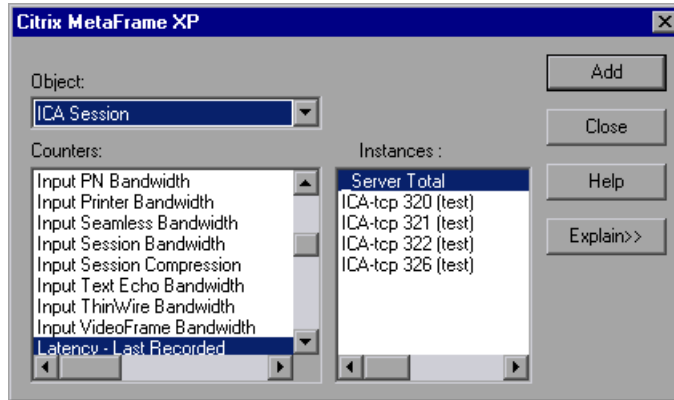
Measurement	Description
Input Audio Bandwidth	This value represents the bandwidth from client to server traffic on the audio mapping channel. This is measured in bps.
Input Clipboard Bandwidth	This value represents the bandwidth from client to server traffic on the clipboard mapping channel. This is measured in bps.
Input COM1 Bandwidth	This value represents the bandwidth from client to server traffic on the COM1 channel. This is measured in bps.

Measurement	Description
Input COM2 Bandwidth	This value represents the bandwidth from client to server traffic on the COM2 channel. This is measured in bps.
Input COM Bandwidth	This value represents the bandwidth from client to server traffic on the COM channel. This is measured in bps.
Input Control Channel Bandwidth	This value represents the bandwidth from client to server traffic on the ICA control channel. This is measured in bps.
Input Drive Bandwidth	This value represents the bandwidth from client to server traffic on the client drive mapping channel. This is measured in bps.
Input Font Data Bandwidth	This value represents the bandwidth from client to server traffic on the local text echo font and keyboard layout channel. This is measured in bps.
Input Licensing Bandwidth	This value represents the bandwidth from server to client traffic on the licensing channel. This is measured in bps.
Input LPT1 Bandwidth	This value represents the bandwidth from client to server traffic on the LPT1 channel. This is measured in bps.
Input LPT2 Bandwidth	This value represents the bandwidth from client to server traffic on the LPT2 channel. This is measured in bps.
Input Management Bandwidth	This value represents the bandwidth from client to server traffic on the client management channel. This is measured in bps.
Input PN Bandwidth	This value represents the bandwidth from client to server traffic on the Program Neighborhood channel. This is measured in bps.
Input Printer Bandwidth	This value represents the bandwidth from client to server traffic on the printer spooler channel. This is measured in bps.

Measurement	Description
Input Seamless Bandwidth	This value represents the bandwidth from client to server traffic on the Seamless channel. This is measured in bps.
Input Text Echo Bandwidth	This value represents the bandwidth from client to server traffic on the local text echo data channel. This is measured in bps.
Input Thinwire Bandwidth	This value represents the bandwidth from client to server traffic on the Thinwire (graphics) channel. This is measured in bps.
Input VideoFrame Bandwidth	This value represents the bandwidth from client to server traffic on the VideoFrame channel. This is measured in bps.
Output Audio Bandwidth	This value represents the bandwidth from server to client traffic on the audio mapping channel. This is measured in bps.
Output Clipboard Bandwidth	This value represents the bandwidth from server to client traffic on the clipboard mapping channel. This is measured in bps.
Output COM1 Bandwidth	This value represents the bandwidth from server to client traffic on the COM1 channel. This is measured in bps.
Output COM2 Bandwidth	This value represents the bandwidth from server to client traffic on the COM2 channel. This is measured in bps.
Output COM Bandwidth	This value represents the bandwidth from server to client traffic on the COM channel. This is measured in bps.
Output Control Channel Bandwidth	This value represents the bandwidth from server to client traffic on the ICA control channel. This is measured in bps.
Output Drive Bandwidth	This value represents the bandwidth from server to client traffic on the client drive channel. This is measured in bps.

Measurement	Description
Output Font Data Bandwidth	This value represents the bandwidth from server to client traffic on the local text echo font and keyboard layout channel. This is measured in bps.
Output Licensing Bandwidth	This value represents the bandwidth from server to client traffic on the licensing channel. This is measured in bps.
Output LPT1 Bandwidth	This value represents the bandwidth from server to client traffic on the LPT1 channel. This is measured in bps.
Output LPT2 Bandwidth	This value represents the bandwidth from server to client traffic on the LPT2 channel. This is measured in bps.
Output Management Bandwidth	This value represents the bandwidth from server to client traffic on the client management channel. This is measured in bps.
Output PN Bandwidth	This value represents the bandwidth from server to client traffic on the Program Neighborhood channel. This is measured in bps.
Output Printer Bandwidth	This value represents the bandwidth from server to client traffic on the printer spooler channel. This is measured in bps.
Output Seamless Bandwidth	This value represents the bandwidth from server to client traffic on the Seamless channel. This is measured in bps.
Output Text Echo Bandwidth	This value represents the bandwidth from server to client traffic on the local text echo data channel. This is measured in bps.
Output Thinwire Bandwidth	This value represents the bandwidth from server to client traffic on the Thinwire (graphics) channel. This is measured in bps.
Output VideoFrame Bandwidth	This value represents the bandwidth from server to client traffic on the VideoFrame channel. This is measured in bps.

- 5 To select additional measurements, click **Add**. A dialog box displaying the Citrix object, its counters, and instances opens. Choose the object whose counters you want to display. LoadRunner displays the object's counters in the Counters pane.



- 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 Citrix MetaFrame dialog box to activate the monitor.

Middleware Performance Monitoring

Using LoadRunner's Middleware Performance monitors, you can monitor the TUXEDO and the IBM WebSphere MQ servers during a scenario 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. LoadRunner's Middleware Performance monitors provide you with information about the middleware performance usage of the TUXEDO and IBM WebSphere MQ servers during a scenario 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 scenario.

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 Controller 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: 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 Controller 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 LoadRunner 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"> • 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. • QDPHIEV(action) – where action is the word “ENABLED” or “DISABLED”, enabling or disabling the generation of the event, respectively.
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"> • 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. • QDPLOEV(action) – where action is the word “ENABLED” or “DISABLED”, enabling or disabling the generation of the event, respectively.
Event - Queue Full	<ul style="list-style-type: none"> • 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. • QDPMAXEV(action) – where action is the word “ENABLED” or “DISABLED”, enabling or disabling the generation of the event, respectively.

Measurement	Set Event Attributes
Event - Queue Service Interval High	<ul style="list-style-type: none"> • QSVCIINT(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. • QSVCIIEV(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.
Event - Queue Service Interval OK	<ul style="list-style-type: none"> • QSVCIINT(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. • QSVCIIEV(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.

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 Controller

After you have installed the MQ Client on the Controller, 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 the IBM WebSphere MQ graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the IBM WebSphere MQ dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. The format of the server name is *<machine name> : <port number>*. Select the platform on which the machine runs, and click **OK**.

- 4 Click **Add** in the Resource Measurements section of the IBM WebSphere MQ dialog box to select the measurements that you want to monitor. The IBM WebSphere MQ Add Measurements dialog box opens.

MQMonitor - Add Measurements

Connection Information

Server: weblog:1415

Client Channel: [Dropdown]

Queue Manager: [Dropdown]

Available Measurements

Object type: < Select object type >

Object name: [Text Box]

Event/Attribute: [Text Box]

Filter system objects.

Add Object Add Remove Alternate Queue

Object	Attribute/Event	Alternate Queue Manager

OK Cancel Help

- 5 In the Connections 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 Controller 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 `<LR_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
Event - Channel Activated (events per second)	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.
Event - Channel Not Activated (events per second)	Event generated when a channel, attempts to become active but inhibited from doing so due to a shortage of queue manager channel slots.
Event - Channel Started (events per second)	Event generated when a channel is started.
Event - Channel Stopped (events per second)	Event generated when a channel is stopped, regardless of source of stoppage.
Event - Channel Stopped by User (events per second)	Event generated when a channel is stopped by a user.
Status - Channel State	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).
Status - Messages Transferred	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.
Status - Buffer Received	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
Status - Buffer Sent	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.
Status - Bytes Received	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.
Status - Bytes Sent	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.

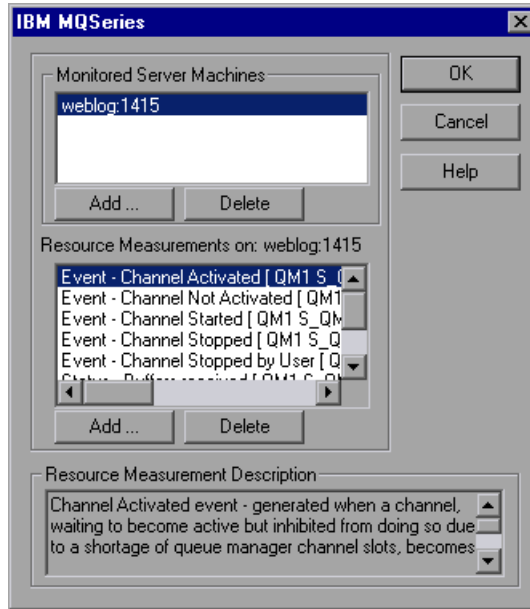
Note: In order to enable the event for a queue, ensure that the attributes for the queue have been set. For more information, refer to “Configuring the Server Environment to Monitor Events,” on page 553.

- 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**. The IBM MQSeries dialog box opens, displaying the name of the monitored server machine, a list of the resource measurements selected, and a description of each measurement.



- 15** Click **OK** in the IBM WebSphere MQ 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 Controller machine, more than one TUXEDO application server can be monitored at a time. However, if TUXEDO 6.5 or below is installed on the Controller 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 Controller 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 Controller 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 the TUXEDO graph in the graph tree, and drag it into the right pane of the Run view.
- 2** Right-click the graph and choose **Add Measurement(s)**, or choose **Monitors > Add Online Measurement**.
- 3** In the Monitored Server Machines section of the TUXEDO dialog box, click **Add** to enter the server name or IP address of the machine you want to monitor. Select the platform on which the machine runs, and click **OK**.
- 4** Click **Add** in the Resource Measurements section of the TUXEDO dialog box to select the measurements that you want to monitor. You will be prompted 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 LoadRunner script. You can also determine the client name from the **lrt_tpinitialize** statement in the recorded script.

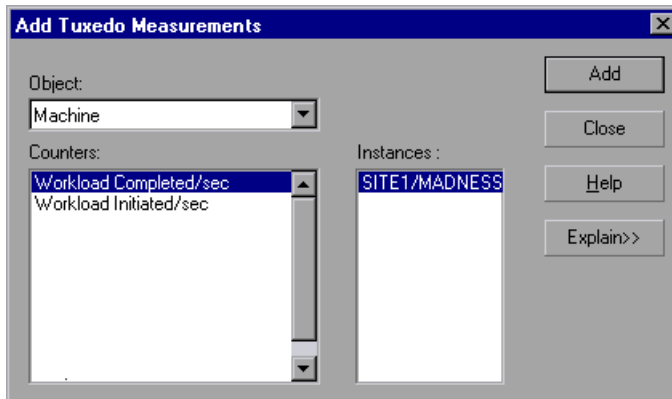
In the following example of a *tpinit.ini* file, the TUXEDO monitor was configured for a server named URANUS using port 65535, and a client named bankapp. The logon user name was Smith and the password was mypasswd.

```
[Logon]
LogonServername=//URANUS:65535
LogonUserName=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 Controller session. Once it connects to an application server, that server is the only one used by the monitor until the Controller 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	Requests per second - How many server requests were handled per second
	Workload per second -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	Workload completed per second - The total workload on all the servers for the machine that was completed, per unit time
	Workload initiated per second - The total workload on all the servers for the machine that was initiated, per unit time
	Current Accessers - Number of clients and servers currently accessing the application either directly on this machine or through a workstation handler on this machine.
	Current Clients - Number of clients, both native and workstation, currently logged in to this machine.
	Current Transactions - Number of in use transaction table entries on this machine.
Queue	Bytes on queue - The total number of bytes for all the messages waiting in the queue
	Messages on queue - The total number of requests that are waiting on queue. By default this is 0.

Monitor	Measurements
Workstation Handler (WSH)	Bytes received per second - The total number of bytes received by the workstation handler, per unit time
	Bytes sent per second - The total number of bytes sent back to the clients by the workstation handler, per unit time
	Messages received per second - The number of messages received by the workstation handler, per unit time
	Messages sent per second - The number of messages sent back to the clients by the workstation handler, per unit time
	Number of queue blocks per second - 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 TUXEDO dialog box to activate the monitor.

33

Troubleshooting Online Monitors

LoadRunner monitors allow you to view the performance of the scenario 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 LoadRunner cannot locate the specified server, make sure that the specified server is available. Perform a “ping” operation by typing `ping <server_name>` from the Controller 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: <code>%net use \\<MachineName>/</code> <code>user:[<Domain>\<RemoteMachineUsername>]</code> 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"> a. Using Explorer or File Manager, give the user READ access to: <ul style="list-style-type: none"> %windir%\system32\PERFCxxx.DAT %windir%\system32\PERFHxxx.DAT 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. b. Using REGEDT32, give the user READ access to: <ul style="list-style-type: none"> HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Perflib and all sub keys of that key. c. Using REGEDT32, give the user at least READ access to: <ul style="list-style-type: none"> HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurePipeServers\winreg
<p>Some Win 2000 counters cannot be monitored from an NT machine.</p>	<p>Run the Controller 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): HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\MSSQLServer\MSSQLServer (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 <i>lrun/bin</i> .
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 21, "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>LoadRunner</i> root folder\ <i>dat\monitors\WebLogicMon.ini</i> file, and search for: [WebLogicMonitor] JVM=javaw.exe Change javaw.exe to java.exe. A window containing trace information opens.

Troubleshooting the Network Delay Monitor

If monitoring is unsuccessful and LoadRunner 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 `<LR root folder\bin\webtrace.exe` from the source machine to determine whether the problem is related to the Controller, 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 Controller 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 Controller. Verify that the source machine is not a UNIX machine, and contact Mercury Interactive's Customer Support with the following information:

- the Controller log file, `drv_log.txt`, located in the temp directory of the Controller machine.
- the `traceroute_server` log file, located on the source machine. Note that in LoadRunner 7.02, this information is located in `<LR root folder\dat\<the latest mdrv log>`. In LoadRunner 7.5, this information is located in `<LR root folder\bin\traceroute_server.log`.

- 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 *<LR 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 scenarios, the network card has little impact on scenario 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 LoadRunner on a WAN or heavy loaded LAN, you may notice some unusual LoadRunner 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 Controller may not be reaching the LoadRunner 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 V

Appendixes

A

Interpreting LoadRunner Online Graphs

LoadRunner online monitor graphs present important information about the performance of your scenario. This appendix describes some of the key online graphs in greater depth and shows how they can be used to identify and pinpoint performance bottlenecks as your scenario is running.


Online Monitoring Graphs

Using the online monitor graphs, you can determine whether transactions transpire within an acceptable amount of time, whether your bandwidth is sufficient to keep download times to a minimum, and whether your hardware and operating system can handle peak load.

Question 1: Do all transactions in my scenario transpire within an acceptable amount of time? Which particular transactions take too long?


Answer: The **Transaction Response Time** graph shows the amount of time it takes for each transaction to be completed. Note that in the graph below, the transaction response time is quick, except for the login transaction. The initial login did not take much time, but subsequent logins were quite slow.

This indicates that the database is unable to process more than one login at a time, which may be due to inefficient database querying.



Question 2: Is bandwidth sufficient to keep download times to a minimum?


Answer: The **Throughput** graph shows the amount of throughput on the Web server during each second of the scenario run. Throughput represents the amount of data received from the server at any given second.



Note that in the above graph, the throughput scales upward as time progresses and the number of users increases, indicating that the bandwidth is sufficient. If the graph were to remain relatively flat as the number of users increased, it would be reasonable to conclude that the bandwidth is constraining the volume of data requested.

Question 3: Can the hardware and operating system handle peak load?

Answer: The **Windows Resources** graph displays Windows resource usage in real-time. You use this graph to monitor the resources used during a scenario and locate a bottleneck on a particular machine.



The *% Total Processor Time* in the above graph shows the amount of data processed by the server. *File Data Operations/sec* shows the rate at which the server is issuing Read and Write operations to file system devices. *Page Faults/sec* counts the number of page faults in the processor, representing virtual memory and caching algorithm opportunities.

It is commonly thought that newer and faster servers can resolve slow download times. However the above graph demonstrates that only a small amount of data is processed by the server. The graph indicates that there is adequate processor capacity, and additional server hardware will not result in increased performance. There are cases, however, in which increased performance can be achieved by optimizing the data file system.

B

Performing Path Translation

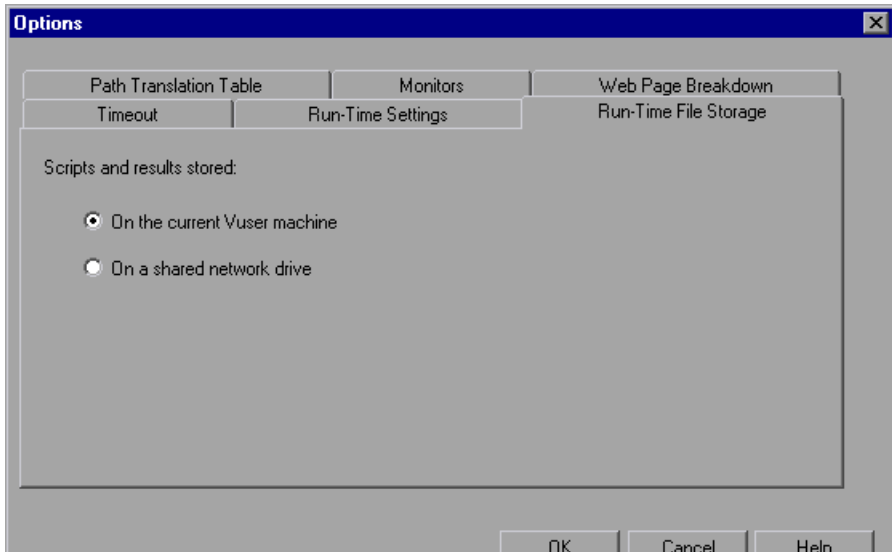
When you run a scenario, LoadRunner gathers run-time data from the participating Vusers. By default, LoadRunner stores the data in temporary files on each Vuser machine. After the scenario, the data is collated in the general results directory.

Alternatively, you can instruct LoadRunner to write the run-time data directly to a shared network drive. (See Chapter 10, “Configuring a Scenario.”) This method is not recommended, since it increases network traffic and necessitates path translation.

Understanding Path Translation

Path Translation is a mechanism used by LoadRunner to convert a remote path name for the Controller. A typical scenario might have the LoadRunner Controller 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 scenario such as this, you need to ensure that all participating machines recognize the same network drive.

You instruct LoadRunner 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 scenario—and their locations. A script's location (path) is always based on the Controller 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 Controller is running on a Windows-based machine named *pc2*, and that a Vuser script is located on a network drive. The Controller 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 Controller 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 Controller) 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:

```
<controller_host><controller_path><remote_path>[<remote_host>]
```

controller_host

The name or type of the machine that is running the Controller. For example, if the Controller 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 Controller (for example, LOADPC1).

The value of *controller_host* can be:

hostname	the name of the machine running the Controller
win	the Controller is running on a Windows-based computer
unix	the Controller is running on a UNIX machine
all	the Controller is running on a Windows-based or a UNIX machine

<i>controller_path</i>	The path of a specific directory—as recognized by the Controller. For example, if the directory <i>scripts</i> is located on the network drive <i>r</i> —as mapped by the Controller—type the path <i>r:\scripts</i> in the <i>controller_path</i> field.
<i>remote_path</i>	<p>The path of a specific directory—as recognized by the remote machine. For example, if the directory <i>scripts</i> is located on the network drive <i>n</i>—as mapped by the remote load generator—type the path <i>n:\scripts</i> in the <i>remote_path</i> field.</p> <p>If a Vuser on the remote UNIX load generator recognizes the above path as <i>/m/tests</i>, you would type this path in the <i>remote_path</i> field.</p>
<i>remote_host</i>	The name or type of the remote load generator. For example, if all the remote machines are UNIX workstations, you could type <i>unix</i> in the <i>remote_host</i> field. The options for the <i>remote_host</i> field are the same as the options for the <i>controller_host</i> field, listed above. The <i>remote_host</i> parameter is optional.

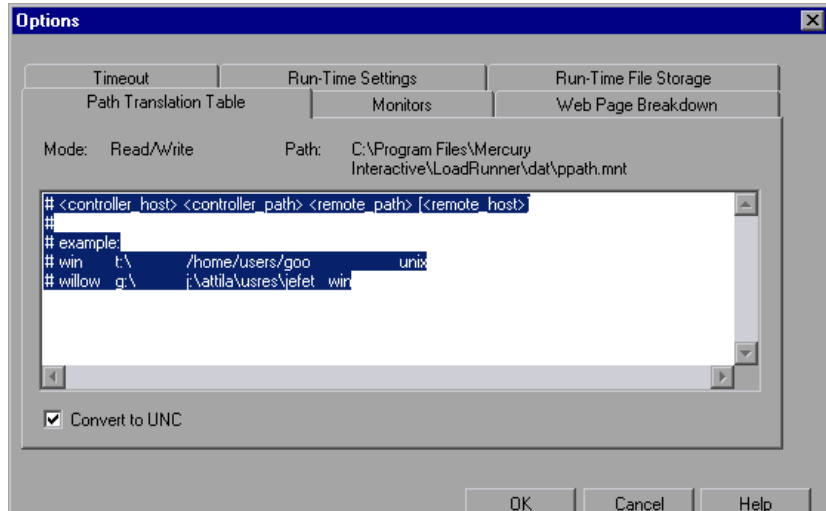
Editing the Path Translation Table

You maintain the Path Translation table using the LoadRunner Controller. LoadRunner saves the Path Translation table as an ASCII file, *ppath.mnt*. This file, stored in *LoadRunner_directory/dat*, has a one-line entry for each network path to translate.

To edit the Path Translation table:

- 1 Start the LoadRunner Controller.

- 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 Controller 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 LoadRunner 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 Controller 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 Controller'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
--------	---------	--------------	-----

C

Working in Expert Mode

Advanced users can fine-tune the LoadRunner 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 LoadRunner Controller Expert mode is intended for support personnel to provide access to system information. When you work in the Expert mode, the Controller dialog boxes contain additional options for fine tuning the Controller operation.

To activate the Expert mode, choose **Tools > Expert Mode**. An active Expert mode is indicated by a check mark.

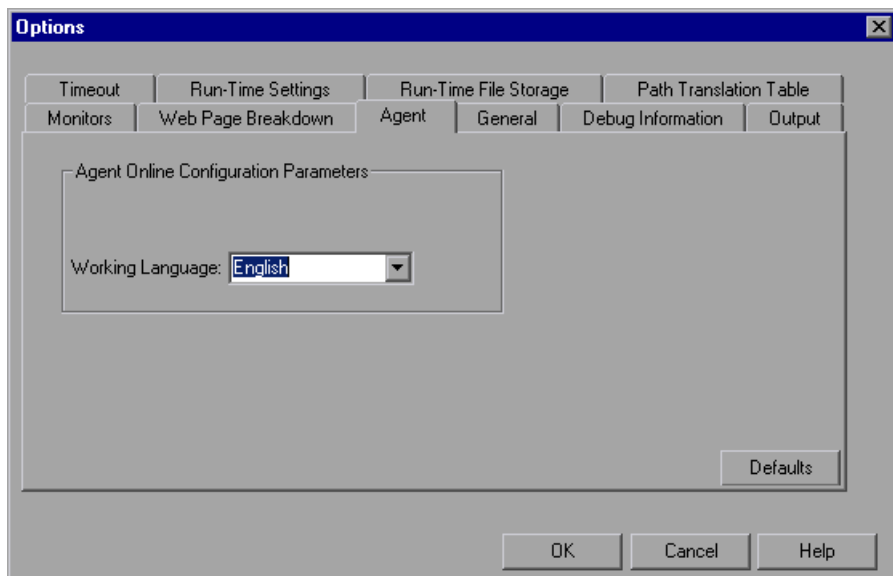
To exit the Expert mode, repeat the above process.

Options - Agent Settings

The Agent settings allow you to customize the behavior of the LoadRunner 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 agent's working language (English or Japanese).
- 4 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 LoadRunner not to collate log files.

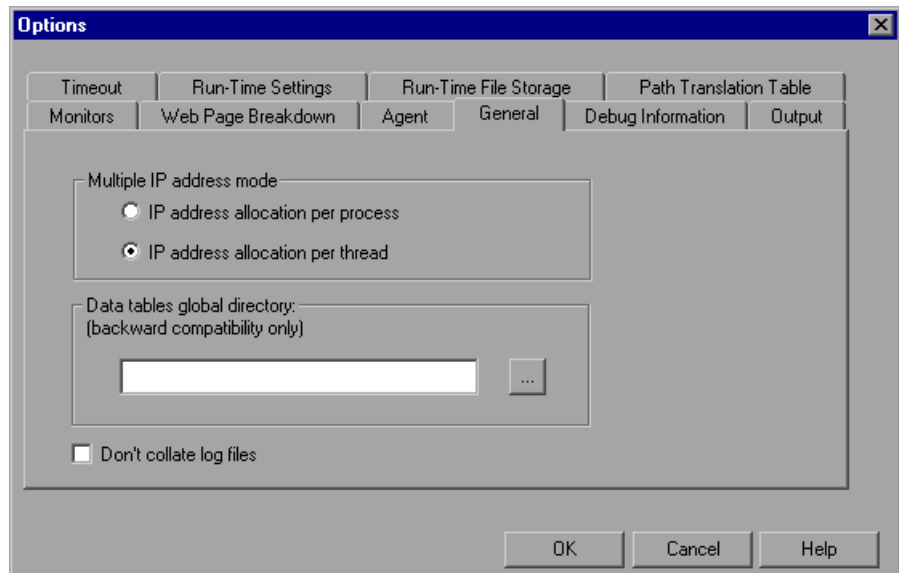
Multiple IP address mode: The mode used to allocate IP addresses when the multiple IP address option is enabled (**Scenario > Enable IP Spoofer**). The Controller can allocate an IP address per process or per thread. Allocation per thread results in a more varied range of IP addresses in a scenario.

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 LoadRunner.

Do not collate log files: Instructs LoadRunner 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 LoadRunner 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

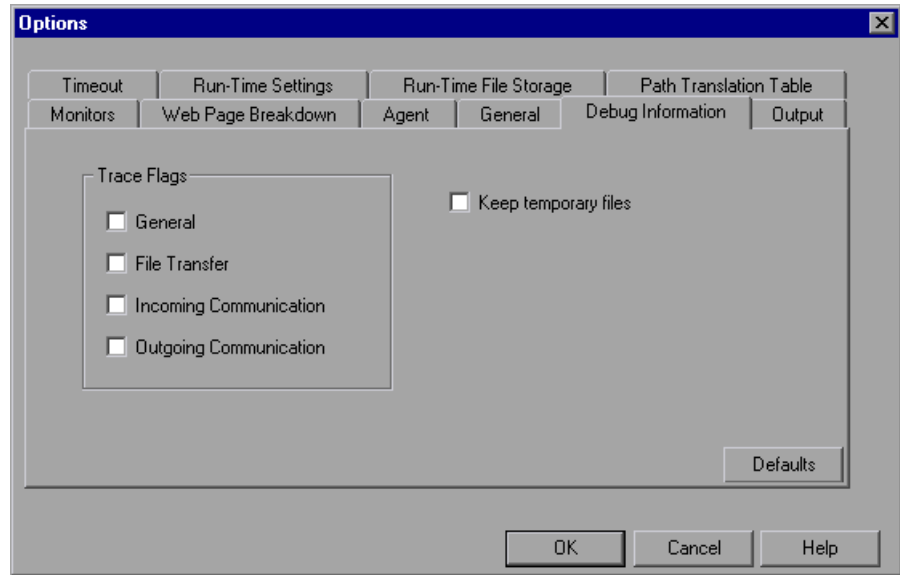
The Debug settings in the Options dialog box allow you to determine the extent of the trace to be performed during scenario execution. The debug information is written to the Output window.

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 LoadRunner agent and Controller 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 LoadRunner agent files are saved in *brr* folders in the TMP or TEMP directory of the agent machine. The Controller files are saved in *lrr* folders in the TMP or TEMP directory of the Controller machine. At the end of the scenario, all these files are automatically deleted. However, using the Debug Information Expert mode settings, you can instruct LoadRunner 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

The Output tab in the Options dialog box allows you to configure how running Vusers are displayed on the Controller machine.

Configuration of the 'Show Vuser' Operation: The following settings are available:

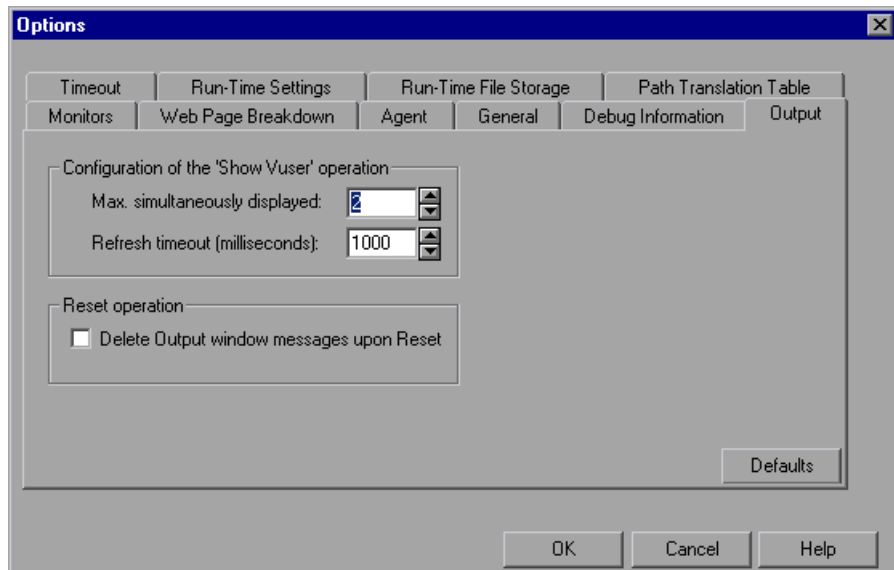
Max simultaneously displayed: Specifies the maximum number of Vuser logs that may be displayed simultaneously, as well as the maximum number of active UNIX, GUI, RTE, or Web Vusers that the Controller should display by opening up Run-Time Viewers on your machine. The default number is 10.

Refresh timeout: Defines how often to refresh the Vuser log. The default is every 1000 milliseconds.

Delete Output window messages upon Reset: Instructs LoadRunner to clear all messages in the Output window when you reset a scenario.

To set the Output settings:

- 1 Choose **Tools > Options**. The Options dialog box appears. Select the **Output** tab.

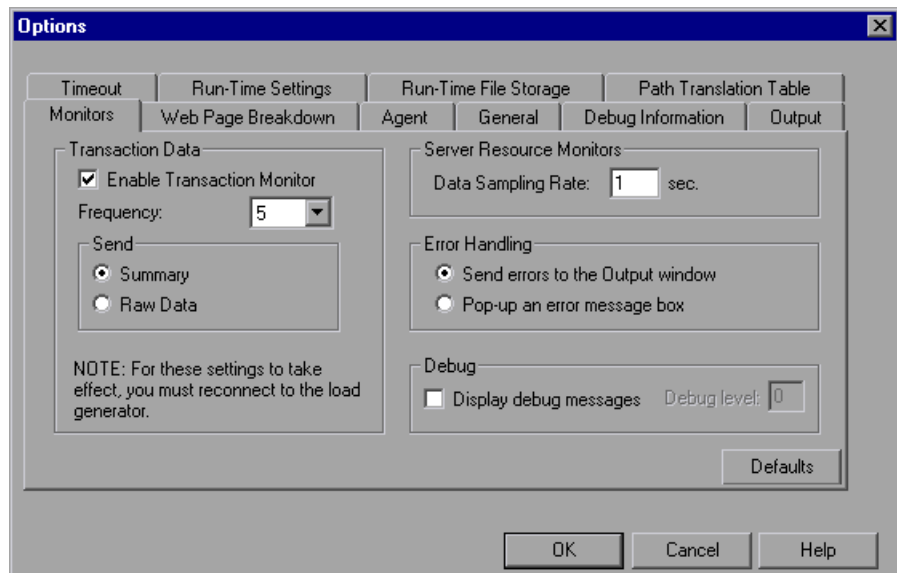


- 2 Specify the maximum number of Vuser logs to be displayed simultaneously, in the **Max. simultaneously displayed** box.
- 3 Specify the frequency at which LoadRunner refreshes the Vuser log, in the **Refresh timeout** box.
- 4 To clear the messages in the Output window when you reset a scenario, select 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 Controller, 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 Controller, 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.

The screenshot shows the 'Load Generator Information' dialog box with the 'UNIX Environment' tab selected. The 'Name' field is set to 'localhost', and the 'Platform' is set to 'UNIX'. The 'Temporary directory' field is empty. The checkbox 'Enable load generator to take part in the scenario' is checked. The 'Login as' section has three checkboxes: 'Name' (unchecked), 'Use lower case for login names' (unchecked), and 'Local user' (unchecked). The 'Shell Settings' section has a checkbox 'Don't use RSH' (unchecked), a 'Default shell' dropdown set to 'csh', and an empty 'Initialization command' field. A 'NOTE' at the bottom states: 'NOTE: This tab will not be relevant for a host running Vusers or monitoring over the firewall.' Buttons for 'OK', 'Cancel', 'Help', and 'Defaults' are visible.

User Limits		Firewall		WAN Emulation	
Status	Run-Time File Storage	UNIX Environment		Run-Time Quota	
Login as: <input type="checkbox"/> Name: <input type="text"/> <input type="checkbox"/> Use lower case for login names <input type="checkbox"/> Local user: <input type="text"/>					
Shell Settings: <input type="checkbox"/> Don't use RSH Default shell: <input type="text" value="csh"/> Initialization command: <input type="text"/>					
NOTE: This tab will not be relevant for a host running Vusers or monitoring over the firewall.					

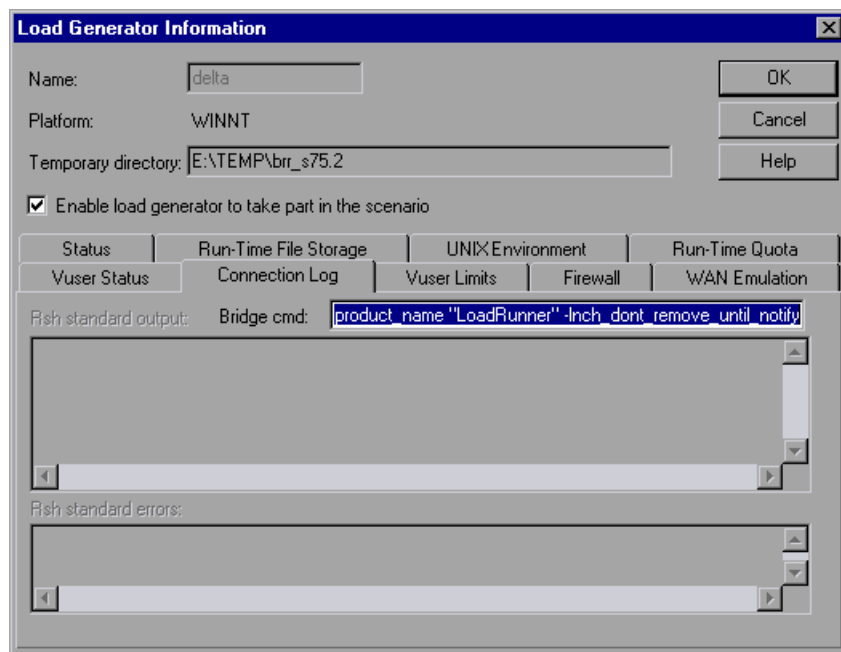
Load Generator Information - Connection Log Settings

The Connection Log tab in the Load Generator dialog box allows you to view the standard output and standard errors generated as the Controller connects to the selected UNIX load generator. You can also change the command that the Controller 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 **Scenario > 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 Controller 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 Controller to the remote bridge in order to connect the UNIX load generator.

D

Troubleshooting the Controller

LoadRunner enables you to test entire applications. If one of the components of the application is not configured properly, LoadRunner scenarios will not run.

This appendix discusses the most common LoadRunner problems:

- ▶ LoadRunner 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

About Troubleshooting

LoadRunner relies heavily upon communication between machines on a network. If communication is not established properly, the Controller will be unable to send commands to remote load generators and the scenario 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 scenario and not your Vuser script, you should verify that your script runs properly on all remote load generators as a stand-alone:

- Test your GUI Vuser scripts on Windows platforms using WinRunner.
- Test your Vuser scripts on UNIX platforms by running them from the command line.
- Test all other types of Vuser scripts on Windows platforms by running them from VuGen, or by running a single user from the Controller.

Note: When a test runs in VuGen, the full browser is used. This differs from a test run in the Controller, 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 Controller. Before running a scenario in the Controller with multiple Vusers, run a single Vuser to ensure the test is bug free.


For more information on running Vuser scripts in stand-alone mode, refer to the appropriate guide for creating Vuser scripts.

LoadRunner Communications

Most communication problems can be solved if you understand your LoadRunner 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 LoadRunner. There are five servers: The LoadRunner Controller, the Web server, the application server, the database server, and the file server which stores the scenario 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 Controller in both directions (send/receive), but with the file server in one direction (send). The Controller must have access to the file server. All Users participating in the scenario 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 scenario will fail.

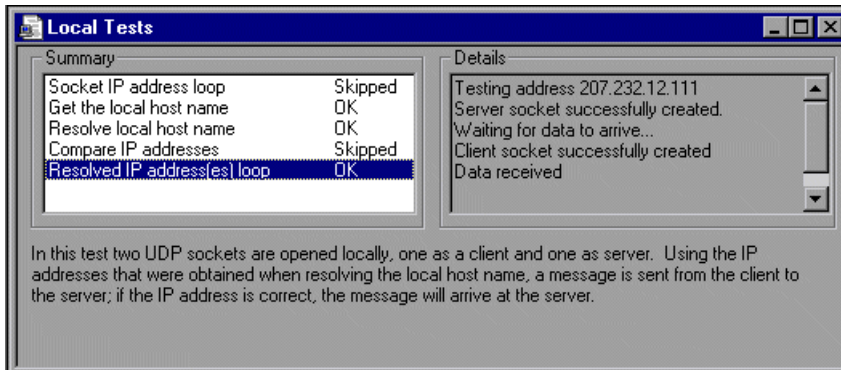
Failure to Communicate with a Load Generator

The most common communication error is the failure of the Controller 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. LoadRunner includes a utility called Hostinfo (hostinfo.exe), located under LoadRunner'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 Controller 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 Controller 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 Controller'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.)

Name	Status	Details
doc9pc	✓ Ready	
goose	✓ Ready	
miro	✓ Ready	
rman	✗ Failed	Connection to host failed. Communication problem: RPC: F
oxygen	✗ Failed	Connection to host failed. Communication problem: RPC: F
jukebox	✗ Failed	Connection to host failed. Communication problem: RPC: F
hammer	✗ Failed	Connection to host failed. Communication problem: RPC: F
steel	⚙️ Connecting	Connection started.

If your scenario uses several domains (for example, Vusers on a different domain than the Controller), the Controller may have trouble communicating with the load generators. This occurs because the Controller uses the short load generator name—not including the domain—by default. To solve this, you must tell the Controller to determine the full load generator names, including the domains.

Modify the *miccomm.ini* file in the Controller 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 Controller with Multiple IP Addresses

If the load generator machine does not recognize the Controller machine by its short name or full name, and the Controller machine has more than one IP address, you can define an alias name for the Controller 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 Controller 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 Controller 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 Controller 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 64.

To use UNIX without RSH:

- 1 On the UNIX Load Generator machine, run the agent daemon by running the following command from `<LoadRunner 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 `<LR_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 LoadRunner 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 Controller, 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 `<LR_root>/bin` directory: `m_agent_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 scenario. 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 LoadRunner scenario will fail if the result path or Vuser script is inaccessible to one or more of the participating machines. Check the following items:

- Path Translation
- Vuser Script
- Result Path

Path Translation: A script's location (path) is always based on the Controller 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 Controller'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 Controller 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, LoadRunner saves files locally and collates them to the Controller 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 Controller 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` Controller 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 scenario will fail. For more information about path translation, see Appendix B, “Performing Path Translation.”

Vuser Script: Make sure that the Vuser script is accessible to all load generators participating in the scenario through path translation and permissions. View or run the Vuser 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 scenario through path translation and permissions. Check the permissions of the result directory files and modify them if necessary.

Failed Vusers or Transactions

LoadRunner Vusers or transactions may fail for a variety of reasons relating to the network, database, or actual script. You can find information about scenario runs from the following sources:

- Run View
- Output Window
- Output File (excluding GUI Vusers)
- Analysis Reports and Graphs

Run View

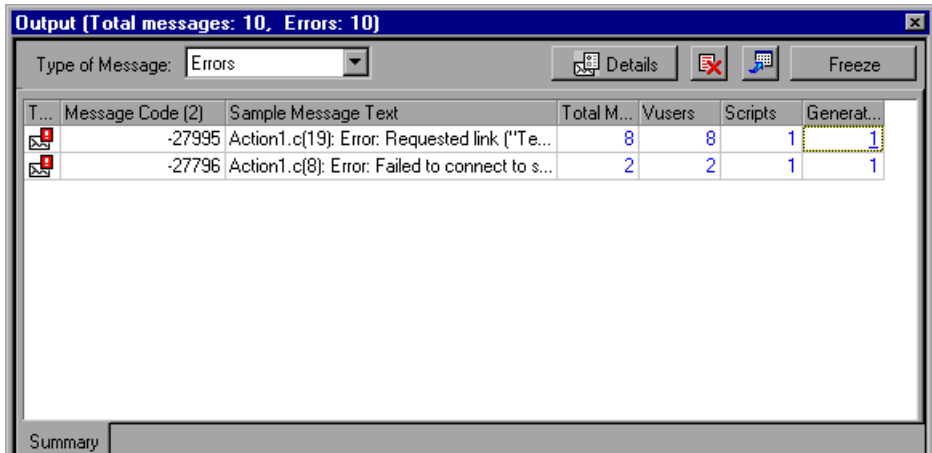
The Run view is part of the LoadRunner Controller. The Scenario Groups pane in the top left-hand corner of the view indicates the status of the Vuser groups during and after the scenario run. During the scenario run, the columns will show a PENDING, INITIALIZING, READY, RUNNING, or RENDEZVOUS status. 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, LoadRunner displays an error status. If a Vuser completes the script execution, LoadRunner 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 13, “Running a Scenario.”

ID	Status	Script	Load Generator	Elapsed Time
1	Done. Failed 1 iteration(s)	test2	localhost	00:01:06
2	Running	test2	localhost	00:01:10
3	Done. Failed 1 iteration(s)	test2	localhost	00:01:09
4	Running	test2	localhost	00:01:09
5	Running	test2	localhost	00:01:09
6	Running	test2	localhost	00:01:09
7	Done. Failed 1 iteration(s)	test2	localhost	00:01:08
8	Running	test2	localhost	00:01:09
9	Running	test2	localhost	00:01:09
10	Done. Failed 1 iteration(s)	test2	localhost	00:01:08

Output Window

View the Output window from the Controller. The output window contains useful information for debugging a scenario. 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 Controller, if you are using automation.



For more information about the Output window, see Chapter 14, “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 scenario
- 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 *LoadRunner Creating Vuser Scripts User's Guide*.

Analysis Reports and Graphs

You can generate graphs and reports to view information about the scenario run. For example, the Scenario Summary report displays a table containing the scenario'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.

Analysis Summary Period: 07/03/2001: 10:06:38 - 07/03/2001 10:06:38

Scenario Name: C:\Sanity_Analysis\Scenario\Scenario.lrs
Results in session: F:\results\Sessions\w\0\res_amazon\res_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 Total 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	2.815	3.575	2.86	12	0	0

For more information on the available graphs and reports, see the *LoadRunner 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.

LoadRunner 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 **LoadRunner agent log file** contains information about communication activity between the LoadRunner agent and the MI Listener.

- ▶ To open the file on Windows machines, right-click the LoadRunner agent icon in the system tray of the LoadRunner agent machine, and select **View Log**. Alternatively, open the latest `<temp_directory>\LoadRunner_agent_startup<unique identifier>.log` file (if the LoadRunner agent is a process), or `<temp_directory>\LoadRunner_agent_service<unique identifier>.log` file (if the LoadRunner 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->LoadRunner->Advanced Settings (or open file `<LR_root>\launch_service\dat\br_Inch_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 LoadRunner agent and the Controller.

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>\LoadRunner_agent_startup<unique identifier>.log` file (if the LoadRunner agent is a process), or `<temp_directory>\LoadRunner_agent_service<unique identifier>.log` file (if the LoadRunner agent is a service), in a text editor.

To increase the logging level, select **Start > Programs > LoadRunner > Advanced Settings > Agent Settings**, or open the `<LR_root>\launch_service\dat\br_Inch_server.cfg` file in a text editor. In the Log section, set AgentExtended to 1.

The **Controller log file** contains information about communication activity between the Controller 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 LoadRunner Agent and MI Listener

If there is a proper connection between the LoadRunner 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 LoadRunner agent machine. `Local_machine_key` is the value set in the Agent Configuration, as described in Chapter 15, "Working with Firewalls." The file will be removed when the LoadRunner agent disconnects from the MI Listener.
- On both UNIX and Windows platforms, the following message will appear in the LoadRunner agent log file: Notify Connected to MI Listener.

Note: The LoadRunner agent tries to connect to the MI Listener machine every Timeout seconds (as defined in the Agent Configuration). After a successful connection, if no Controller has connected through this MI Listener to the agent after another Timeout, the LoadRunner will disconnect from the Controller.

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 LoadRunner agent machine.

In both Windows and UNIX, the message Disconnected from MI Listener will appear in the LoadRunner agent log file.

UNIX Connection Errors

After installing the *m_agent_daemon* as described in Chapter 15, "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 <LR_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.

LoadRunner 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 LoadRunner 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 H, "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 H, “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 LoadRunner agent machine (see Appendix H, “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 LoadRunner agent's CA List.
- ▶ The MI Listener's certificate has a future verification date.

See Appendix H, "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 LoadRunner agent's certificate does not have a signature that is included in the MI Listener's CA List.
- ▶ The LoadRunner agent's certificate has a future verification date.

See Appendix H, "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 LoadRunner agent's certificate has expired.

General Connection Errors

These errors can occur when using all configurations.

If no errors appear both in the LoadRunner 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 `<LR_Installation>\dat\br_Inch_server.cfg` file on the LoadRunner agent machine, is set to 1.

Verifying Connection Between the Controller and Agent through the MI Listener

When there is a successful connection between the LoadRunner agent and the MI Listener, and the Controller machine fails to connect, you should check the following:

- ▶ The **Name** field in the Load Generators dialog in the Controller 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 Controller, 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 Controller machine recognizes the LoadRunner agent machine (e.g., by using the ping utility). If this fails, there is a configuration problem in the system not related to LoadRunner, and it must be solved before the connection can be made.
- ▶ Make sure that the Controller 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).

E

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, or SQL Server monitors by editing the *res_mon.dft* file found in the LoadRunner/dat directory.

To change the default counters:

- 1** Open a new scenario and click the **Run** tab.
- 2** For each of the monitors, select the counters you want to measure.
- 3** Save the scenario and open the scenario *.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.

F

Configuring Multiple IP Addresses

When you run a scenario, 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 Controller

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.

LoadRunner'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:

- ▶ **Application Deployment Solution:** Citrix ICA
- ▶ **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

LoadRunner 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 scenarios. For information about adding IP addresses on UNIX machines, see “Configuring Multiple IP Addresses on UNIX” on page 634.

The following procedure summarizes how to add new IP addresses to a load generator:

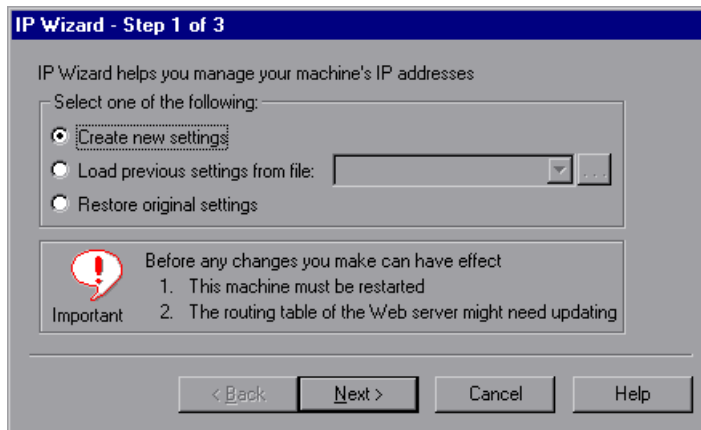
- 1** Run the IP Wizard on the load generator machine to add a specified number of IP addresses. Manually configure the new IP addresses for UNIX load generator machines.
- 2** Restart the machine.
- 3** Update the server’s routing table with the new addresses, if necessary.
- 4** Enable this feature from the Controller. Refer to “Enabling Multiple IP Addressing from the Controller” on page 636.

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 from the LoadRunner program group.




- 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.



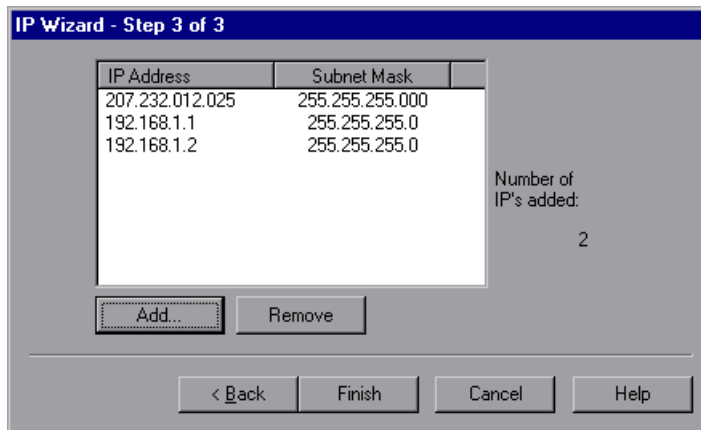
- 5 To check the server's routing table directly after adding the addresses, enter the server IP address. Refer to "Updating the Routing Table" on page 635 for more information.
- 6 Click **Next** to see a list of the machine's IP address(es). Click **Add** to define the range of addresses.




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 635 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 Controller

Once you define multiple IP addresses, you set an option to tell the Controller to use this feature.

To enable multiple IP addressing from the Controller:

- 1 In the Controller Design view, select **Scenario > Enable IP Spoofer**.

Note: You must select this option before connecting to a load generator.

- 2 Use the **General Options** of the Controller Expert Mode to specify how the Controller should implement this feature.

For more information, refer to Appendix C, “Working in Expert Mode.”

G

Controller Command Line Arguments

When you invoke the Controller from the command line, you can pass arguments to instruct the Controller how to behave. By passing arguments in the command line, you configure Controller scenario settings without the need to manually define them using the Controller UI.

This appendix describes:

- Invoking the Controller from the Command Line
- TestDirector Arguments
- Run-Time Arguments

About Controller Command Line Arguments

When invoked, the Controller checks all of the received arguments and sets its start-up environment accordingly. If no arguments are passed, the Controller uses its default settings.

For example, you can instruct Controller to Connect to TestDirector on start-up, save results to a directory other than the directory defined in the scenario, and invoke Analysis upon scenario termination.

Invoking the Controller from the Command Line

To invoke the Controller, type `wlrun` in the command line, followed by the arguments. Each argument should be preceded by a dash. **Note that the arguments are case-sensitive.** For example:

```
wlrun -TestPath C:\LoadRunner\scenario\Scenario.lrs -Run
```

When you invoke the Controller from the command line, the following rules apply:

- ▶ If the Controller is invoked with no arguments in the command line, the Controller uses its default settings.
- ▶ The Controller will always overwrite results.
- ▶ The controller will automatically terminate upon scenario termination and results will be collected. If you don't want the controller to automatically terminate upon scenario termination, add the flag `-DontClose` to the command line.
- ▶ The Controller launched through the command line behaves normally except when using the `-Run` option. Using the `-Run` option, dialogs and message boxes that usually open and require the user to close them in a usual launch, do not open in a command line launch.
- ▶ The Controller's settings are loaded from `wlrun5.ini`, located in Windows directory.

TestDirector Arguments

These arguments define the LoadRunner integration with TestDirector. For more information about the LoadRunner TestDirector integration, refer to Chapter 12, "Managing Scenarios Using TestDirector."

ConnectToTD	Specifies whether the Controller should connect to TestDirector on startup (0/1 or ON/OFF)
TDServer	TestDirector server name. Must be a machine where TestDirector is installed
TDDB	TestDirector database name, for example, "lrun"
UserName	User name for connecting to TestDirector
Password	Password corresponding to the user name
TestPath	Path to scenario in TestDirector database. For example, "[TD]\Subject\LoadRunner\Scenario1" If path includes blank spaces, use quotation marks.

TestId	Test ID (used by TestDirector only)
ResultCleanName	For use with ResultCycle only. Example: "Res1"
ResultCycle	TestDirector cycle. For example, "LR_60_SP1_247"

Note: The **ResultCycle** and **ResultCleanName** arguments are required if you wish to store the results within the TestDirector database.

Run-Time Arguments

These arguments specify the run-time related scenario settings. For more information on scenario settings, refer to Chapter 11, "Preparing to Run a Scenario."

TestPath	Path to the scenario, for example, C:\LoadRunner\scenario\Scenario.lrs This argument can also be used for a scenario residing in a TestDirector database. For example, "[TD]\Subject\LoadRunner\Scenario1" If the path includes blank spaces, use quotation marks.
Run	Runs the scenario, dumps all output messages into res_dir\output.txt and closes Controller
InvokeAnalysis	Instructs LoadRunner to invoke Analysis upon scenario termination. If this argument is not specified, LoadRunner uses the scenario default setting.
ResultName	Full results path. For example, "C:\Temp\Res_01"
ResultCleanName	Results name. For example, "Res_01"
ResultLocation	Results directory. For example, "C:\Temp"

Note: If the scenario doesn't specify a results directory, and one of the results arguments was not passed, the scenario will not run.

H

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 LoadRunner agent, it always sends its certificate as well (this is the server-side certificate). The LoadRunner agent can be configured to authenticate the certificate which it received, as described in Chapter 15, “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 LoadRunner 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 15, “Working with Firewalls.” If the LoadRunner agent owns a certificate, it sends it to the

MI Listener for the same authentication process. If the LoadRunner agent does not own a certificate, the communication might not be continued.


An SSL CA list and an SSL Certificate are included in each LoadRunner installation. This certificate is the same for all LoadRunner 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 <LoadRunner 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 (`cacert.cer`), and the CA Private Key (`capvk.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 <LoadRunner 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.

Index

A

- Acrobat Reader xi
- activating rstatd 289
- Add Destination Machines for Network Delay Monitoring dialog box 312
- Add Group dialog box 52
- Add Load Generator dialog box 62
- Add Machine dialog box
 - DB2 monitor 433
- Add Oracle Measurements dialog box 451
- Add SAP Monitor Measurements dialog box 476
- Add Script dialog box 108
- Add Sybase Measurements dialog box 460
- Add TUXEDO Measurements dialog box 565
- Add Vusers dialog box 58
- adding measurements 246
- agent
 - daemon 606
 - defined 7
 - settings 590
 - summary window 193
- Antara FlameThrower
 - monitor 292–304
- Apache
 - Add Measurements dialog box 324
 - monitor 324
- application
 - analyzing 20–22
 - configuration 21
 - usage model 22
- Application Deployment Solutions monitors 539–549
 - Citrix MetaFrame XP monitor 540–549

Ariba

- monitor 344
- Monitor Configuration dialog box 344

ASP

- monitor 387

ATG Dynamo

- monitor 348
- Resources dialog box 348

B

BEA WebLogic

- monitor 401
- Monitor Configuration dialog box 403

Books Online xi

BroadVision

- monitor 355
- Monitor Configuration dialog box 356

C

CA 641

Check Point FireWall-1

- monitor 319
- SNMP Resources dialog box 319

Citrix MetaFrame XP monitor 540–549

ColdFusion

- dialog box 365
- monitor 365

collating scenario results 149

command line

- arguments 637
- options, Vuser script 81

- configuring
 - load generator 61–63
 - load generator settings 64–79
 - scenario 137–144
 - script 80–83
 - script (goal-oriented scenario) 108–112
 - script (Percentage Mode) 92–95
 - Vusers 56
- connecting
 - to database 607
 - to TestDirector 154
- Connection Log tab 597
- Context Sensitive Help xii
- Controller 38
 - defined 6
 - invoking 35
 - managing scenario files 42
 - overview 38
 - quick tour 35–46
 - running from the command line 637
- Controller window
 - Design view 40
 - icon bar 40
 - menu bar 38
 - Output window 40
 - Run view 40
 - Status bar 38
 - title bar 38
 - Toolbar 38
- controller_host 585
- controller_path 586
- converting a scenario
 - to Percentage mode 88
 - to Vuser Group mode 95
- counters, for stress testing 624
- creating
 - goal-oriented scenario 97–112
 - manual scenario 49–85
 - manual scenario in Percentage Mode 87–95
 - Vuser groups 50–55
 - Vuser scripts 16
- cross result analysis, naming result sets 146
- custom queries
 - Oracle monitor 454

D

- Data Points graph (online) 267
- database
 - connecting to 607
- Database Server Resource monitors 431–466
 - DB2 monitor 432
 - Oracle monitor 449
 - SQL Server monitor 456
 - Sybase monitor 459
- DB2
 - monitor 432
 - Monitor Configuration dialog box 433
- debug
 - information settings 592
 - level 223
- default counter, changing 623
- Details button 189
- digital certificate
 - agent configuration settings 206
 - MI Listener configuration settings 209
 - overview 641
- disabling
 - Vuser groups 54
 - Vuser scripts (goal-oriented) 112
 - Vuser scripts (manual scenario) 95
- disconnecting from TestDirector 156
- documentation set xii
- Done.Failed - Vuser state
 - Scenario Groups pane 182
- Done.Passed - Vuser state
 - Scenario Groups pane 182
- Down - Vuser state
 - Scenario Groups pane 182
- Duration 119

E

- Edit Scenario Goal dialog box 101
 - Load Behavior tab 101
 - Load Preview 101
 - Scenario Settings tab 101
- EJB
 - monitor 482–498
- ERP Server Resource monitors 475–480

- Error - Vuser state
 - Running Vusers graph 266
 - Scenario Groups pane 182
- error handling 223
- Error Statistics graph 266
- Ethernet-bus based network 575
- Execution Notes dialog box 192
- Exiting - Vuser state
 - Scenario Groups pane 182
- Expert mode 589–598
 - agent settings 590
 - connecting to UNIX load generator 597
 - debug settings 592
 - general settings 591
 - monitor settings 595
 - output settings 594

F

- files, Vuser script 82
- Finished - Vuser state
 - Running Vusers graph 266
- firewalls
 - configuring agent to operate over firewall 205
 - Firewall Server monitors 319–322
 - installation configurations 196
 - installing MI_Listener 208
 - monitors over 237–247
 - network monitoring 315
 - preparing for data collection 243
 - running Vusers 195–211
 - troubleshooting 614
- Fujitsu INTERSTAGE
 - monitor 368
 - SNMP Resources dialog box 368
- Function Reference xi

G

- goal-oriented scenario 97–112
 - assigning load generators to scripts 106
 - assigning percentage of target to scripts 106

- goal-oriented scenario (cont'd)
 - defining goals 100
 - Select Scenario Type 36
- Gradual Exiting - Vuser state
 - Scenario Groups pane 182
- Graph Configuration dialog box 227
- graph time 226
- graphs, *See* online graphs
- Group Information dialog box 55
- GUI Vusers, defined 9

H

- hardware
 - checking communications 600
 - selecting for testing 28–30
- Hits per Second graph 274
- hme0 device 634
- hostid, IP address component 632
- Hostinfo utility 602
- hosts file 603
- HP, configuring IP addresses 635
- HTTP
 - Response per Second graph 275

I

- IBM WebSphere MQ monitor 552–563
- IIS monitor 327
- increasing number of Vusers 613
- initialization quota 68
- Initialize 173
- Initializing - Vuser state
 - Scenario Groups pane 182
- installing
 - LoadRunner, *See the LoadRunner Installation Guide*
 - Monitors over Firewall 238
- interpreting online graphs 579–582
- invoking the Controller 35
- IP addresses
 - adding to a load generator 629
 - class 632
 - configuring multiple 627–636
 - configuring on HP 635
 - configuring on Linux 635

- IP addresses (cont'd)
 - configuring on Solaris 634
 - enabling from the Controller 636
 - hostid 632
 - IP Wizard 629
 - load generator machine 627
 - netid 632
 - submask 632
- iPlanet (NAS)
 - dialog box 374, 376
 - monitor 370
- iPlanet (SNMP)
 - dialog box 334
 - monitor 334
- iPlanet/Netscape
 - Add Measurements dialog box 330
 - monitor 329
- J**
- J2EE
 - monitor 522–538
- Java
 - Java Performance monitors 481–520, 521–538
- JProbe
 - dialog box (monitor) 500
 - monitor 498–501
- L**
- Linux
 - configuring IP addresses 635
- lists 16
 - load generator list 61–63
 - rendezvous list 127–135
 - script list 80–83
 - script list (goal-oriented scenario) 108–112
 - script list (Percentage Mode) 92–95
 - Vuser group list 50–55
- load balancing 107
- load generator configuration
 - checking Controller communication 601
 - connecting load generators 62
 - disabling load generators 62
- load generator configuration (cont'd)
 - disconnecting load generators 62
 - enabling load generators 62
 - Expert mode 596
 - firewall 70
 - initializing quota 68
 - limiting of Vusers 69
 - run-time files 65
 - selecting load generators 58
 - UNIX shell 66
 - WAN emulation 72
- Load Generator Information dialog box 64
 - Firewall tab 70
 - Run-Time File Storage tab 65
 - Run-Time Quota tab 68
 - Status tab 64
 - Unix Environment tab 66
 - Vuser Limits tab 69
 - Vuser Status tab 71
 - WAN Emulation Advanced options 75
 - WAN Emulation tab 73
- load generators 6
 - adding 63
 - adding an IP address 629
 - balancing 107
 - configuration 61–63
 - defined 6
 - modifying 63
 - multiple IP addresses 591
 - selecting 58
 - selecting (Percentage Mode) 90
 - setting attributes 64–79
 - viewing load generator details 63
- Load Generators window 62
- load testing, defined 3
- LoadRunner
 - application testing 3
 - emulating human users with Vusers 5
 - implementation planning 25
 - overview 3–13
 - testing process, *See* testing process
 - working with LoadRunner 6
- LoadRunner Controller User's Guide (Windows) 3
- lr_user_data_point 267

M

- manual scenario
 - creating 49–85
 - defined 36
 - Percentage Mode 87–95
- Measurement Configuration dialog box
 - Configuration tab 232
 - Description tab 232
- measurement frequency, setting 247
- Media Player Client
 - monitor 473
- MI_Listener 208
 - installing 243
- Microsoft
 - ASP monitor 387
 - IIS monitor 327
- middleware
 - response time measurements 26
 - system configuration 21
- Middleware Performance monitors 551–568
 - IBM WebSphere MQ monitor 552–563
 - TUXEDO monitor 564–568
- monitors
 - Application Deployment Solutions 539–549
 - database server resources 431–466
 - ERP Server Resources 475–480
 - Firewall Server 319–322
 - Java Performance 481–520, 521–538
 - Middleware Performance 551–568
 - network 307–317
 - online 215–218
 - run-time 266
 - streaming media 467–473
 - system resources 279–306
 - transaction 268
 - Web application server resources 343–430
 - Web resources 273–278
 - Web server resources 323–333
- monitors over firewall 237–247
 - adding and removing measurements 246
 - configuring measurement frequency 247

- monitors over firewall (cont'd)
 - configuring properties 244
 - installing 238
 - installing MI_Listener 243
 - overview 238
 - preparing for data collection 243
- MQ monitor 552–563
- MS Active Server Pages dialog box 387
- MS IIS dialog box 327
- MS SQL Server dialog box 456
- multiple IP addresses 591
 - connecting to Controller 604
 - enabling 591

N

- naming results 146
- netid, IP address component 632
- Network
 - Breakdown dialog box 317
 - Delay Time dialog box 312
 - Delay Time graph 316
 - Monitor Settings for Defined Path dialog box 312
 - Network Delay options 227
- Network monitor 307–317
 - configuring 312
 - determining bottlenecks 308
 - monitoring over a firewall 315
 - on UNIX 309
 - overview 307
- network segment delay, viewing 317
- New Monitored Server Properties dialog box 244
- New Scenario dialog box 36

O

- online graphs 225
 - configuring 225
 - customizing display view 222
 - data point 267
 - exporting 236
 - interpreting 579–582
 - merging two graphs 229
 - modifying a measurement scale 233

- online graphs (cont'd)
 - opening graphs 220
 - remote monitoring 249
 - sampling rate 223
 - viewing data offline 236
 - x-axis style 226
 - y-axis style 226
- online monitors 215–218
 - changing default counters 623
 - configuring graphs 225
 - configuring measurements 231
 - debugging 223
 - display type 226
 - error handling 223
 - graph time 226
 - graphs 230
 - interpreting online monitoring
 - graphs 579
 - line color 232
 - pausing 231
 - show/hide lines 234
 - starting 218
 - viewing data offline 236
- online transaction monitoring
 - adding transactions 270
 - graphs 268
 - setup 269
- online Web server resource monitoring
 - using a proxy server 342
- Open a New Graph dialog box 221
- Open Scenario from TestDirector project dialog box 157
- Options dialog box
 - Agent tab 590
 - Debug Information tab 593
 - General tab 591
 - Path Translation Table tab 586
 - Run-Time File Storage tab 143
 - Run-Time Settings tab 139
 - Timeout tab 141
 - Web Page Breakdown tab 271
- Oracle
 - custom queries 454
 - Logon dialog box 451
 - monitor 449

- Oracle9iAS HTTP monitor 389
 - Server Monitor Configuration dialog box 389
- output file 611
- Output window 187–190
 - clearing 189
 - debugging information 610
 - drilling down on log information 189
 - filtering messages 188
 - refreshing 190
 - saving messages to a file 189
 - sorting messages 188
 - viewing 40
 - viewing message detail 189
- Overlay Graphs dialog box 229

P

- packets 308
- Pages Downloaded per Second graph 277
- path translation
 - debugging file locations 607
 - defined 583
 - editing the Path Translation Table 586
 - examples 588
 - scenario configuration 144
 - script path 85
 - using the Path Translation Table 585
- pausing
 - monitors 231
 - Vusers 174
- Pending - Vuser state
 - Scenario Groups pane 182
- Percentage Mode
 - creating a scenario 87–95
- Percentage mode
 - assigning load generators to scripts 90
 - assigning percentage of total Vusers to scripts 90
 - converting a scenario to Percentage Mode 88
 - converting to Vuser Group mode 95
 - defining the total number of Vusers 88

- performance analysis, defined 6
 - performance measurement scope 25
 - planning load tests 19–33
 - Proxy Server 342
- R**
- Ramp Down 120
 - Ramp Up 118
 - Ready - Vuser state
 - Running Vusers graph 266
 - Scenario Groups pane 182
 - Real Server dialog box 470
 - RealPlayer
 - Client monitor 472
 - Server monitor 470
 - Refresh button 81
 - registry, modifying 613
 - relative script paths 83
 - Remote Agent Dispatcher (Process) 7
 - Remote Performance monitoring 249–263
 - configuring graph measurements 262
 - configuring graph settings 259
 - configuring the user settings 251
 - connecting to the monitor 254
 - installing the monitor 250
 - refreshing graphs 261
 - scaling graphs 259
 - selecting graphs 256
 - remote_host 586
 - remote_path 586
 - removing measurements 246
 - rendezvous 127–135
 - attributes 129
 - defined 5
 - disabling Vusers 133
 - enabling Vusers 133
 - information 135
 - manually releasing Vusers 175
 - overview 127
 - setting release policy 132
 - setting the attributes 129
 - setting timeout policy 132
 - Vuser state 182
 - Rendezvous Information dialog box 135
 - renumbering Vusers 174
 - response time measurement
 - end-to-end 25
 - GUI 26
 - middleware-to-server 26
 - network and server 25
 - server 26
 - results 146
 - collating 149
 - directory file structure 148
 - files for debugging 607
 - location in TestDirector project 160
 - naming 146
 - specifying location for 146
 - Retries per Second graph 278
 - routing table 635
 - rsh
 - checking Controller connection 605
 - connection for UNIX network
 - monitor 310
 - running UNIX without 605
 - rstatd process
 - activating 289
 - resource monitors 289
 - RTE Vuser scripts
 - in the Controller 12
 - Run/Stop Vusers dialog box 176
 - Running - Vuser state
 - Running Vusers graph 266
 - Scenario Groups pane 182
 - running over firewall 195–211
 - Run-Time graphs 265–270
 - interpreting 579
 - Run-Time Settings
 - group 59
 - script 59
 - shared 60
 - Run-Time settings
 - in the Controller 110
 - setting in the Controller 59, 80
 - Run-Time Viewer
 - view replay from Controller 173

S

sampling rate 223

SAP

monitor 476

Monitor Logon dialog box 476

Save Scenario to TestDirector project dialog box 158

scenario

closing 43

collating results 149

configuring 137–144

conversion to Percentage mode 88

conversion to Vuser Group mode 95

creating a goal-oriented scenario 97–112

creating a manual scenario 49–85

creating a manual scenario in Percentage Mode 87–95

creating a new scenario 42

defined 5

defining scenario goals 100

goal-oriented scenario defined 36

managing scenario files 42–43, 44

manual scenario defined 36

opening 42

opening from TestDirector 157

overview 50

overview of running 44–46

preparing to run 145–151

result directory 148

running 167–175

saving 43

saving to TestDirector 158

Scenario Goal pane 100

Scenario Scripts pane 100

Scenario Start dialog box 115

scheduling 117

summary information 151

viewing output messages 187

scenario configuration

path translation 144

run-time file location 142

run-time settings 138

specifying results location 146

timeout intervals 139

scenario execution 167–175

activating additional Vusers 176

controlling individual Vuser groups 170

controlling individual Vusers 173

delaying 115, 117

limiting duration 117

loading Vuser groups 170

loading Vusers 174

manually releasing Vusers from a rendezvous 175

messages 187

monitoring active Vusers 182

overview 44–46, 167

pausing Vuser groups 171

running scenarios unattended 169

running Vuser groups 171

stopping Vuser groups 172

Schedule Builder 113–125

creating a schedule 116

deleting a schedule 117

modifying a schedule 117

renaming a schedule 117

scenario execution 117

Schedule Builder dialog box 117

selecting a schedule 116–125

Vuser groups 121

Schedule Definition dialog box

Duration tab 119

Ramp Down tab 120

Ramp Up tab 118

Schedule Definition dialog box (Group)

Duration tab 123

Ramp Down tab 124

Ramp Up tab 122

Start Time tab 121

Script Information dialog box 109

script paths, relative 83

scripts, *See* Vuser scripts

server monitors

adding and removing measurements 246

configuring properties 244

setting a measurement frequency 247

Server Monitors dialog box 244

server routing table 635

- Set Results Directory dialog box
 - local or remote location 147
 - TestDirector project 160
 - settings
 - agent 590
 - debug 592
 - general 591
 - load generator 64–79
 - measurement frequency 247
 - monitors 595
 - output 594
 - timeout 139
 - show/hide a measurement
 - online monitors 231
 - Transaction monitor 234
 - SilverStream
 - Add Measurements dialog box 395
 - monitor 394
 - SiteScope monitor 304–306
 - Sitraka JMonitor 501–520
 - SNMP
 - Resources monitor 290–292
 - SNMP Resources
 - dialog box 290
 - Solaris
 - configuring IP addresses 634
 - SQL Server
 - dialog box 456
 - monitor 456
 - SSL
 - agent configuration settings 206
 - MI Listener configuration settings 209
 - overview 641
 - Status bar 38
 - Stopped - Vuser state
 - Scenario Groups pane 182
 - stopping
 - Vuser groups 172
 - Vusers 173
 - Streaming Media monitors 467–473
 - Media Player Client monitor 473
 - RealPlayer Client monitor 472
 - RealPlayer Server monitor 470
 - Windows Media Server monitor 468
 - Summary Information dialog box 151
 - Support Information xii
 - Support Online xii
 - Sybase
 - Logon dialog box 460
 - monitor 459
 - System Resource monitors 279–306, –306
 - Antara FlameThrower monitor 292–304
 - SiteScope monitor 304–306
 - SNMP Resources monitor 290–292
 - UNIX Resources monitor 286–288
- T**
- TCP/IP setup 602
 - test objectives
 - defining 23–24
 - summary of common objectives 31–33
 - TestDirector
 - adding Vuser scripts 161
 - command line arguments 638
 - connecting to 154
 - disconnecting from 156
 - integration 153–163
 - managing scripts 153
 - opening a scenario 157
 - saving results to TestDirector 160
 - saving scenarios to TestDirector 158
 - TestDirector Connection dialog box 154
 - testing process
 - analyzing test results 18
 - creating the scenario 16
 - creating the Vuser scripts 16
 - monitoring a scenario 18
 - planning the test 16
 - running the scenario 17
 - Throughput graph 274
 - timeout
 - Controller settings 139
 - Toolbar 38
 - Total Transactions per Second (Passed) graph 268
 - Transaction monitor 265–270

- transactions 5
 - deciding which to define 27
 - defined 5
 - failed 609
 - Transaction Response Time graph 268
 - Transactions dialog box 184
 - Transactions per Second (Failed, Stopped) graph 268
 - Transactions per Second (Passed) graph 268
- troubleshooting
 - Controller 599–621
 - firewalls 614
 - monitors 569–576
 - network considerations 574
- TUXEDO
 - monitor 564–568
 - monitor measurements 567

U

- UNIX
 - activating rstatd 289
 - connection to load generator 597
 - Resources monitor 286–288
 - rsh 605
 - shell 605
 - without rsh 605
- Unix Kernel Statistics dialog box 286
- User-Defined Data Points graph 267

V

- viewing Vusers 181–193
 - agent summary 193
 - Output window 187
 - overview 181
- Virtual users, *See* Vusers
- Vuser
 - group list 50–55
- Vuser groups
 - adding Vusers to 58
 - creating 50–55
 - deleting 54
 - enabling/disabling 54
 - loading 170

- Vuser groups (cont'd)
 - modifying 54
 - pausing 171
 - running 171
 - scheduling 121
 - stopping 172
- Vuser scripts 5
 - adding (goal-oriented scenario) 108
 - adding (Percentage Mode) 92
 - adding from TestDirector 161
 - assigning load generators to (goal-oriented scenario) 106
 - assigning load generators to (Percentage mode) 90
 - assigning percentage of defined target to (goal-oriented scenario) 106
 - assigning percentage of total Vusers to (Percentage mode) 90
 - command line options 81
 - configuring 80–83
 - creating 16
 - defined 5
 - deleting (goal-oriented scenario) 112
 - deleting (Percentage Mode) 95
 - editing 82
 - enabling/disabling (goal-oriented scenario) 112
 - enabling/disabling (manual scenario) 95
 - files 82
 - modifying script details 80, 82
 - selecting for group 59
 - selecting for load generator 53
 - selecting for scenario 37
- Vuser states
 - Run-Time graphs 266
- Vusers
 - activating additional during scenario execution 176
 - adding to a group 58
 - configuring 56
 - deciding how many to run 27
 - defined 5
 - emulating maximum user load 127
 - error, warning, and notification messages 187

- Vusers (cont'd)
 - GUI Vusers 9
 - loading 174
 - monitoring 182
 - pausing 174
 - renumbering 174
 - RTE Vusers 12
 - running 173
 - status in Scenario Groups pane 182
 - stopping 173
 - types of 8
 - viewing 181–193
 - User Information dialog box 82
 - User Log 190
 - User script log 175
 - User window 45
 - Vusers with Error Statistics graph 266
- Vusers window 57
- W**
- WAN emulation 72
 - advanced options 75
 - configuring the settings 73
 - excluding IPs 76
 - predefined profile settings 74
 - setting up 73
 - stopping and restarting 78
- Web Application Server Resource monitors 343–430
 - Ariba monitor 344
 - ATG Dynamo monitor 348
 - BroadVision monitor 355
 - ColdFusion monitor 365
 - Fujitsu INTERSTAGE monitor 368
 - iPlanet (NAS) 370
 - Microsoft ASP 387
 - Oracle9iAS HTTP monitor 389
 - SilverStream monitor 394
 - WebLogic (JMX) monitor 401
 - WebLogic monitor 397
 - WebSphere (EPM) monitor 421
 - WebSphere monitor 407
- Web page breakdown, enabling 271
- Web Resource monitors 273–278
 - Web Server Resource monitors 323–333
 - Apache monitor 324
 - iPlanet (SNMP) monitor 334
 - iPlanet/Netscape monitor 329
 - Microsoft IIS monitor 327
 - WebLogic
 - (JMX) monitor 401
 - (SNMP) Resources dialog box 398
 - monitor 397
 - WebSphere
 - (EPM) monitor 421
 - (EPM) Monitor Configuration dialog box 429
 - monitor 407
 - Monitor Configuration dialog box 410
- Windows
 - Media Server - Add Measurements dialog box 468
 - Media Server monitor 468
 - Resources dialog box 281
 - Resources monitor 281–286
- Windows Resource monitors
 - Windows Resources monitor 281–286

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.



MERCURY INTERACTIVE

Mercury Interactive Corporation

1325 Borregas Avenue
Sunnyvale, CA 94089 USA

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

Customer Support: support.mercuryinteractive.com



LRCTRUG7. 6/ 03