

OPTIMIZE

MERCURY DIAGNOSTICS™ FOR J2EE, .NET & ERP/CRM

VERSION 4.1

Installation and Configuration Guide

MERCURY™

BUSINESS TECHNOLOGY OPTIMIZATION

Mercury Diagnostics

Installation and Configuration Guide

Version 4.1

Document Release Date: June 23, 2006

MERCURY™

Mercury Diagnostics Installation and Configuration Guide Version 4.1

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: United States: 5,511,185; 5,657,438; 5,701,139; 5,870,559; 5,958,008; 5,974,572; 6,137,782; 6,138,157; 6,144,962; 6,205,122; 6,237,006; 6,341,310; 6,360,332; 6,449,739; 6,470,383; 6,477,483; 6,549,944; 6,560,564; 6,564,342; 6,587,969; 6,631,408; 6,631,411; 6,633,912; 6,694,288; 6,738,813; 6,738,933; 6,754,701; 6,792,460 and 6,810,494. Australia: 763468 and 762554. Other patents pending. All rights reserved.

Mercury, Mercury Interactive, the Mercury logo, the Mercury Interactive logo, LoadRunner, WinRunner, SiteScope and TestDirector are trademarks of Mercury Interactive Corporation and may be registered in certain jurisdictions. The absence of a trademark from this list does not constitute a waiver of Mercury's intellectual property rights concerning that trademark.

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

Mercury provides links to external third-party Web sites to help you find supplemental information. Site content and availability may change without notice. Mercury makes no representations or warranties whatsoever as to site content or availability.

Mercury Interactive Corporation
379 North Whisman Road
Mountain View, CA 94043
Tel: (650) 603-5200
Toll Free: (800) TEST-911
Customer Support: (877) TEST-HLP
Fax: (650) 603-5300

© 2004 - 2006 Mercury Interactive Corporation, All rights reserved

If you have any comments or suggestions regarding this document, please send them by e-mail to documentation@mercury.com.

Table of Contents

Welcome to This Guide	xiii
How This Guide Is Organized	xiii
Who Should Read This Guide	xiv
Mercury Diagnostics Online Documentation.....	xv
Additional Online Resources.....	xvii
Documentation Updates	xvii
Typographical Conventions.....	xviii

PART I: INTRODUCTION

Chapter 1: Preparing to Install Mercury Diagnostics	3
About Mercury Diagnostics 4.1	4
Mercury Diagnostics Components.....	5
Licensing Mercury Diagnostics	7
Supported Application Servers	7
Mercury Diagnostics Data Flow	8
Recommended Deployment Configuration	9
Host System Requirements for the Diagnostics Components	14
Upgrading from Earlier Versions of Diagnostics.....	18
Other Pre-installation Considerations	19
Recommended Order of Installation.....	20
Planning the Installation	21

PART II: INSTALLATION OF THE MERCURY DIAGNOSTICS SERVER

Chapter 2: Installing the Mercury Diagnostics Server.....29
Installing the Diagnostics Server on a Windows Machine30
Installing the Diagnostics Server on a UNIX Machine41
Starting and Stopping the Diagnostics Server52
Verifying the Diagnostics Server Installation54
Licensing the Diagnostics Server in Commander Mode.....55
Configuring the Diagnostics Server55
Determining the Version of the Diagnostics Server56
Uninstalling the Diagnostics Server.....57
Troubleshooting Diagnostics Server Issues58

Chapter 3: Licensing Mercury Diagnostics59
About Mercury Diagnostics Licensing59
Types of Licenses60
Licensing the Diagnostics Server in Commander Mode.....60

PART III: INSTALLATION AND CONFIGURATION OF THE MERCURY DIAGNOSTICS PROBES

Chapter 4: Installing the Diagnostics Collector.....67
About Installing the Diagnostics Collector68
Installing the Diagnostics Collector on a Windows Machine68
Installing the Diagnostics Collector on a UNIX Machine76
Silent Installation of the Diagnostics Collector83
Configuring the Active System Property Files84
Verifying the Diagnostics Collector Installation87
Starting and Stopping the Diagnostics Collector.....89
Determining the Version of the Diagnostics Collector90
Uninstalling the Diagnostics Collector.....90

Chapter 5: Installing the Mercury Diagnostics Probe for .NET.....91
About the Mercury Diagnostics Probe for .NET91
Installing the .NET Probe92
Verifying the .NET Probe Installation106
Configuring the .NET Probe108
Determining the Version of the .NET Probe109
Uninstalling the .NET Probe109

Chapter 6: Installing the Mercury Diagnostics Probe for J2EE.....	111
About Installing and Configuring the J2EE Probe	112
Installing the J2EE Probe on a Windows Machine	113
Installing the J2EE Probe on a UNIX Machine	130
Installing the J2EE Probe on a z/OS Mainframe	146
Installing the J2EE Probe Using the Generic Installer	149
Silent Installation of the J2EE Probe	150
Chapter 7: Post-Installation Configuration of the	
Mercury Diagnostics Probe for J2EE	153
Configuring the Probe Manually	154
Configuring the Application Server	157
Configuring the Probe for Advanced Situations	158
Deploying BEA JRockit Mission Control.....	158
Verifying the J2EE Probe Installation.....	158
Determining the Version of the J2EE Probe.....	160
Upgrading to a Newer Version of the J2EE Probe	161
Uninstalling the J2EE Probe	164
Chapter 8: Running the JRE Instrumenter	165
About the JRE Instrumenter	165
Running the JRE Instrumenter.....	166
Chapter 9: Configuring the Application Servers	171
About Configuring the Application Server	172
Configuring WebSphere Application Servers.....	173
Configuring WebLogic Application Servers	188
WebLogic 6.1	189
Configuring the Oracle9i Application Server	198
Configuring the JBoss Application Server	201
Configuring the SAP NetWeaver Application Server	204
Configuring a Generic Application Server	206
Adjusting the Heap Size for the J2EE Probe in the Application	
Startup Script.....	208

PART IV: SETTING UP INTEGRATION WITH OTHER MERCURY PRODUCTS

Chapter 10: Setting Up Mercury Business Availability Center 6.x to Use Diagnostics211
About Setting Up Mercury Business Availability Center 6.x to use Diagnostics212
Specifying the Diagnostics Server Details214
Enabling the Diagnostics View220
Enabling Diagnostics Monitoring in Business Process Monitor (BPM) Related Views222
Assigning Permissions for Diagnostics Users224
Accessing the Diagnostics Pages in Windows 2003226

Chapter 11: Installing LoadRunner 8.1 and the LoadRunner Diagnostics AddIn227
Understanding the LoadRunner Diagnostics AddIn227
Installing Mercury LoadRunner228
Installing the Mercury LoadRunner Diagnostics AddIn228
Configuring LoadRunner for Mercury Diagnostics230

Chapter 12: Setting Up LoadRunner 8.1 to use Mercury Diagnostics231
About Setting Up LoadRunner to Use Mercury Diagnostics231
Specifying the Diagnostics Server Details232
Configuring LoadRunner Scenarios to use Mercury Diagnostics234

Chapter 13: Setting Up Performance Center 8.1 to Use Mercury Diagnostics235
About Setting Up Performance Center to Use Mercury Diagnostics235
Specifying the Diagnostics Server Details236
Configuring Performance Center Load Tests to Use | Mercury Diagnostics238

PART V: ADVANCED COMPONENT CONFIGURATION

Chapter 14: Advanced Diagnostics Server Configuration	241
Synchronizing Time Between Diagnostics Components.....	242
Configuring the Diagnostics Server for a Large Installation.....	247
Overriding the Default Diagnostics Server Host Name.....	250
Changing the Default Diagnostics Server Port	250
Configuring the Diagnostics Server for Multi-Homed Environments.....	252
Reducing Diagnostics Server Memory Usage	256
LoadRunner / Performance Center Diagnostics Server Assignments	257
Configuring Diagnostics Using the Diagnostics Server Configuration Pages.....	258
Configuring the Diagnostics Server for the LoadRunner Offline File	259
Chapter 15: Advanced .NET Probe Configuration	263
Automatically Discovering ASP.NET Applications.....	264
Discovering the Classes and Methods in an Application	265
Customizing the Instrumentation for ASP.NET Applications	267
Restarting IIS.....	270
Enabling the Mercury Diagnostics Probe for .NET	271
Enabling the Mercury Diagnostics Profiler for .NET.....	272
Configuring Latency Trimming and Throttling	272
Configuring Depth Trimming.....	276
Configuring the .NET Probe for Light-Weight Memory Diagnostics	278
Disabling Logging.....	280
Overriding the Default Probe Host Machine Name.....	280
Modifying the .NET Probe Configuration File	281

Chapter 16: Advanced J2EE Probe and Application Server Configuration	305
Advanced Configuration Directory.....	306
Configuring the Probe to Work with Mercury Products	307
Configuring the Probes for Multiple Application Server JVM Instances	312
Disabling the J2EE Diagnostics Profiler.....	322
Specifying Probe Properties as Java System Properties.....	322
Controlling Probe Logging.....	323
Setting the Probe Host Machine Name.....	325
Controlling Automatic Method Trimming on the Probe.....	326
Controlling Probe Throttling	329
Configuring a Probe for a Proxy Server.....	332
Configuring Reverse HTTP for a Probe in MMS.....	332
Configuring the Application Server for Deep Diagnostics Allocation Capture.....	334
Diagnostics Probe Administration Page.....	334
J2EE Probe Security and User Management.....	337
Using BEA JRockit Mission Control	341
Chapter 17: Probe Instrumentation	343
About Instrumentation and Capture Points Files	343
Understanding Points in the Capture Points File	345
Understanding Diagnostics Layers.....	352
Defining Points Using Code Snippets.....	356
Chapter 18: Using Solution Templates	359
About Solution Templates.....	359
Understanding Solution Template Data	360
Viewing Solution Template Information.....	361

PART VI: ADVANCED COMPONENT COMMUNICATION

Chapter 19: Enabling HTTPS Between Diagnostics Components	373
About Configuring HTTPS Communications	374
Enabling Incoming HTTPS Communication for Diagnostics Components	375
Enabling Outgoing HTTPS Communication from Diagnostics Components	379
Enabling HTTPS Communications for the Mercury Business Availability Center Server	382

Chapter 20: Configuring Diagnostics Components for HTTP Proxy	385
Enabling HTTP Proxy Communications for the Diagnostics Servers	386
Enabling HTTP Proxy Communications for the J2EE Probe	387
Enabling HTTP Proxy Communications for a .NET Probe	387
Chapter 21: Configuring Diagnostics Components to Work with a Firewall	389
Overview of Configuring Diagnostics for a Firewall	390
Collating Offline Analysis Files over a Firewall	393
Installing and Configuring the Mercury MI Listener	393
Configuring the Diagnostics Server in Mediator mode to Work with a Firewall	394
Configuring LoadRunner and Performance Center to Work with Diagnostics Firewalls	400

PART VII: CONFIGURING DIAGNOSTICS METRICS COLLECTORS

Chapter 22: Overview of Diagnostics Metrics Collectors	407
About Metrics Capture	407
Configuring Metric Collector Entries	408
Chapter 23: Configuring System Metrics Capture	413
About System Metrics	413
System Metrics Captured by Default	414
Configuring System Metric Collector	415
Capturing Custom System Metrics	416
Enabling z/OS System Metrics Capture	423
Chapter 24: Configuring JMX Metric Capture	425
About JMX Metrics	425
Configuring WebSphere for JMX Metric Collection	426
Configuring JMX Metric Collectors	428
Capturing Custom JMX Metrics	428

PART VIII: APPENDIXES

Appendix A: Diagnostics Server Administration Page	435
Accessing the Diagnostics Server Administration Page	435
Configuring Diagnostics Using the Diagnostics Server Configuration Pages	437

Appendix B: Product Security & User Privileges	443
About Product Security and User Privileges	444
Understanding User Privileges	445
Accessing Diagnostics Using Default User Names	446
Assigning User Privileges.....	447
Tracking User Administration Activity	451
Changing Your Password	452
Appendix C: Using the System Health Monitor	455
Introducing the System Health Monitor	456
Accessing the System Health Monitor	456
Understanding the System Health Monitor.....	459
Viewing Component Status and Host Configuration Tooltips	464
Viewing Detailed Component Information	466
Viewing Troubleshooting Tips	470
Viewing Log Information for the Whole System.....	470
Customizing the System Health Monitor Display	471
Filtering the System Health Monitor Component Map by Customer.....	473
Controlling the Refresh Rate of the System Health Monitor	474
Creating and Using System Health Monitor Snapshots	475
Appendix D: Diagnostics Data Management	477
About Diagnostics Data.....	478
Custom View Data.....	478
Performance History Data	480
Data Storage Size & Data Retention	483
Pre-Installation Data Management Considerations.....	485
Backing Up Diagnostics Data	486
Handling Diagnostics Data when Upgrading Diagnostics	489
Appendix E: Diagnostics Component Configuration and Communication Diagrams	495
Diagnostics Component Configuration and Communication	496
Appendix F: Diagnostics Upgrade Paths for Mercury Products	497
Overview of Product Upgrade Paths	497
General Upgrade Recommendations	498
Product Upgrade Paths.....	499

Appendix G: Troubleshooting Mercury Diagnostics	503
Component Installation Interrupted on a Solaris Machine	503
J2EE Probe Fails to Operate Properly.....	504
J2EE Probe Throttling Seems Excessive	504
Error During WAS Startup with Mercury Diagnostics Profiler for J2EE.....	505
Appendix H: Using UNIX Commands	507
Index.....	509

Table of Contents

Welcome to This Guide

Welcome to the *Mercury Diagnostics Installation and Configuration Guide*.

This guide describes how to install and configure the Mercury Diagnostics components. This guide also explains how to set up other Mercury products for integration with Mercury Diagnostics.

Mercury Diagnostics 4.1 enables you to monitor and diagnose the performance of your enterprise applications. The product gathers the metrics that you want to monitor and presents them in graphs and reports that enable you to see how your application is processing the data and how it is utilizing your systems resources. When you discover an anomaly or an issue that you want more information about, Mercury Diagnostics enables you to drill down to the graphs to find the transactions and methods that are contributing to the performance results you are seeing.

How This Guide Is Organized

This guide contains the following parts:

Part I Introduction

Provides the information and instructions to plan and prepare for the installation and configuration of the Diagnostics components.

Part II Installation of the Mercury Diagnostics Server

Describes how to install and configure the Mercury Diagnostics Server.

Part III Installation and Configuration of the Mercury Diagnostics Probes

Describes the processes for installing and configuring the Mercury Diagnostics Probes.

Part IV Setting Up Integration with Other Mercury Products

Describes the necessary steps to set up LoadRunner, Performance Center, and Mercury Business Availability Center for integration with Mercury Diagnostics.

Part V Advanced Component Configuration

Describes advanced configuration of the Mercury Diagnostics Server, and the Mercury .NET and J2EE Probes.

Part VI Advanced Component Communication

Describes how to set up your Mercury Diagnostics platform using different communication channels.

Part VII Configuring Diagnostics Metrics Collectors

Describes metrics capture and how to configure the metric collectors.

Part VIII Appendixes

Describes more detailed and advanced information about selected topics.

Who Should Read This Guide

This guide is intended for the people responsible for installing and configuring Mercury Diagnostics.

This document assumes that you are moderately knowledgeable about enterprise application development and highly skilled in enterprise system and database administration.

Mercury Diagnostics Online Documentation

Your Mercury Diagnostics application comes with the following documentation:

- ▶ **Mercury Diagnostics User's Guide:** Explains how to use Mercury Diagnostics to analyze the performance of your enterprise applications. You access this guide online from the **Help** button in Diagnostics or from the help menu in the integrated Mercury product. You access the PDF version of this guide from the Windows Start menu (**Start > Programs > Mercury Diagnostics Server > User's Guide**), from the **Docs** directory on the Mercury Diagnostics installation CD, or from the Diagnostics Server installation directory.
- ▶ **Mercury Diagnostics Installation and Configuration Guide:** Explains how to install and configure the Mercury Diagnostics components. You access the PDF version of this guide from the Windows Start menu (**Start > Programs > Mercury Diagnostics Server > Installation Guide**), from the **Docs** directory on the Mercury Diagnostics installation CD, or from the Diagnostics Server installation directory.
- ▶ **Readme:** Provides last-minute technical and troubleshooting information about Mercury Diagnostics. The file is located in the Mercury Diagnostics installation CD root directory.
- ▶ **Mercury Diagnostics Probe for J2EE Installation Quick Start:** Provides the basic instructions for installing the Mercury Diagnostics Probe for J2EE and is available in the **Docs** directory on the Probe installation CD or from the Windows Start menu (**Start > Programs > Mercury Diagnostics J2EE Probe > QuickStart**).
- ▶ **Mercury Diagnostics Profiler for J2EE Installation and User's Guide:** Describes how to install, configure and use the Mercury Diagnostics Profiler for J2EE. You access this guide online by clicking **Help** in the Mercury Diagnostics Profiler for J2EE. You access the PDF version of this guide from the **Docs** directory on the Probe installation CD.

- ▶ **Mercury Diagnostics Profiler for .NET Installation and User's Guide:**
Describes how to install, configure and use the Mercury Diagnostics Profiler for .NET. You access this guide online by clicking **Help** in the Mercury Diagnostics Profiler for .NET. You access the PDF version of this guide from the **Docs** directory on the Probe installation CD.

Note: The information in the Mercury Diagnostics Profiler guides is also included in the **Mercury Diagnostics Installation and User's Guide**.

- ▶ **J2EE Technical Practice Documents:** The following J2EE Technical Practice Documents are available in the **Docs** directory on the Mercury Diagnostics installation CD.
 - ▶ Advanced Instrumentation for Mercury Diagnostics for J2EE
 - ▶ Performance Impact Analysis of Mercury Diagnostics Probe for J2EE

Additional Online Resources

Customer Support Web Site uses your default Web browser to open the Mercury Customer Support Web site. This site enables you to browse the Mercury Support Knowledge Base and add your own articles. You can also 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.mercury.com>.

Mercury Home Page uses your default Web browser to access Mercury's Web site. This site provides you with the most up-to-date information on Mercury and its products. This includes new software releases, seminars and trade shows, customer support, educational services, and more. The URL for this Web site is <http://www.mercury.com>.

Documentation Updates

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

To download updated documentation:

- 1** In the Customer Support Web site, click the **Documentation** link.
- 2** Under **Please Select Product**, select either **Business Availability Center**, **LoadRunner** or **Performance Center**.

Note that if one of these items does not appear in the list, you must add it to your customer profile. Click **My Account** to update your profile.

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

Typographical Conventions

This guide uses the following typographical conventions:

UI Elements	This style indicates the names of interface elements on which you perform actions, file names or paths, and other items that require emphasis. For example, “Click the Save button.”
<i>Arguments</i>	This style indicates method, property, or function arguments and book titles. For example, “Refer to the <i>Mercury User’s Guide</i> .”
<Replace Value>	Angle brackets enclose a part of a file path or URL address that should be replaced with an actual value. For example, <MyProduct installation folder>\bin .
Example	This style is used for examples and text that is to be typed literally. For example, “Type Hello in the edit box.”
CTRL+C	This style indicates keyboard keys. For example, “Press ENTER.”
Function_Name	This style indicates method or function names. For example, “The wait_window statement has the following parameters:”
[]	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. In a programming example, an ellipsis is used to indicate lines of a program that were intentionally omitted.
	A vertical bar indicates that one of the options separated by the bar should be selected.

Part I

Introduction

1

Preparing to Install Mercury Diagnostics

This chapter provides you with the information and instructions that will help you to plan and prepare for the installation and configuration of the Diagnostics components.

This chapter describes:	On page:
About Mercury Diagnostics 4.1	4
Mercury Diagnostics Components	5
Licensing Mercury Diagnostics	7
Supported Application Servers	7
Mercury Diagnostics Data Flow	8
Recommended Deployment Configuration	9
Host System Requirements for the Diagnostics Components	14
Upgrading from Earlier Versions of Diagnostics	18
Other Pre-installation Considerations	19
Recommended Order of Installation	20
Planning the Installation	21

About Mercury Diagnostics 4.1

Mercury Diagnostics is a composite application triage and diagnostics solution that is designed to help you improve the performance of your enterprise applications on J2EE, .NET, Oracle and SAP platforms throughout the application lifecycle. Mercury Diagnostics has been integrated with Mercury's Application Delivery and Application Monitoring solutions to provide you with the insight and information that you need to build, develop, test, and monitor applications that perform efficiently and effectively.

Mercury Diagnostics for J2EE provides you with the capability to monitor the performance of your applications that run on most of the J2EE-compliant application servers. You install the Mercury Diagnostics Probe for J2EE on the application server instances where the application that you want to monitor runs. The J2EE Probe collects performance metrics on the servlets, JSPs, EJBs, JNDI, JDBC, JMS, and Struts method calls that are performed by your application. Custom business logic methods can be captured by creating a custom instrumentation specification file for your application.

Mercury Diagnostics for .NET provides you with the capability to monitor the performance of your applications that run on the Microsoft .NET Framework. The Mercury Diagnostics Probe for .NET uses runtime instrumentation to capture method latency information from specified applications. By default, the .NET Probe captures methods from ASP, ADO, and MSMQ. Custom business logic methods can be captured by creating a custom instrumentation specification file for your application.

Mercury Diagnostics for ERP/CRMS provides you with the capability to monitor the performance of Oracle 10g Database and SAP R/3 systems .

Mercury Diagnostics for SAP NetWeaver provides an end-to-end analysis of SAP Enterprise Portal transactions starting from portal pages and iViews, while maintaining the business context of services provided through the portal. Diagnostics also supports tracing and diagnosis of Java WebDynPro applications hosted on the Portal.

The Probes send the performance metrics to the Diagnostics Server where the events are filtered and aggregated. The Diagnostics Server displays the processed metrics using graphs and tables of the customizable user interface. You can then drill down further into the performance metrics to find the layers, server requests, and methods with the worst performance to understand the underlying cause of the observed behavior.

Mercury Diagnostics Components

Mercury Diagnostics consists of the following components:

- Mercury Diagnostics Probes
- Mercury Diagnostics Collector
- Diagnostics Servers

Mercury Diagnostics Probes

The Mercury Diagnostics Probes are installed on the machine that hosts your application. The Probes are responsible for capturing events from your application and sending the performance metrics to a Diagnostics Server. The Probe captures events such as method invocations, collection sites, the beginning and end of business transactions and server transactions.

The Mercury Diagnostics Probe works with many of Mercury's Diagnostics products such as LoadRunner, Mercury Business Availability Center, Performance Center, and Deep Diagnostics, to cover different parts of the lifecycle.

For a diagram showing the data flow, see the Mercury Diagnostics Data Flow diagram on page 8.

To enable the Probe to monitor your application, you must first configure the Probe and the application environment. For more information, see Chapter 4, "Introducing Installation and Configuration of Diagnostics Probes."

Mercury Diagnostics Collector

The Diagnostics Collector gathers data from external active systems. The Collector can be configured to collect performance data from two types of active systems:

- SAP R/3 system
- Oracle 10g Database

During the installation of the Collector, you can choose to monitor either one or both of these active systems. After the installation, you define instances of Oracle 10g and SAP R/3 systems to be monitored. Each one of these instances is represented as a Probe (either Oracle Probe or SAP R/3 Probe) belonging to a Probe group in the Mercury Diagnostics UI.

Diagnostics Server

The Mercury Diagnostics Server is responsible for working with the Probes and the other Mercury products to capture, process, and present the performance metrics for your application. The Diagnostics Server pre-processes and aggregates the raw data that it receives from the Probes, and formats the information so that it can be displayed in the views of the user interface. Since the Diagnostics Server does the processing for the Probes, the Probes have a negligible effect on the performance of the applications that they monitor.

In a distributed environment, where you have more than one Diagnostics Server, you must designate one of them to have the role of Diagnostics Server in Commander mode. The other Diagnostics Servers in the deployment must be installed as Diagnostics Servers in Mediator mode.

The Diagnostics Server in Commander mode is responsible for the command and control functions between the various Diagnostics components and the components of the other Mercury products with which Diagnostics is working. The Diagnostics Server in Commander mode keeps track of the location and status of the other Diagnostics components, and is the communication hub between the other components.

The Diagnostics Server in Commander mode is also responsible for displaying the performance information for the monitored applications in the charts and graphs of the Diagnostics views. If you are using Diagnostics with other Mercury products you may also be able to access the Diagnostics views from the user interface of the other products. For more information about accessing Diagnostics, see the *Mercury Diagnostics User's Guide*.

Licensing Mercury Diagnostics

To be able to see the performance metrics for your applications in the Diagnostics views, you must obtain a valid license for the Diagnostics Server in Commander mode. The license is a node-locking license that unlocks the Diagnostics views that are displayed by the Diagnostics Server in Commander mode and removes the load restriction from the Profiler views that are accessed from the Diagnostics Probes. For more information on obtaining a license and other licensing issues, see Chapter 3, “Licensing Mercury Diagnostics.”

Supported Application Servers

Mercury Diagnostics supports the monitoring of most J2EE-based application servers, including: WebLogic, WebSphere, Oracle9iAS, Sun Java System, JBoss, and more.

Note: Mercury Diagnostics monitors only those methods that are invoked by the application server, and are instrumented by the probe. For more information about instrumenting additional methods, see Chapter 16, “Advanced J2EE Probe and Application Server Configuration.”

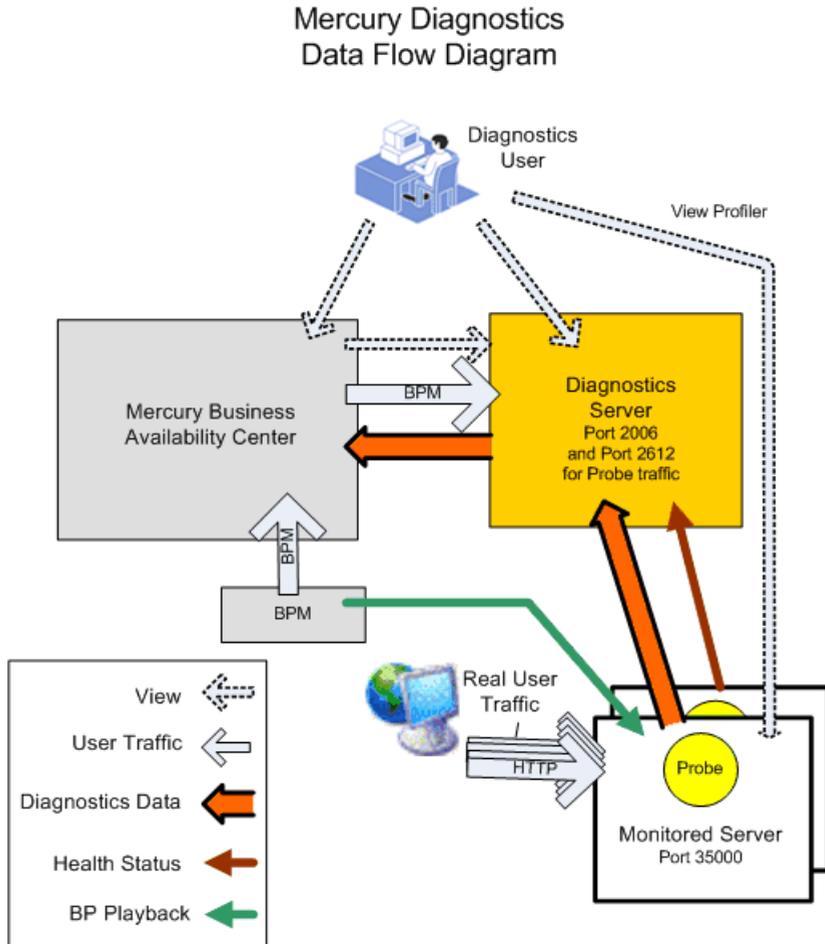
For .NET-based application servers, Mercury Diagnostics supports the Microsoft IIS .NET Framework.

Mercury Diagnostics Data Flow

The following diagram illustrates the Diagnostics data flow when Diagnostics is integrated with Mercury Business Availability Center.

Virtual-user (BP) transactions are sent from the Mercury Business Availability Center server to the Mercury Diagnostics Server.

For real-time reports, real-user transactions are aggregated every 30 seconds, and sent to the Mercury Diagnostics Server.



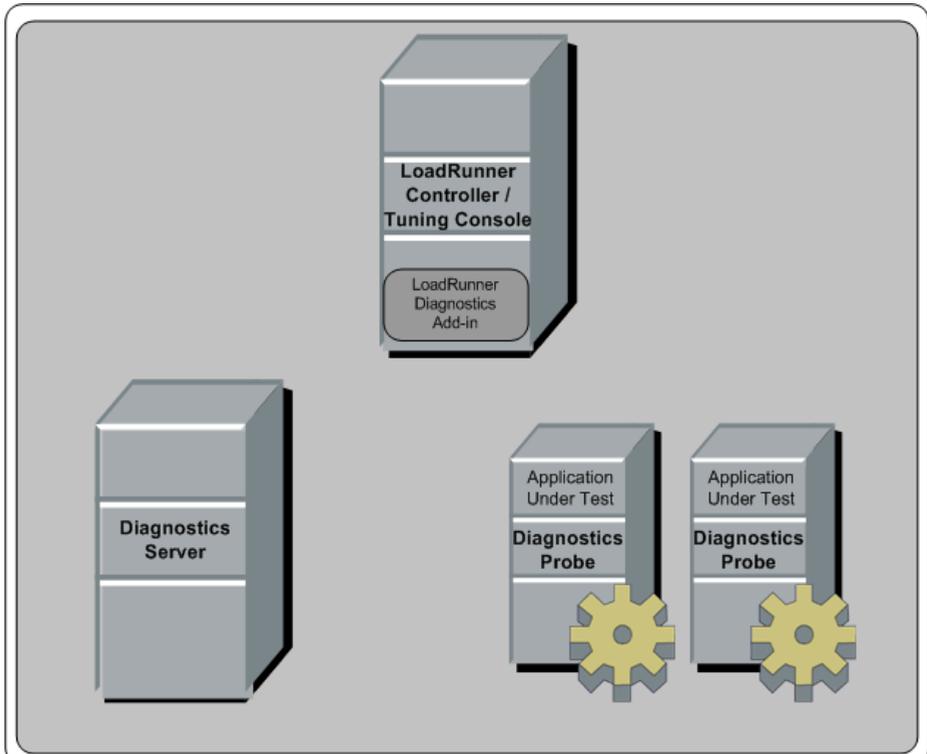
Recommended Deployment Configuration

This section provides you with a recommended deployment configuration for LoadRunner 8.1, Performance Center 8.1, or Mercury Business Availability Center 6.x and Mercury Diagnostics. These diagrams have been provided as a reference for you as you review the order of installation. More detailed instructions about the installation and configuration of the Diagnostics components are provided in the chapters of this guide that follow.

Note: A component communication diagram that shows the communication paths between the Diagnostics components and the ports that are used can be found in Appendix E, “Diagnostics Component Configuration and Communication Diagrams.”

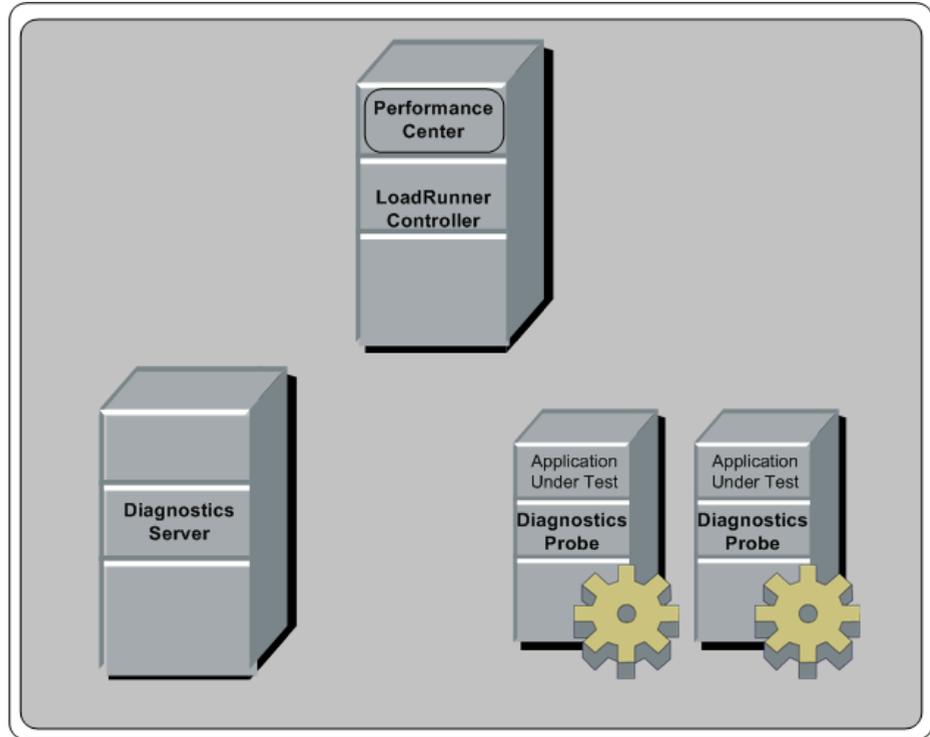
Deployment with LoadRunner

The following diagram shows a typical deployment configuration for 10 to 15 VMs, with a moderate load using a total of three host machines: one for the LoadRunner Controller and one for the Diagnostics Server. The deployment that will work best for you is dependent upon the number of VMs that you are monitoring and the size of the load that you are running.



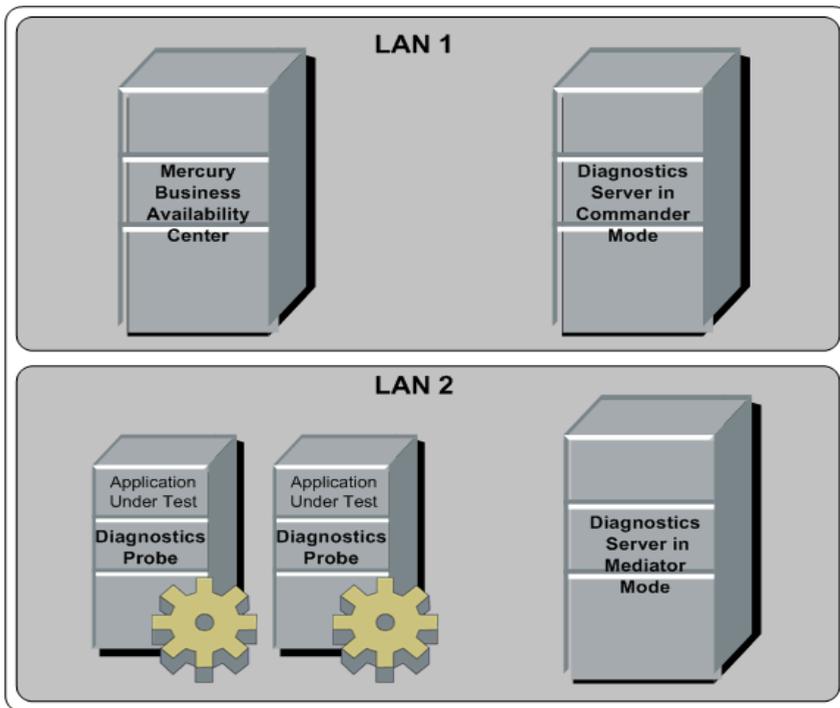
Deployment with Performance Center

The following diagram shows a suggested architecture for deploying Mercury Diagnostics with Performance Center:

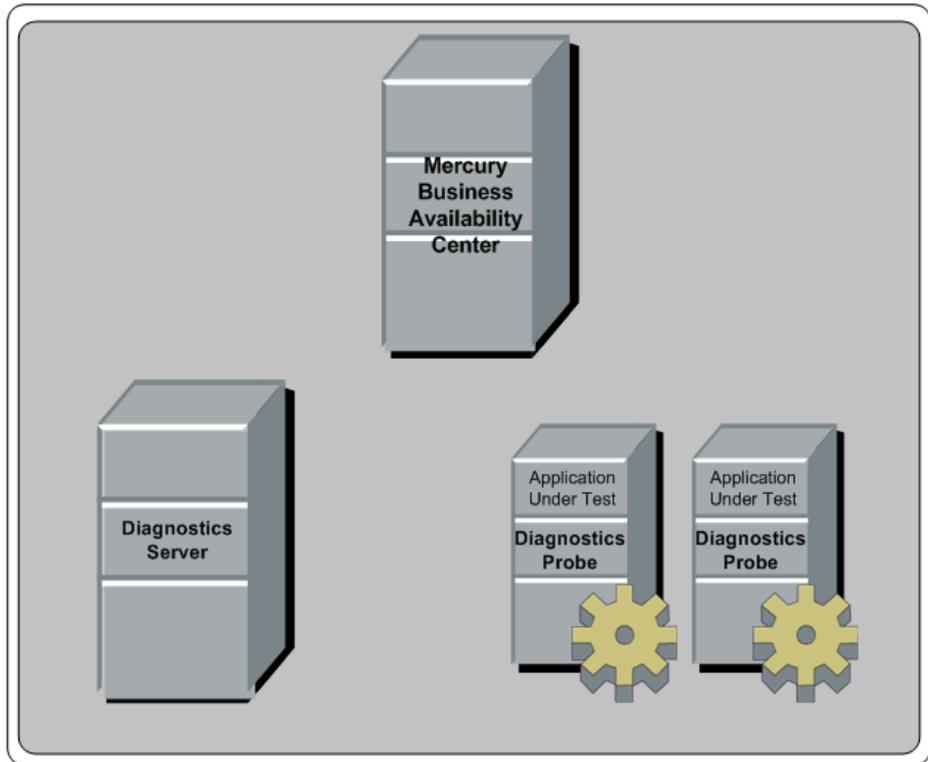


Deployment for Mercury Business Availability Center

The following diagram shows a typical deployment for a distributed environment where one of the Diagnostics Servers has been designated as the Diagnostics Server in Commander mode. The deployment that will work best for you, depends the number of VMs that you are monitoring, their workload, and the custom instrumentation that you defined. An alternate deployment example for a smaller deployment has been provided in the next diagram.



The following diagram shows an optional deployment configuration that you may want to consider if you have a smaller deployment where you are monitoring fewer VMs. In this example, two hosts are used—one for the single Diagnostics Server and one for the Mercury Business Availability Center Server.



Host System Requirements for the Diagnostics Components

The following section describes the recommended system configurations for hosting the components of the Mercury Diagnostics. Refer to the deployment diagram in the previous section to understand the component hosts that are described in this section.

When you select the machines that will host the Diagnostics components, you must make sure that the system configuration of the machines supports the processing load and the number of applications that you will be monitoring.

Requirements for the Mercury Diagnostics Server Host

The system configuration requirements for the host of the Diagnostics Server depends upon the number of probes and Diagnostics Server in Mediator mode that are reporting to it.

Note: When a Diagnostics Server is designated as the Diagnostics Server in Commander mode, each Diagnostics Server in Mediator mode that reports to it counts as one half of a probe when determining the system configuration requirements.

For example if a Diagnostics Server in Commander mode had 20 Diagnostics Server in Mediator mode reporting to it, in addition to 5 probes, the table below indicates that it needs `-Xmx700M` as a parameter to its JVM.

The following table lists the required system configuration for the host of the Diagnostics Server:

Platform	Item	Up to 10 Probes	Up to 20 Probes	Up to 30 Probes
Windows	CPU	1x 1.0 GHz	1x 2.0 GHz	2x 2.4 GHz
Windows	Memory	768 MB	1 GB	2 GB
Solaris	CPU	1x Ultrasparc 2	2x Ultrasparc 2	2x Ultrasparc 3
Solaris	RAM	1 GB	1.5 GB	2 GB
All	Java Heap	350 M	700 M	1400 M
All	Disk	TBD	TBD	TBD

Note: For configuration considerations related to the Diagnostics performance data that is stored on the host for the Diagnostics Server in Commander mode, see “Pre-Installation Data Management Considerations” on page 485.

Requirements for the Mercury Diagnostics Probe for J2EE Host

The overhead that the Mercury Diagnostics Probe for J2EE (J2EE Probe) imposes on the system being monitored is extremely low. The following are the recommendations for free memory and disk space that will support the Probe's processing:

Platform:	All Platforms
Memory:	50 MB Additional RAM
Free Hard Disk Space:	200 MB Additional Space

Note: The additional memory must be allocated to the max heap for the JVM by adding `-Xmx???m` to the java settings in the application's start up script.

For information on setting the max heap for the J2EE Probe, see "Adjusting the Heap Size for the J2EE Probe in the Application Startup Script" on page 208.

Requirements for the Host Running the J2EE Diagnostics Profiler User Interface

The user interface for the J2EE Diagnostics Profiler is presented in a web browser using a JAVA applet that requires JRE 1.4 or above to be installed on the machine that displays UI. This machine must be able to access the J2EE Diagnostics Profiler URL: `http://<probe_host>:<probeport>/profiler`. The Probes are assigned to the first available port beginning at 35000.

Requirements for the Mercury Diagnostics Probe for .NET Host

The overhead that the Mercury Diagnostics Probe for .NET (.NET Probe) imposes on the system being monitored is extremely low. The following are the recommendations for free memory and disk space that will support the Probe's processing:

Platform:	All Platforms
Memory:	60 MB Additional RAM
Free Hard Disk Space:	200 MB Additional Space

Requirements for the Host for the .NET Diagnostics Profiler User Interface

The user interface for the .NET Diagnostics Profiler is presented using DHTML/XML/XSLT/JavaScript technology that requires IE6. The machine that is to be used to present the UI must be able to access the .NET Diagnostics Profiler URL: <http://<probehost>:<probeport>/profiler>. The Probes are assigned to the first available port beginning at 35000.

Requirements for the Diagnostics Collector Host

The Collector can be installed on any machine that can interact with the host machines of the SAP or Oracle application from which it is collecting data.

The following table lists the required system configuration for the host of the Diagnostics Collector:

Platform	Item	Up to 50 Probes
Windows	CPU	1x 1.0 GHz
Windows	Memory	1 GB
Solaris	CPU	1x Ultrasparc 2
Solaris	RAM	1 GB
All	Java Heap	350 M
All	Disk	TBD

Upgrading from Earlier Versions of Diagnostics

If you are installing Mercury Diagnostics into an environment where a previous version of the product was installed, or in an environment where other Mercury products need to be upgraded so that the features of Diagnostics can be accessed, follow the instructions in Appendix F, “Diagnostics Upgrade Paths for Mercury Products.” These instructions guide you to the appropriate instructions for upgrading your current Mercury products and the Diagnostics components.

Other Pre-installation Considerations

Note: Before you install any of the Diagnostics components on a Windows machine, make sure that the **Start > Settings > Control Panel > Administrative Tools > Services** window is not open.

LoadRunner and Performance Center Host Machines

- ▶ If LoadRunner 8.1 is already installed, ensure that the Controller/Console and main Mercury LoadRunner window are closed before you install the LoadRunner Diagnostics Add-in.
- ▶ The LoadRunner Diagnostics Add-in is not required for Performance Center.
- ▶ The time and time-zone settings of the host machines for the Diagnostics components must be consistent. You will encounter time-difference problems if the time is not properly set.

Mercury Diagnostics Server

- ▶ The performance metrics for Mercury Diagnostics cannot be displayed until the Diagnostics Server in Commander mode has been licensed with a valid license. For more information on obtaining a license and other licensing issues, see Chapter 3, “Licensing Mercury Diagnostics.”

Mercury Diagnostics Probe for J2EE

- ▶ The J2EE Probe must be installed on the same machine as the Java application under test.
- ▶ The Mercury Diagnostics Profiler for J2EE will operate in an unlicensed mode with load restrictions until it is able to connect to a Diagnostics Server in Commander mode that has been properly licensed. For more information on obtaining a license and other licensing issues, see Chapter 3, “Licensing Mercury Diagnostics.”

Mercury Diagnostics Probe for .NET

- ▶ The .NET Probe must be installed on the same machine as the .NET application under test.
- ▶ The Mercury Diagnostics Profiler for .NET will operate in an unlicensed mode with load restrictions until it is able to connect to a Diagnostics Server in Commander mode that has been properly licensed. For more information on obtaining a license and other licensing issues, see Chapter 3, “Licensing Mercury Diagnostics.”

Recommended Order of Installation

Careful planning and preparation for installing the components of Mercury Diagnostics can help you to avoid complications and errors, and complete the installation and configuration steps quickly.

Before you start, review the following information to get an overview of the entire installation and configuration process.

Note: The order of the installation presented here is the recommended order of installation for the products and components. Deviation from this recommended order of installation could increase the complexity of the installation process and produce unpredictable results.

1 Check the system requirements and installation considerations.

See “Host System Requirements for the Diagnostics Components” on page 14.

2 Install the Mercury Diagnostics Server.

See Chapter 2, “Installing the Mercury Diagnostics Server.”

3 Install and configure the Mercury Diagnostics Probe(s).

- ▶ For a J2EE environment, see Chapter 6, “Installing the Mercury Diagnostics Probe for J2EE.”
- ▶ For a .NET environment, see Chapter 5, “Installing the Mercury Diagnostics Probe for .NET.”
- ▶ For Oracle and SAP R3 environments, see Chapter 4, “Installing the Diagnostics Collector.”

4 For J2EE Probes, configure the application server to work with the probes.

For more information, see Chapter 9, “Configuring the Application Servers.”

5 If Mercury Diagnostics is integrated with Mercury LoadRunner, Performance Center or Business Availability Center, it must be set up to use Mercury Diagnostics.

- ▶ For Mercury Business Availability Center, see Chapter 10, “Setting Up Mercury Business Availability Center 6.x to Use Diagnostics.”
- ▶ For LoadRunner integration, install the LoadRunner Diagnostics Add-in (see Chapter 11, “Installing LoadRunner 8.1 and the LoadRunner Diagnostics AddIn”) and set up LoadRunner to use Mercury Diagnostics (see Chapter 12, “Setting Up LoadRunner 8.1 to use Mercury Diagnostics”).
- ▶ For Performance Center, see Chapter 13, “Setting Up Performance Center 8.1 to Use Mercury Diagnostics.”

Planning the Installation

Before you start installing the Diagnostics components, you should carefully plan the configuration of the Diagnostics components and the machines that will host them. You should also consider the location of the component hosts within your network topography.

The tables in the following sections have been provided to help you gather the information that you must provide during the installation of the Diagnostics components.

Note: When you are installing Mercury Business Availability Center with Diagnostics, when entering the names for the hosts of the Diagnostics components it is strongly recommended that you use fully qualified host names, that is, the machine name and the domain name.

Diagnostics Server

Information Required	Where to find it	Value
For a Diagnostics Server in Commander mode, the location of the Mercury Diagnostics license that has been generated for the machine that will host the server.	Contact your Mercury support person to request a license and place it in a folder where it can be accessed from the Diagnostics Server installer.	
For a Diagnostics Server in Mediator mode, the URL for the Diagnostics Server in Commander mode.	After the Diagnostics Server in Commander mode has been installed, available from the System Health Monitor. (See Appendix C, “Using the System Health Monitor.”)	
Will the Diagnostics Server be used in a Mercury Managed Services environment or in the Mercury Business Availability Center environment?		

J2EE Probe► **Mercury Product, Diagnostics Server, and Deep Diagnostics Information**

Information Required	Where to find it	Value
Name of Mercury product(s) that will use the Probe	Choose according to product license. <ul style="list-style-type: none"> • Performance Center 8.1 / LoadRunner 8.1 • Mercury Business Availability Center 6.x • Deep Diagnostics 	
Diagnostics Server in Commander mode URL	System Health Monitor (See Appendix C, "Using the System Health Monitor.")	
Diagnostics Server Event Host	System Health Monitor (See Appendix C, "Using the System Health Monitor.")	
Diagnostics Server Event Port	System Health Monitor (See Appendix C, "Using the System Health Monitor.")	Default value: 2612
Deep Diagnostics install directory (Only if you will be using Deep Diagnostics.)	The location where the Deep Diagnostics Server was installed.	
Deep Diagnostics Application Definition Name (Only if you will be using Deep Diagnostics.)	See the Deep Diagnostics User Documentation for the location of this information.	

► Probe and Application Server Information

Information Required	Where to find it	Value
<p>Probe name</p>	<p>A <i>unique</i> string; Created by user. Note: The name of the Probe should indicate the probe type, to help you distinguish between the different types of probes</p>	<p>For example: J2EEProbe1</p>
<p>Probe Group</p>	<p>Used to relate Probes so that their metrics can be reported as a logical group. This is user defined at the time that the Probe is installed.</p>	<p>Default value: DefaultProbeGroup</p>
<p>Type of application server that the J2EE Probe will be monitoring</p>	<p>The host system administrator.</p>	
<p>Application Server configuration properties</p>	<p>The host system administrator. The details will vary according to the application server that you are using.</p>	
<p>Location of the JRE executable</p>	<p>The host system administrator. This depends upon the type of application server you are configuring. See “Running the JRE Instrumenter” on page 165.</p>	

.NET Probe► **Diagnostics Server Information**

Information Required	Where to find it	Value
URL of Diagnostics Server in Commander mode	System Health Monitor URL that the Probe will use to register with the Diagnostics Server. This is not required for using the.NET Diagnostics Profiler in a standalone mode.	
Diagnostics Server Event Host	System Health Monitor This is not required for using the.NET Diagnostics Profiler in a standalone mode.	
Diagnostics Server Event Port	System Health Monitor Default value: 2612 This is not required for using the.NET Diagnostics Profiler in a standalone mode.	

► Probe and Port Information

Information Required	Where to find it	Value
Probe Group	Used to relate probes so that their metrics can be reported as a group. This is user defined at the time that the Probe is installed. Default value: DefaultProbeGroup	
Web Port Min	System Administrator. The lowest port number in a range of ports on the Probe host that can be assigned to the Probe. Default value: 35000	
Web Port Max	System Administrator The highest port number in a range of ports on the Probe host that can be assigned to the Probe. Default value: 35100	

Part II

Installation of the Mercury Diagnostics Server

2

Installing the Mercury Diagnostics Server

This chapter describes how to install the Mercury Diagnostics Server on Windows and UNIX machines.

This chapter describes:	On page:
Installing the Diagnostics Server on a Windows Machine	30
Installing the Diagnostics Server on a UNIX Machine	41
Starting and Stopping the Diagnostics Server	52
Verifying the Diagnostics Server Installation	54
Licensing the Diagnostics Server in Commander Mode	55
Configuring the Diagnostics Server	55
Determining the Version of the Diagnostics Server	56
Uninstalling the Diagnostics Server	57
Troubleshooting Diagnostics Server Issues	58

Installing the Diagnostics Server on a Windows Machine

The following steps provide detailed instructions for installing the Diagnostics Server on a Windows machine.

To install the Diagnostics Server on a Windows machine:

1 Launch the installation program for the Diagnostics Server.

Insert the Mercury Diagnostics for Windows installation CD into the CD-ROM drive of the Diagnostics Server host machine, and double-click **setup.exe** in the root folder of the CD.

Note: To launch the installer from a different location, locate **\Diagnostics_Server\DiagnosticsServerSetupWin.exe** on the CD, copy it to the new location, and then run it.

The installer displays the Mercury Diagnostics main installation menu.

Click **Mercury Diagnostics Server** to launch the installer.

2 Accept the software license agreement.

A dialog box opens displaying the license agreement.

Read the agreement and select **I accept the terms of the license agreement**.

Click **Next** to continue.

3 Specify the location where the Diagnostics Server is to be installed.

In the **Installation Directory Name** box, type the name of the directory where you want to install the Diagnostics Server. Alternatively, accept the default directory, or click **Browse** to navigate to another directory.

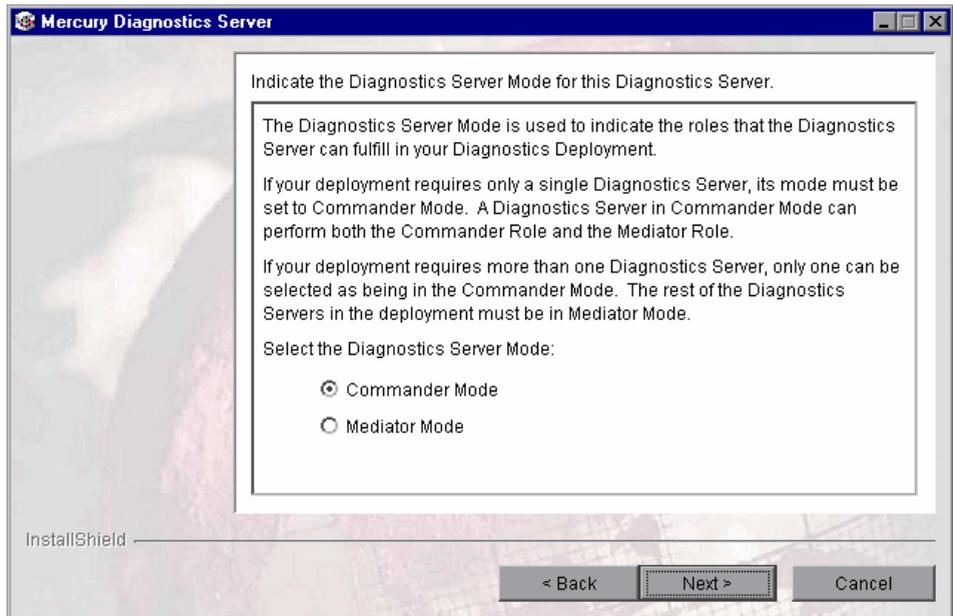
Click **Next** to continue.

Note: If you specified a directory that contains an existing Diagnostics Server installation, the installation process backs up any existing property settings in a directory called **etc.old**.

4 Indicate the mode of the Diagnostics Server that you are installing.

The deployment that you are setting up may consist of one or many Diagnostics Servers. If there is only one Diagnostics server in your deployment, it is configured in Commander mode and can perform both Commander and Mediator roles. If there is more than one Diagnostics Server in a deployment, one must be configured in Commander mode, and all the rest in Mediator mode.

- ▶ If this is the only Diagnostics Server in your deployment, select **Commander Mode**.
- ▶ If there is more than one Diagnostics Server in your deployment, and the one that you are currently installing is to be configured in Commander mode, then select **Commander Mode**. Otherwise select **Mediator Mode**.



Click **Next** to continue.

Next Step: At this stage, the installation differs depending on whether you are installing the Diagnostics Server in Commander mode or in Mediator mode.

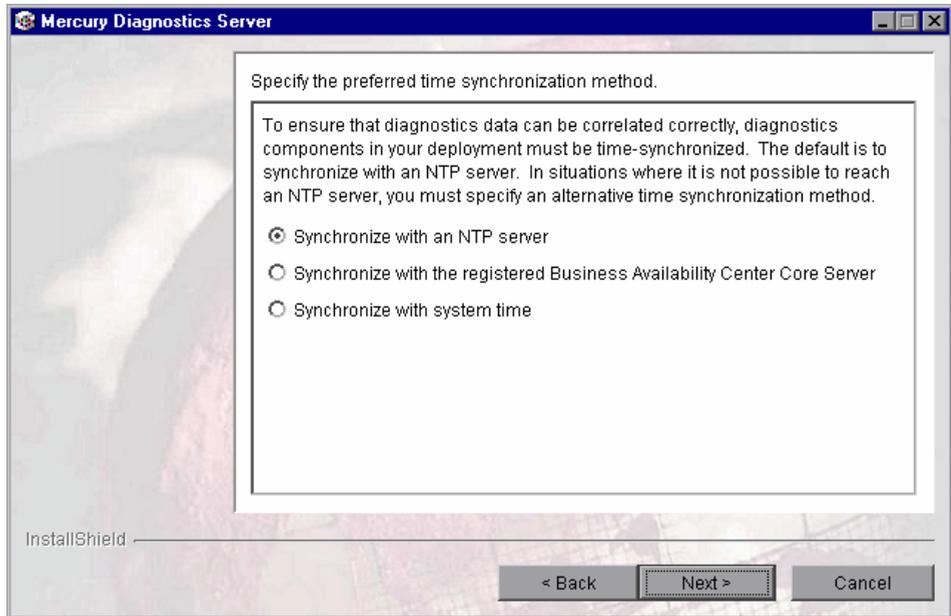
- ▶ To install the Diagnostics Server in Commander mode, continue with “Installing the Diagnostics Server in Commander Mode,” below.
 - ▶ To install the Diagnostics Server in Mediator mode, continue with “Installing the Diagnostics Server in Mediator Mode” on page 37.
-

Installing the Diagnostics Server in Commander Mode

If you are installing the Diagnostics Server in Commander mode, continue n as follows:

1 Select a time synchronization method.

In order for diagnostics data to be correlated properly, all the components in the Diagnostics deployment must be time-synchronized.



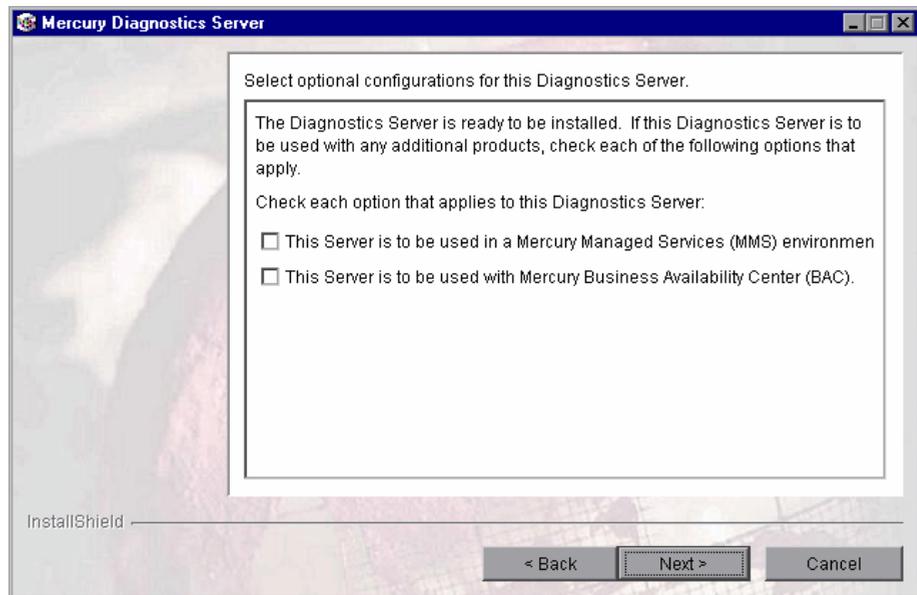
Select one of the following time synchronization methods:

- ▶ **Synchronize with an NTP server.** This option applies only if the Diagnostics Server can access an NTP Server outside the firewall. This is the default method.
- ▶ **Synchronize with the registered Mercury Business Availability Center Core Server.** If the Diagnostics Server is to work in a Mercury Business Availability Center environment, select this option to synchronize with the Mercury Business Availability Center core server.
- ▶ **Synchronize with system time.** Select this option if the Diagnostics Server is to work in an environment other than Mercury Business Availability Center and there is no access to an NTP server.

Click **Next** to continue.

2 Indicate whether the Diagnostics Server will be used in a Mercury Managed Services environment or with Mercury Business Availability Center.

Select the options that apply to this Diagnostics Server.



Click **Next** to continue.

Next Step: If you selected only the **Mercury Managed Services** option, skip to step 4.

3 Mercury Business Availability Center environment only: Provide the path to the Probe installers.

Note: You need to have the Mercury Diagnostics installation CD available for this step.

To be able to download the Diagnostics Probe installers from the Diagnostics Configuration page in Mercury Business Availability Center, you need to specify the path to the directory on the Mercury Diagnostics installation CD where these installers are located (**\Diagnostics_Probes**).

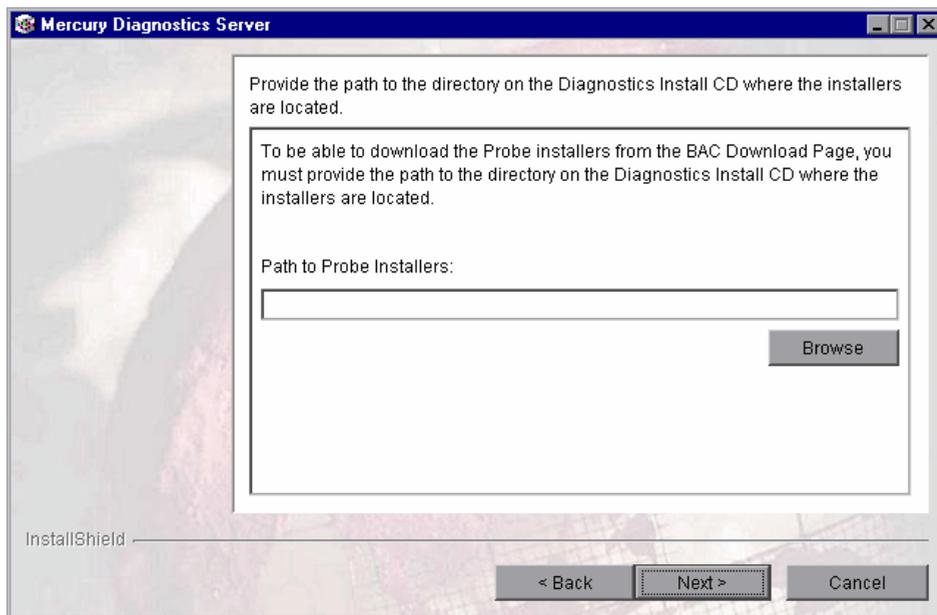
The Probe installers are copied to the Diagnostics Server installation directory, which Mercury Business Availability Center can access.

Note: You can choose to skip this step and always access the Probe installers directly from the product CD.

Alternatively, you can perform this step manually at a later stage, by copying the probe installers (**<Mercury Diagnostics installation CD>/Diagnostics_Probes**) to the Diagnostics Server installation directory (**<diag_server_install_dir>/html/opal/downloads**).

To be able to download the Diagnostics Collector installer from the Diagnostics Downloads page in Mercury Business Availability Center, locate the **\Diagnostics_Collector\CollectorSetupWin.exe** on the Diagnostics CD, and copy it to the following folder of the Diagnostics Server installation directory: **<diag_server_install_dir>/html/opal/downloads**. For more information, see Chapter 4, “Installing the Diagnostics Collector.”

In the **Path to Probe Installers** box, type the path, or click **Browse** to navigate to the **Diagnostics_Probes** directory on the installation CD.



Click **Next** to continue.

4 Review the pre-installation summary.

The installation settings that you selected are displayed in a read-only window. Review the information to make sure that you are satisfied.

Note: The estimated total size of the Diagnostics Server in Commander mode installation does not include the size of the Probe installers, if they were made available for Mercury Business Availability Center.

To select different installation settings, click **Back**.

To begin installing the Diagnostics Server, click **Next**.

5 Close the installation wizard.

When the installation has finished running, a message is displayed, confirming that the Diagnostics Server was successfully installed. Click **Finish** to exit the installation.

Note: On Windows machines, the Diagnostics Server attempts to start automatically. The Diagnostics Server will not start if any other applications are using the default Diagnostics Server ports. For instructions on starting the Diagnostics Server manually, see “Starting and Stopping the Diagnostics Server” on page 52.

6 Upload the Diagnostics license file.

In order to view performance metrics, you must upload a valid Mercury Diagnostics license file after you have installed the Diagnostics Server in Commander mode.

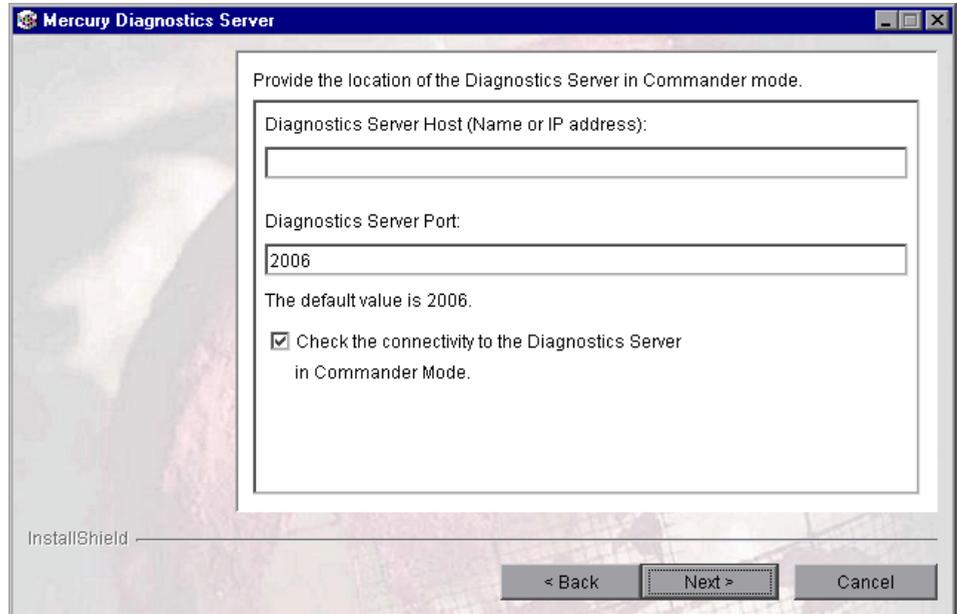
For instructions on requesting and uploading a valid license file, see Chapter 3, “Licensing Mercury Diagnostics.”

Installing the Diagnostics Server in Mediator Mode

If you are installing the Diagnostics Server in Mediator mode, continue the installation as follows:

1 Provide the location of the Diagnostics Server in Commander mode.

Provide the information that enables the Diagnostics Server in Mediator mode to connect to the Diagnostics Server in Commander mode.



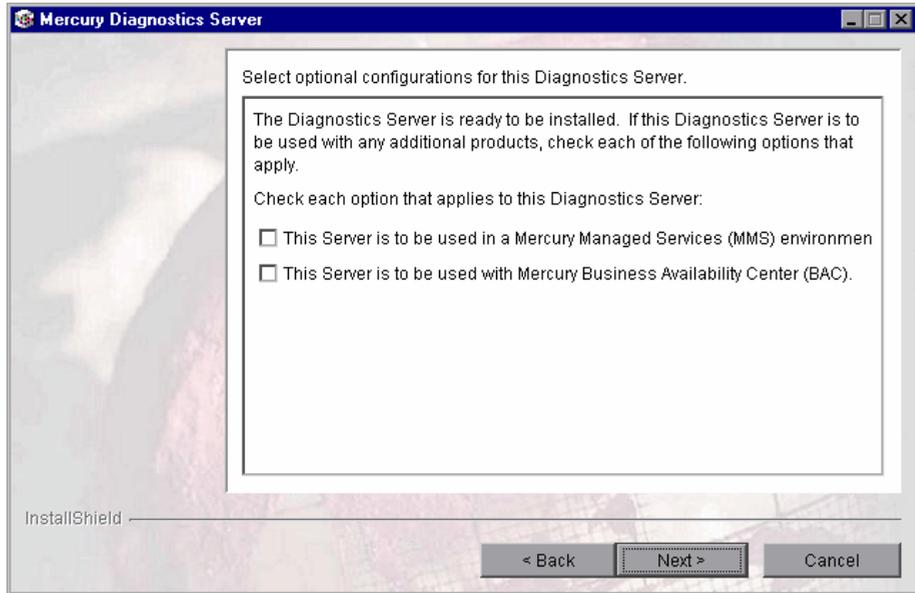
- a** In the **Diagnostics Server Host** box, type the host name or IP address of the host for the Diagnostics Server in Commander mode.
- b** In the **Diagnostics Server Port** box, type the port number of the Diagnostics Server in Commander mode. The default port for the Diagnostics Server in Commander mode is **2006**. If you have changed the port of the Diagnostics Server in Commander mode since it was installed, you should specify that port number instead of the default. For information on changing the Diagnostics Server port, see “Changing the Default Diagnostics Server Port” on page 250.
- c** If you want to make sure that the Diagnostics Server is running and accessible from the installation host, select **Check the connectivity to the Diagnostics Server in Commander Mode**. (This is selected by default.)

Click **Next** to continue.

Note: If you selected **Check the connectivity to the Diagnostics Server in Commander Mode** and connectivity problems were encountered, the installer provides you with the results of the connectivity check. If you do not want to address these problems at this stage, clear the **Check the connectivity to the Diagnostics Server in Commander Mode** check box, proceed with the installation, and address the problem later.

2 Indicate whether the Diagnostics Server will be used in a Mercury Managed Services environment or with Mercury Business Availability Center.

Select that options that apply to this Diagnostics Server.

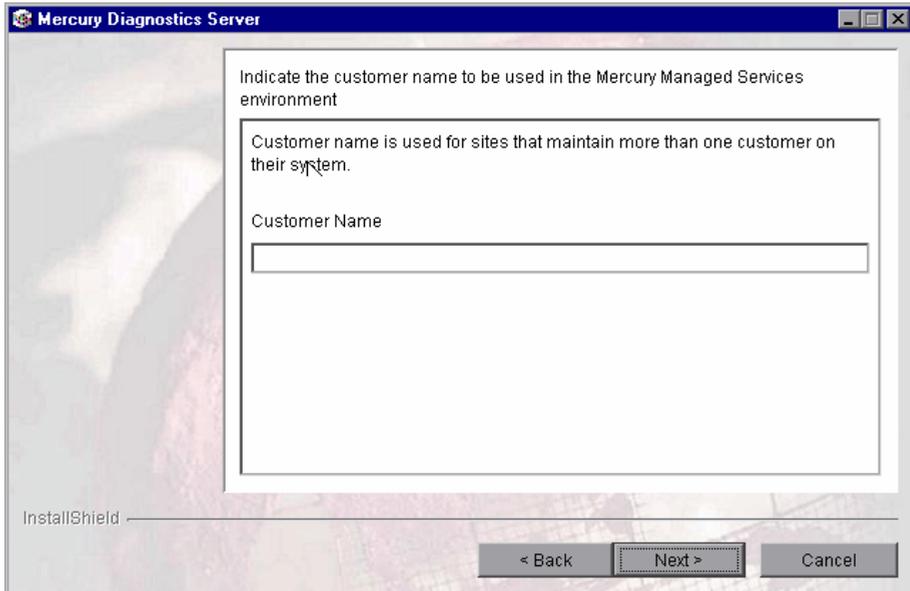


Click **Next** to continue.

Next Step: If you selected only the **Mercury Business Availability Center** option, skip to step 4.

- 3** If you indicated that the Diagnostics Server is to be used in a Mercury Managed Services environment, provide the customer name for the Mercury Managed Services environment.

Enter a unique name for the customer in the Mercury Managed Services environment.



Click **Next** to continue.

- 4** Review the pre-installation summary.

The installation settings that you selected are displayed in a read-only window. Review the information to make sure that you are satisfied.

To select different installation settings, click **Back**.

To begin installation, click **Next**.

5 Close the installation wizard.

When the installation has completed, review the post-installation summary information displayed to make sure that the installation completed successfully, and click **Finish**.

Note: On Windows machines, the Diagnostics Server attempts to start automatically. The Diagnostics Server will not start if any other applications are using the default Diagnostics Server ports. For instructions on starting the Diagnostics Server manually, see “Starting and Stopping the Diagnostics Server” on page 52.

Installing the Diagnostics Server on a UNIX Machine

This section provides instructions for installing the Diagnostics Server in most UNIX environments using either a graphical installation or a console mode installation.

Notes:

- ▶ The following instructions are for component installation on a Solaris machine. These same instructions should apply for the other certified UNIX platforms.
 - ▶ The following instructions assume an understanding of UNIX console screens and commands. For more information about UNIX screens and commands, see Appendix H, “Using UNIX Commands.”
-

To install the Diagnostics Server on a UNIX machine in console mode:

1 Launch the installation program for the Diagnostics Server.

- a** On the Mercury Diagnostics installation CD for UNIX, locate the **Diagnostics_Server** directory.
- b** Copy the installer, **DiagnosticsServerSetup<platform>.bin**, to the machine where the Diagnostics Server is to be installed.
- c** Change the mode of the installer file to make it executable.
- d** Run the installer.

- To run the installer in the graphical mode enter the following at the UNIX command prompt:

```
./DiagnosticsServerSetupSolaris.bin
```

The installer screens that are displayed in a graphical installation are similar to those documented for the Windows installer.

Continue with step 2 of “Installing the Diagnostics Server on a Windows Machine” on page 30.

- To run the installer in console mode enter the following at the UNIX command prompt:

```
./DiagnosticsServerSetupSolaris.bin -console
```

The installer continues as shown in the next step.

2 Accept the software license agreement.

The software license agreement is displayed.

Read the agreement. As you read, you can press **ENTER** to move to the next page of text, or type **q** to jump to the end of the license agreement.

Accept the terms of the agreement, and select **Next** to continue with the installation.

3 Specify the location where the Diagnostics Server is to be installed.

Note: If you specify a directory that contains an existing Diagnostics Server installation, the installation process backs up any existing property settings in a directory called **etc.old**.

At the **Directory Name** prompt, accept the default installation location shown in brackets, or enter the path to a different location.

Select **Next** to continue with the installation.

4 Indicate the mode of the Diagnostics Server that you are installing.

```

nm
-----
Indicate the Diagnostics Server Mode for this Diagnostics Server.

The Diagnostics Server Mode is used to indicate the roles that the Diagnostics
Server can fulfill in your Diagnostics Deployment.

If your deployment requires only a single Diagnostics Server, its mode must be
set to Commander Mode. A Diagnostics Server in Commander Mode can perform both
the Commander Role and the Mediator Role.

If your deployment requires more than one Diagnostics Server, only one can be
selected as being in the Commander Mode. The rest of the Diagnostics Servers in
the deployment must be in Mediator Mode.

Select the Diagnostics Server Mode:

[X] 1 - Commander Mode
[ ] 2 - Mediator Mode

To select an item enter its number, or 0 when you are finished: [0] 1

[X] 1 - Commander Mode
[ ] 2 - Mediator Mode

To select an item enter its number, or 0 when you are finished: [0]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1] _
  
```

The Diagnostics deployment that you are creating may consist of one or many Diagnostics Servers. If there is only one Diagnostics Server in your deployment, it is installed in Commander mode and can perform both Commander and Mediator roles. When there is more than one Diagnostics Server in a deployment, one is configured in Commander mode, and all the rest in Mediator mode.

- ▶ If this is the only Diagnostics Server in your deployment, select **Commander Mode**.
- ▶ If there is more than one Diagnostics Server in your deployment, and the one that you are currently installing is to be configured in Commander mode, then select **Commander Mode**. Otherwise select **Mediator Mode**.

Select **Next** to continue with the installation.

Next Step: At this stage, the installation differs depending on whether you are installing the Diagnostics Server in Commander mode or in Mediator mode.

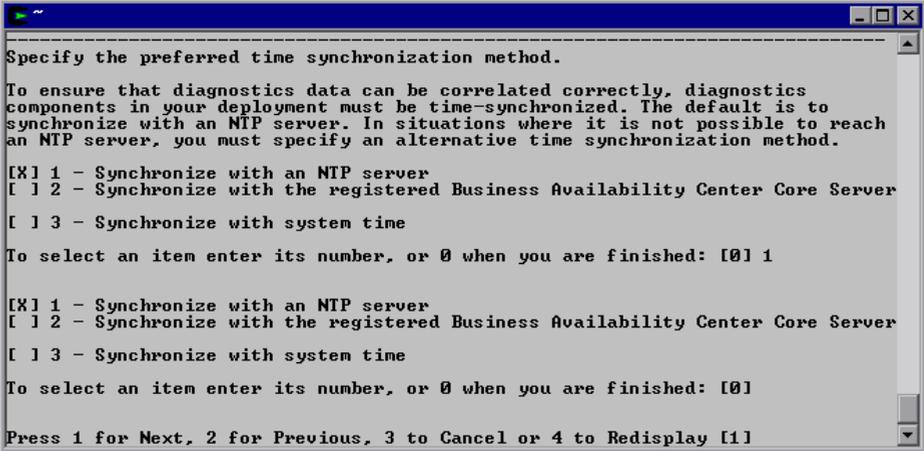
- ▶ To install the Diagnostics Server in Commander mode, continue with “Installing the Diagnostics Server in Commander Mode,” below.
 - ▶ To install the Diagnostics Server in Mediator mode, continue with “Installing the Diagnostics Server in Mediator mode” on page 49.
-

Installing the Diagnostics Server in Commander Mode

If you are installing the Diagnostics Server in Commander mode, continue as follows:

1 Select a time synchronization method.

In order for diagnostics data to be correlated properly, all the components in your deployment must be time-synchronized.



```

Specify the preferred time synchronization method.

To ensure that diagnostics data can be correlated correctly, diagnostics
components in your deployment must be time-synchronized. The default is to
synchronize with an NTP server. In situations where it is not possible to reach
an NTP server, you must specify an alternative time synchronization method.

[X] 1 - Synchronize with an NTP server
[ ] 2 - Synchronize with the registered Business Availability Center Core Server
[ ] 3 - Synchronize with system time

To select an item enter its number, or 0 when you are finished: [0] 1

[X] 1 - Synchronize with an NTP server
[ ] 2 - Synchronize with the registered Business Availability Center Core Server
[ ] 3 - Synchronize with system time

To select an item enter its number, or 0 when you are finished: [0]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

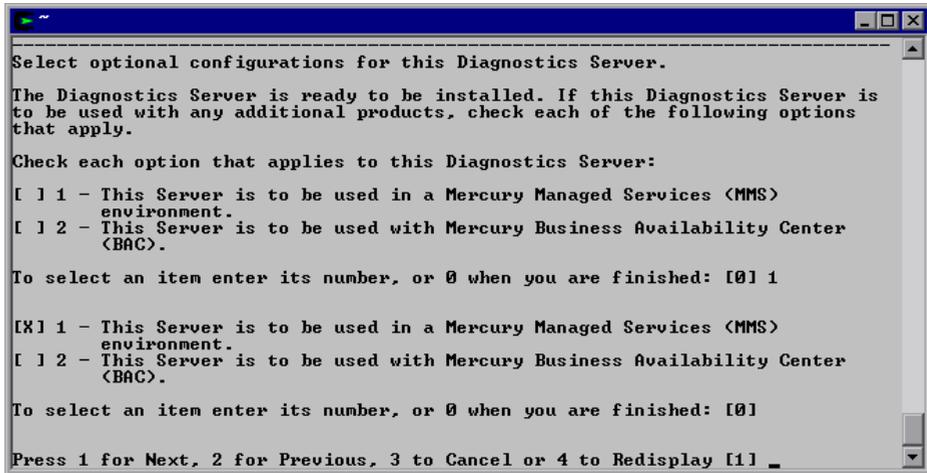
```

The following are the available synchronization options:

- **Synchronize with an NTP server.** This is the default method. This option applies only if the Diagnostics Server can access an NTP Server outside the firewall.
- **Synchronize with the registered Mercury Business Availability Center core server.** If the Diagnostics Server is to work in a Mercury Business Availability Center environment, select this option to synchronize with the Mercury Business Availability Center core server.
- **Synchronize with system time.** Select this option, if the Diagnostics Server is to work in an environment other than Mercury Business Availability Center and there is no access to an NTP server.

Select a time-synchronization option, and select **Next** and continue with the installation.

2 Indicate whether the Diagnostics Server will be used in a Mercury Managed Services environment or with Mercury Business Availability Center.



Select the options that apply to this Diagnostics Server, and select **Next** to continue with the installation.

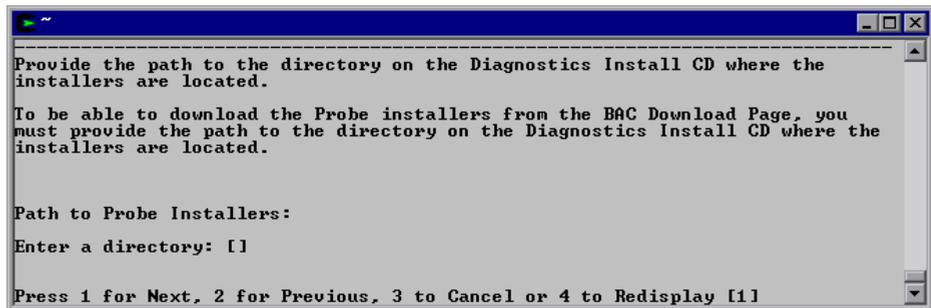
Next Step: If you selected only the Mercury Managed Services (MMS) option, skip to step 4.

3 Mercury Business Availability Center environment only: Provide the path to the Probe installers.

Notes:

- ▶ You need to have the Mercury Diagnostics installation CD available for this step.
 - ▶ You can skip this step at this stage, and enter the path manually when you configure the Diagnostics Server, or you can access the Probe installers directly from the product CD.
 - ▶ If you are using the Diagnostics Server with LoadRunner or Performance Center, or in Standalone mode, you should skip this step.
-

To be able to download the Diagnostics Probe installers from the Diagnostics Configuration page in Mercury Business Availability Center, you need specify the path to the directory on the Mercury Diagnostics installation CD where the Probe installers are located (**\Diagnostics_Probes**). The Probe installers are copied to the Diagnostics Server installation directory, which Mercury Business Availability Center can access.



Note: The installers for the Diagnostics Collectors are not stored on the same CD as the Probe installers. For this reason, the Collector Installers are not automatically copied by the installers.

If you would like to be able to download the Collector installers from Mercury Business Availability Center, you must copy the installers to <diag_server_install_dir>/html/opal/downloads.

At the **Enter a directory** prompt, type the path to the Probe installers on the Diagnostics Installer CD, and select **Next** to continue with the installation.

4 Review the pre-installation summary.

The installation settings that you selected are displayed. Review the information to make sure that you are satisfied.

Note: The estimated total size of the Diagnostics Server in Commander mode installation does not include the size of the Probe installers, if they were made available for Mercury Business Availability Center.

To change your settings, select **Previous** to return to the previous prompts.

To start the installation of the Diagnostics Server, select **Next**.

5 Close the installation wizard.

When the installation has completed, review the post-installation summary information to make sure that the installation completed successfully.

Select **Finish** to exit the installation.

6 Upload the Diagnostics license file.

In order to view performance metrics, you must upload a valid Mercury Diagnostics license file after you have installed the Diagnostics Server in Commander mode.

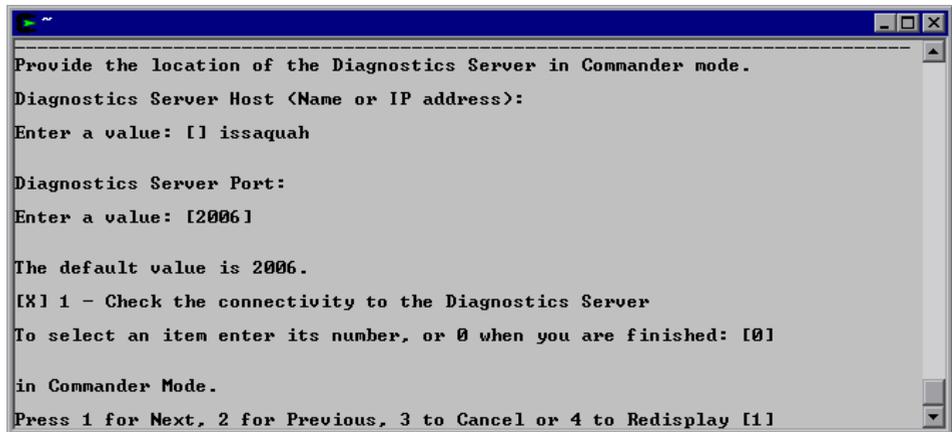
For instructions on requesting and uploading a valid license file, see Chapter 3, “Licensing Mercury Diagnostics.”

Installing the Diagnostics Server in Mediator mode

If you are installing the Diagnostics Server in Mediator mode, continue the installation as follows:

1 Provide the location of the Diagnostics Server in Commander mode.

Provide the information that enables the Diagnostics Server in Mediator mode to connect to the Diagnostics Server in Commander mode.



- a Enter the host name for the Diagnostics Server in Commander mode.
- b Enter the port for the Diagnostics Server in Commander mode.

The default port for the Diagnostics Server in Commander mode is **2006**. If you have changed the port since the Diagnostics Server was installed, you should specify that port number here instead of the default. For information on changing the Diagnostics Server port, see “Changing the Default Diagnostics Server Port” on page 250.

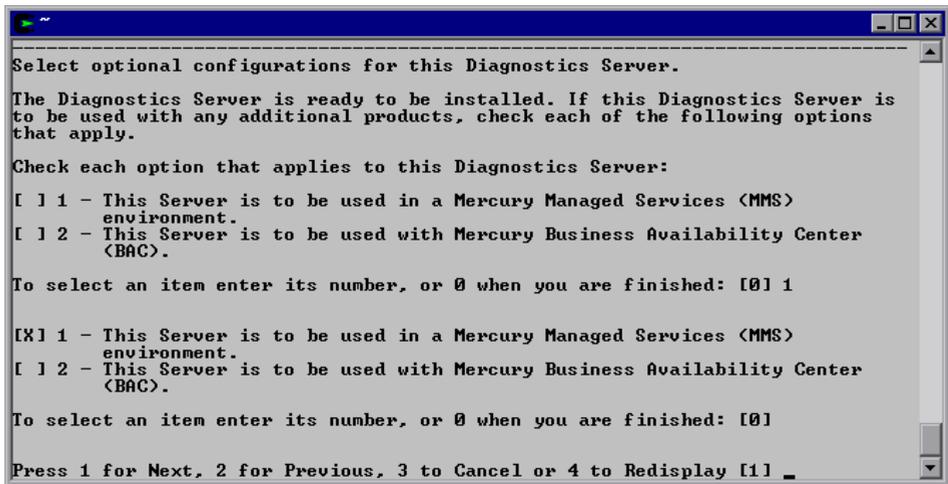
- c To allow the installer to check the connectivity to the host and port that you specified, select **Check the Connectivity to the Diagnostics Server**.

Select **Next** to continue with the installation.

If you instructed the installer to perform the test for connectivity, it tests the connectivity at this point. If there are negative results, it reports these before proceeding with the next installation step.

If you do not want to address these problems at this stage, clear the **Check the connectivity to the Diagnostics Server** option so that the installation can proceed.

- 2 Indicate whether the Diagnostics Server will be used in a Mercury Managed Services environment or with Mercury Business Availability Center.

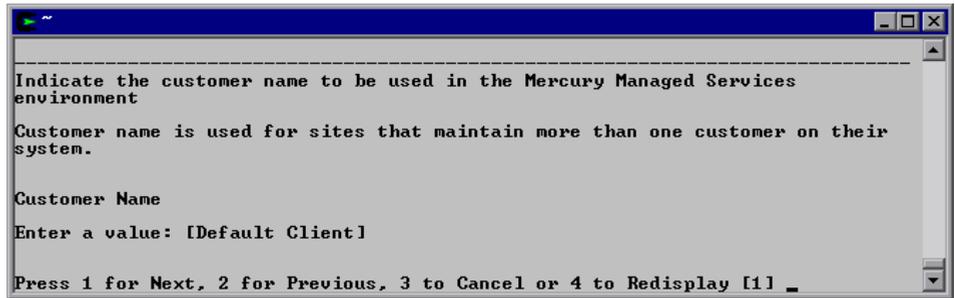


Select the options that apply to this Diagnostics Server, and select **Next** to continue with the installation.

Next Step: If you selected only the **Mercury Business Availability Center** option, skip to step 4.

3 If you indicated that the Diagnostics Server is to be used in a Mercury Managed Services environment, provide the customer name.

Enter a unique name for the customer in the Mercury Managed Services environment.



Select **Next** to continue with the installation.

4 Review the pre-installation summary.

The installation settings that you selected are displayed. Review the information to make sure that you are satisfied.

To change your settings, select **Previous** to return to the previous prompts.

To start the installation of the Diagnostics Server, select **Next**.

5 Close the installation wizard.

When the installation has completed, review the post-installation summary information to make sure that the installation completed successfully.

Select **Finish** to exit the installation.

Starting and Stopping the Diagnostics Server

Instructions for a Windows Machine

To start the Diagnostics Server on a Windows machine:

Select **Start > Programs > Mercury Diagnostics Server > Start Mercury Diagnostics Server**.

To stop the Diagnostics Server on a Windows machine:

Select **Start > Programs > Mercury Diagnostics Server > Stop Mercury Diagnostics Server**.

Instructions for a Solaris Machine

The *nanny* is a process that runs as a daemon to ensure that the Diagnostics Server is always running. The following procedures start and stop the Diagnostics Server using the nanny.

To start the Diagnostics Server on a Solaris machine:

- 1 Ensure that the **M_LROOT** environment variable is defined as the root directory of the Diagnostics Server.

For example, in *tcsh*, you could enter the following:

```
$> setenv M_LROOT <diagnostics_server_install_dir>/nanny/solaris
```

If the **M_LROOT** environment variable is not defined as the root directory, you will encounter the following error:

```
Warning : MDRV: cannot find lrun root directory . Please check your  
M_LROOT  
Unable to format message id [-10791]  
m_agent_daemon ( is down )
```

- 2 Change directories to **<diagnostics_server_install_dir>/nanny/solaris/bin**.
- 3 Run **m_daemon_setup** with the **-install** option, as in the following example:

```
$> ./m_daemon_setup -install
```

To stop the Diagnostics Server on a Solaris machine:

- 1** Change directories to `<diagnostics_server_install_dir>/nanny/solaris/bin`.
- 2** Run `m_daemon_setup` with the `-remove` option, as in the following example:

```
$>./m_daemon_setup -remove
```

Instructions for a UNIX Machine

The following procedures start and stop the Diagnostics Server without using the nanny.

To start the Diagnostics Server on a UNIX machine:

Run `<diagnostics_server_install_dir>/bin/server.sh`.

To stop the Diagnostics Server on a UNIX machine:

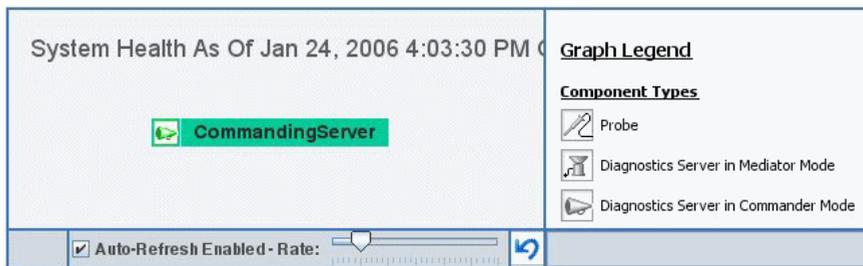
Terminate the process using a utility such as “kill.”

Verifying the Diagnostics Server Installation

To verify that the Diagnostics Server has been installed correctly, and that it has started properly, use the System Health Monitor. (For instructions on starting the System Health Monitor, see the *Mercury Diagnostics User's Guide*.)

According to the recommended installation sequence, after you install the Diagnostics Server you can use the System Health Monitor to verify that the Diagnostics Server was installed and started.

The Diagnostics Server in Commander mode should be displayed on the System Health Monitor as shown in the following example:



When a Diagnostics Server is deployed in Commander mode, it has both commanding and mediating responsibilities. The Diagnostics Server in Commander mode is represented in the System Health Monitor component map by an icon labelled **Commanding Server**.

When a Diagnostics Server is deployed in Mediator mode, it is represented by a single icon labelled **server-<name of host>**.

The System Health Monitor is part of the Diagnostics Server component. If you are unable to access the System Health Monitor after the Diagnostics Server has been installed, then either you are entering an incorrect URL or the Diagnostics Server did not start. For instructions on starting the Diagnostics Server, see “Starting and Stopping the Diagnostics Server” on page 52.

You can leave the System Health Monitor displayed in your browser to verify the progress of the component installations, and to identify and troubleshoot any problems that you encounter as you proceed through the rest of the component installations.

Licensing the Diagnostics Server in Commander Mode

After you have installed the Diagnostics Server in Commander mode you must upload a valid Diagnostics license file to the installation directory for the Diagnostics Server before you will be able to view the performance metrics in the views of the Diagnostics user interface.

For instructions on requesting a license file and uploading the file to the Diagnostics Server installation directory, see Chapter 3, “Licensing Mercury Diagnostics.”

Configuring the Diagnostics Server

The Diagnostics Server is installed with a default configuration that enables it to perform effectively in most situations. You may encounter situations where changing the configuration enables better Diagnostics Server performance, or allows it to work in unusual situations.

For information about configuring the Diagnostics Server, see Chapter 14, “Advanced Diagnostics Server Configuration.”

Determining the Version of the Diagnostics Server

When you request support, it is useful to know the version of the Diagnostics Server.

To determine the version of the Diagnostics Server:

Locate the version file, `<diagnostics_server_install_dir>\version.txt`. This file contains the version number, as well as the build number.

Note: This version number is also available at the bottom of the maintenance page of the Diagnostics Server, in the details of the Diagnostics Server component in System Health, and in the main Diagnostics Server log file, `<diagnostics_server_install_dir>\log\server.log`.

The version of the Diagnostics Server and the version of the Mercury LoadRunner Diagnostics AddIn must be exactly the same. For more information, see “J2EE Probe Fails to Operate Properly” on page 504.

Uninstalling the Diagnostics Server

The following section contains instructions for uninstalling the Diagnostics Server from a Windows or Solaris machine.

Uninstalling the Diagnostics Server From a Windows Machine

- ▶ Uninstall the Diagnostics Server by selecting **Start > Programs > Mercury Diagnostics Server > Uninstall Mercury Diagnostics Server**.
- ▶ Alternatively, you can execute **uninstaller.exe** which is located in the **<diagnostics_server_install_dir>_uninst** directory.

During the uninstallation process, a message asks if you want to remove specific files. Do the following:

- ▶ If you want to completely uninstall the Diagnostics Server, as well as any property settings, click **Yes** or **Yes to All**.
- ▶ If you plan on reinstalling the Diagnostics Server, and want to keep the custom property settings of the Diagnostics Server you are uninstalling, then you need to back up the property files located in the **etc** directory to a new location.

If you have backed up these files, click **Yes** or **Yes to All**.

If you have not yet backed up these files, select **No** or **No to All**.

Uninstalling the Diagnostics Server From a Solaris Machine

You can uninstall the Diagnostics Server in console mode or graphical mode.

To uninstall the Diagnostics Server:

- 1 Stop the Diagnostics Server. For instructions, see “Starting and Stopping the Diagnostics Server” on page 52.
- 2 Change the directory to the root directory.
- 3 Enter the following at the UNIX command prompt:

► **In console mode:**

```
<diagnostics_server_install_dir>/jre/bin/java -jar  
<diagnostics_server_install_dir>/_uninst/uninstall.jar -console
```

► **In graphical mode:**

```
<diagnostics_server_install_dir>/jre/bin/java -jar  
<diagnostics_server_install_dir>/_uninst/uninstall.jar
```

Troubleshooting Diagnostics Server Issues

To troubleshoot Diagnostics Server problems, use the System Health Monitor. For more information, see “Using the System Health Monitor” on page 455.

3

Licensing Mercury Diagnostics

This chapter describes how to license Mercury Diagnostics.

This chapter describes:	On page:
About Mercury Diagnostics Licensing	59
Types of Licenses	60
Licensing the Diagnostics Server in Commander Mode	60

About Mercury Diagnostics Licensing

Diagnostics components can be installed before you have received the license. They can also be started so that they will begin to monitor your applications and process the performance metrics. However, the performance metrics that you can view the using the views of the Diagnostics user interface or the tabs in the Profilers is load-limited until you have provided a valid license file.

Mercury Diagnostics is licensed using a file that you upload to the Diagnostics Server in Commander mode. The license uses node-locking based on the MAC address for the host of the Diagnostics Server.

After you have installed the Diagnostics Server in Commander mode, you must ask your Mercury Customer Support representative to generate a license file for the Diagnostics deployment using the MAC address of the Diagnostics Server host which you provide to them.

When you receive the license file, you must store it in a location accessible by the host of the Diagnostics Server. You then upload the file using the License Upload area of the Diagnostics Server License Management page.

The Diagnostics Server in Mediator mode and Probes associated with the licensed Diagnostics Server in Commander mode are licensed based upon the license installed for the Diagnostics Server in Commander mode.

Types of Licenses

There are two types of licenses that can be issued for Diagnostics:

- A *perpetual* license is generated without an expiration date.
- A *custom* license is generated with a built in expiration date. The expiration date is set so that the license will be valid for a specific length of time. This type of license allows for a trial period before you purchase a perpetual license.

Licensing the Diagnostics Server in Commander Mode

You obtain your Diagnostics license from you Mercury Support contact, and then upload it to the Diagnostics Server installation directory.

To license your Diagnostics deployment:

- 1** Determine the MAC address for the host of the Diagnostics Server in Commander mode.
- 2** Request a license from your Mercury Support contact.
- 3** Upload the license file of the Diagnostics Server in Commander mode.
- 4** License the other Diagnostics components. When the Probes and Diagnostics Server in Mediator mode first connect with the Diagnostics Server in Commander mode they are licensed based upon the license installed on the Diagnostics Server in Commander mode.

Determine the MAC Address for the Diagnostics Server Host

The Diagnostics License is a node-locked license that is keyed off of the MAC address for the host of the Diagnostics Server in Commander mode.

To determine the MAC address for the Diagnostics Server Host:

- 1 Open the License Management page for the Diagnostics Server in Commander mode. From the Diagnostics Server administration menu select **Configure Diagnostics -> Licensing**.

The License Management page opens.

MERCURY™

License Management

System Information

Attribute	Value
MAC Address:	00:0D:56:70:50:6B, 00:0D:56:70:50:6A, 00:10:18:06:42:D8
<input type="button" value="Refresh"/>	

License Information

Attribute	Value
License:	Diagnostics Server License
Type:	Custom
Registered To:	Developer at Mercury
Issue Date:	Mar 17, 2006
Expiration:	Jul 1, 2006
Maximum Java Probes:	Unlimited
Maximum .NET Probes:	Unlimited
Maximum SAP Probes:	Unlimited
MAC Address:	Any

License Upload

Note: You may only upload files ending in ".lic". The uploaded file will be renamed to "DiagnosticsServer.lic".

License File:

Starting Mercury Diagnostics Server "server-helena", version 4.1.36.418

- 2 Note the MAC Address value displayed in the System Information section of the page. Click **Refresh** to make sure that the Diagnostics Server has discovered the current MAC Address information.

Be sure to let your Mercury Support contact know if the MAC Address value is **Unknown** when you request the license file. Do not look up the MAC Address using another method.

Note: You must use the MAC Address that is displayed in the System Information section of the License Management page when requesting a license. If the Diagnostics Server cannot determine a MAC address to display on the License Management page, it will not be able to find the MAC address when it attempts to validate a MAC address based license.

Request a License From Your Mercury Support Contact

Your Mercury Support contact can generate a license file for your Diagnostics Server host using the MAC Address value that you found on the License Management page.

To request a Diagnostics License:

- 1 Contact Mercury Support and provide them with the MAC Address that was displayed on the License Management page.
- 2 The Mercury Support contact generates a license file and sends it to you.

Upload the License File

When you receive the license file that was generated for your Diagnostics deployment you must upload the file to the installation directory for the Diagnostics Server in Commander mode using the License Management page of the Server's Maintenance pages.

To upload the Diagnostics license:

- 1 Store the license file in a directory that can be accessed from the License Management page for the Diagnostics Server in Commander mode.

Note: The name of the file to be uploaded must end with the extension **.lic**.

- 2 In the License Upload section of the License Management page, type the path to the location where the license file was stored or click **Browse** to navigate to the license file location.

License Upload

Note: You may only upload files ending in ".lic". The uploaded file will be renamed to "DiagnosticsServer.lic".

License File:

- 3 Click **Upload** to apply the license file to the Diagnostics Server.

The file is renamed by the upload process and stored in the proper location in the Diagnostics Server install directory.

Note: Do not attempt to copy the license file directly to the Diagnostics Server installation directory. Always upload the file using the License Upload section of the License Management screen.

Licensing the Other Diagnostics Components

The Mercury Diagnostics Profiler for .NET and Mercury Diagnostics Profiler for J2EE operate in an unlicensed mode with load restrictions until the probe is able to connect to a Diagnostics Server in Commander mode that has been properly licensed.

The Diagnostics Server in Mediator mode and Probes associated with the licensed Diagnostics Server in Commander mode are licensed based upon the license installed for the Diagnostics Server in Commander mode

When the Probes and Diagnostics Server in Mediator mode first connect with the Diagnostics Server in Commander mode they are licensed based upon the license installed on the Diagnostics Server in Commander mode.

Part III

Installation and Configuration of the Mercury Diagnostics Probes

4

Installing the Diagnostics Collector

This chapter provides instructions for installing a Diagnostics Collector on Windows and UNIX machines.

This chapter describes:	On page:
About Installing the Diagnostics Collector	68
Installing the Diagnostics Collector on a Windows Machine	68
Installing the Diagnostics Collector on a UNIX Machine	76
Silent Installation of the Diagnostics Collector	83
Configuring the Active System Property Files	84
Verifying the Diagnostics Collector Installation	87
Starting and Stopping the Diagnostics Collector	89
Determining the Version of the Diagnostics Collector	90
Uninstalling the Diagnostics Collector	90

About Installing the Diagnostics Collector

The Diagnostics Collector gathers data from external active systems. The Collector can be configured to collect performance data from two types of active systems:

- ▶ SAP R/3 system
- ▶ Oracle 10g Database

During the installation of the Collector, you can choose to monitor either one or both of these active systems. After the installation, you define instances of Oracle 10g and SAP R/3 systems to be monitored. Each one of these instances is represented as a Probe (either Oracle Probe or SAP R/3 Probe) belonging to a Probe group in the Mercury Diagnostics UI.

Installing the Diagnostics Collector on a Windows Machine

The following steps provide detailed instructions for installing the Collector on a Windows machine. These instructions also apply when you are installing the Collector on a UNIX machine using the graphical installer.

Note: The Collector can be installed on any machine. It does not necessarily have to be installed on the host machine of the SAP or Oracle application. For Collector host requirements, see “Requirements for the Diagnostics Collector Host” on page 18.

Launching the Installation

The installation can be launched directly from the installation CD, or from the Diagnostics Downloads page in Mercury Business Availability Center.

To launch the installation from the installation CD:

- 1 Insert the Mercury Diagnostics for Windows installation CD into the CD-ROM drive. If the installer does not launch automatically, double-click **setup.exe** in the root folder of the CD.

Note: To install the Collector from a different location, locate the executable file **CollectorSetupWin.exe** file in the **Diagnostics_Collector** directory on the CD, copy it to the new location, and then launch it.

The installer displays the Mercury Diagnostics main installation menu.

- 2 To start the installation for the Diagnostics Collector, click **Mercury Diagnostics Collector**.

Continue with “Running the Installation,” below.

For Mercury Business Availability Center Users

Note: To be able to download the Collector installer from the Diagnostics Downloads page in Mercury Business Availability Center, locate the **\Diagnostics_Collector\CollectorSetupWin.exe** on the Diagnostics CD, and copy it to the following folder of the Diagnostics Server installation directory: **<diag_server_install_dir>/html/opal/downloads**.

To start the installation from the Mercury Business Availability Center downloads page:

- 1** From the top menu in Mercury Business Availability Center, select **Admin > Diagnostics**, and click the **Downloads** tab.
- 2** On the Downloads page, click the appropriate link to launch the Collector installation for Windows.

Continue with “Running the Installation,” below.

Running the Installation

After you have launched the installation, the software license agreement opens.

To install the Collector:

- 1** **Accept the software license agreement.**

Read the agreement and select **I accept the terms of the license agreement**.

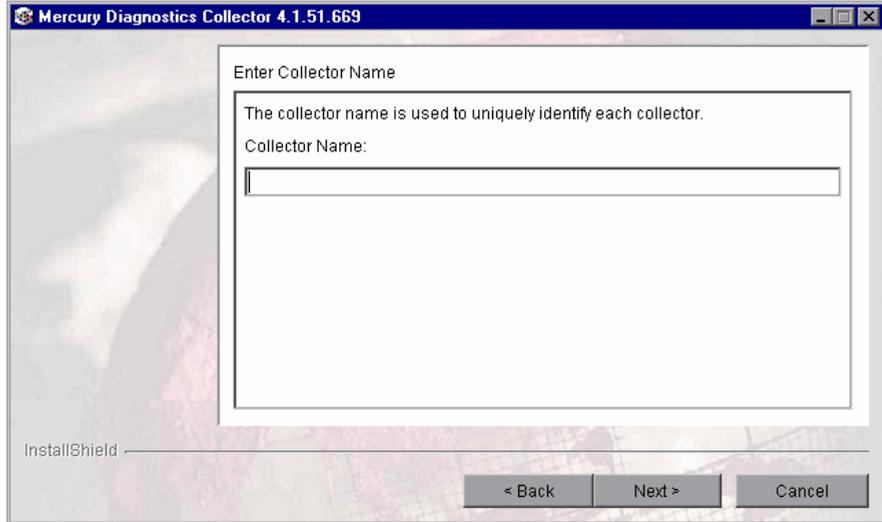
Click **Next** to continue.

- 2** **Specify the location where you want to install the Collector.**

In the **Installation Directory Name** box, type the name of the directory where you want to install the Collector. Alternatively, accept the default directory, **C:\MercuryDiagnostics\Collector**, or click **Browse** to navigate to another directory.

Click **Next** to continue.

3 Assign a unique name to the Collector.

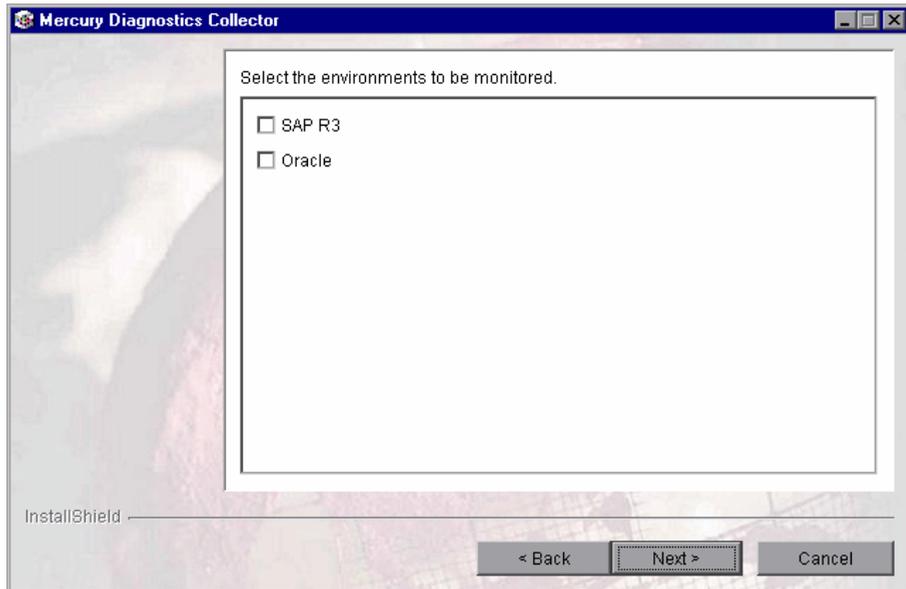


Assign a name to the Collector, that will uniquely identify this specific Collector in the System Health Monitor. You use the System Health Monitor to verify the Collector installation and configuration. For more information, see “Verifying the Diagnostics Collector Installation” on page 87.

The following characters can be used in the name: - , _ , and all alphanumeric characters.

Click **Next** to continue.

4 Select the environment to monitor.



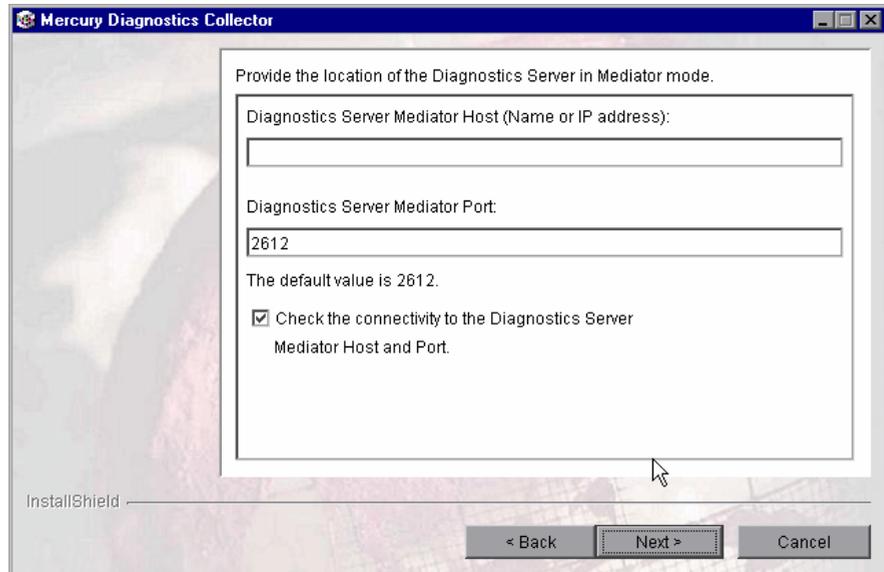
Select the options that apply to this Collector. You can select either one or both options.

- To collect data in an SAP R3 environment, select **SAP R3**.
- To collect data on an Oracle 10g database server, select **Oracle**.

Click **Next** to continue.

5 Provide information about the Diagnostics Server in Mediator mode.

Provide the details that enables communication with the Diagnostics Server in Mediator mode.



If there is only one Diagnostics Server in the Diagnostics deployment where the Collector will run, enter the host name of the Diagnostics Server and its event port information.

If there is more than one Diagnostics Server in the deployment, enter the information for the Diagnostics Server in Mediator mode that is to receive the events from the Collector.

- a** In the **Diagnostics Server Mediator Host** box, type the host name or IP address of the host for the Diagnostics Server in Mediator mode.

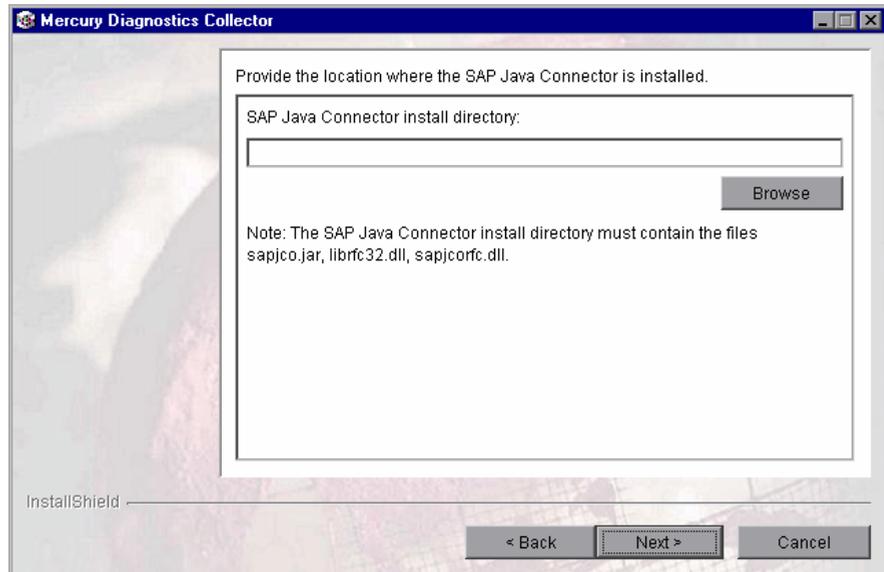
Note: You should specify the fully qualified host name. In a mixed OS environment, where UNIX is one of the systems, this is essential for proper network routing.

- b** In the **Diagnostics Server Mediator Port** box, type the port number where the Diagnostics Server is listening for Collector communication. The default port number is **2612**. If you have changed the port since the Diagnostics Server was installed, specify that port number instead of the default.
- c** If you want to make sure that the Diagnostics Server is running and accessible from the installation host, select **Check the connectivity to the Diagnostics Server Mediator Host and Port**.

Click **Next** to continue.

Note: If you selected **Check the connectivity to the Diagnostics Server Mediator Host and Port** and connectivity problems were encountered, the installer provides you with the results of the connectivity check. If you do not want to address these problems at this stage, clear the **Check the connectivity to the Diagnostics Server Mediator Host and Port** check box, proceed with the installation, and address the problem later.

6 If you selected SAP R3 in step 4, provide the location of the SAP Java Connector.



In the **SAP Java Connector install directory** box, enter the name of the directory where the SAP Java Connector is installed.

This directory must contain the following files:

- sapjco.jar
- librfc.dll
- sapjcorfc.dll

Notes:

- If you do not know the SAP Java Connector directory name, contact your SAP representative.
 - If any of the files is missing from this directory, contact your SAP representative.
-

7 Review the pre-installation summary.

The installation settings that you selected are displayed in a read-only window. Review the information to make sure that you are satisfied.

To select different installation settings, click **Back**.

To begin installation, click **Next**.

8 Close the installation wizard.

When the installation has completed, a message is displayed, confirming that the Collector was successfully installed. Click **Finish** to exit the installation.

9 Configure the XML files for your active systems.

In step 4 you selected the active systems that will be monitored. For each of these active systems, you need to configure properties that enable the Collector host and the active system host to communicate.

For instructions on configuring the relative active system properties, see “Configuring the Active System Property Files” on page 84.

10 Verify that Collector was installed properly, and is running.

Using the System Health Monitor, you can verify that the Collector is running as it should be. For details see, “Verifying the Diagnostics Collector Installation” on page 87.

Installing the Diagnostics Collector on a UNIX Machine

The following instructions provide you with the steps necessary to install the Collector in Solaris and Linux environments, using either a graphical installation or a console mode installation.

The installation screens that are displayed in a graphical installation are the same as those documented for the Windows installation in “Installing the Diagnostics Collector on a Windows Machine” on page 68.

Launching the Installation

The installer can be launched directly from the installation CD, or from the Diagnostics Downloads page in Mercury Business Availability Center.

To launch the installation from the installation CD:

Copy the UNIX installer from the **Diagnostics_Collector** directory on the Diagnostics Collector installation CD to the machine where the Collector is to be installed.

Continue with “Running the Installation,” below.

For Mercury Business Availability Center users:

To start the installation from the Mercury Business Availability Center downloads page:

- 1** From the top menu in Mercury Business Availability Center, select **Admin > Diagnostics**, and click the **Downloads** tab.
- 2** On the Downloads page, click the link to the installer that is appropriate for your environment.
- 3** Save the installer on the machine where the Collector is to be installed.

Note: The Collector installer is available in Mercury Business Availability Center only if you provided the path to the Probe installers directory when you installed the Diagnostics Server in Commander mode. See “Installing the Diagnostics Server in Commander Mode” on page 32.

Continue with “Running the Installation,” below.

Running the Installation

After you have copied the installer to the machine where the Collector is to be installed, you are ready to run the installation.

Note: The following instructions assume an understanding of UNIX console screens and commands. For more information about UNIX screens and commands, see Appendix H, “Using UNIX Commands.”

To install the Collector:

1 Run the installer.

Where necessary, change the mode of the installer file to make it executable.

- ▶ To run the installer in the graphical mode, enter the following at the UNIX command prompt:

```
./<installer>
```

The installer displays the same screens that are displayed for the Windows installer.

Continue with “Installing the Diagnostics Collector on a Windows Machine” on page 68.

- ▶ To run the installer in console mode, enter the following at the UNIX command prompt:

```
./<installer> -console
```

The installer continues in console mode, and displays the software license agreement.

Continue with step 2.

2 Accept the software license agreement.

Read the license agreement and accept the terms.

Select **Next** to continue with the installation.

3 Specify the location where you want to install the Collector.

At the **Directory Name** prompt, accept the default installation location shown in brackets, or type the path to a different installation location.

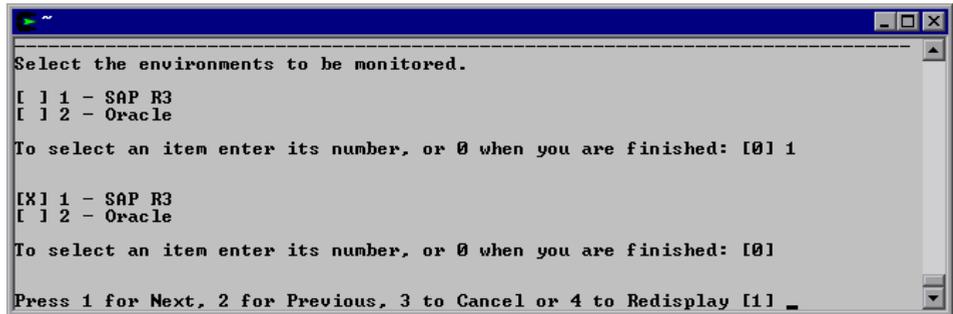
Select **Next** to continue with the installation.

4 Assign a unique name to the Collector.

Assign a name to the Collector, that will uniquely identify this specific Collector in the System Health Monitor. You use the System Health Monitor to verify the Collector installation and configuration. For more information, see “Verifying the Diagnostics Collector Installation” on page 87.

The following characters can be used in the name: -, _, and all alphanumeric characters.

Select **Next** to continue.

5 Specify the active systems that the Collector will monitor.

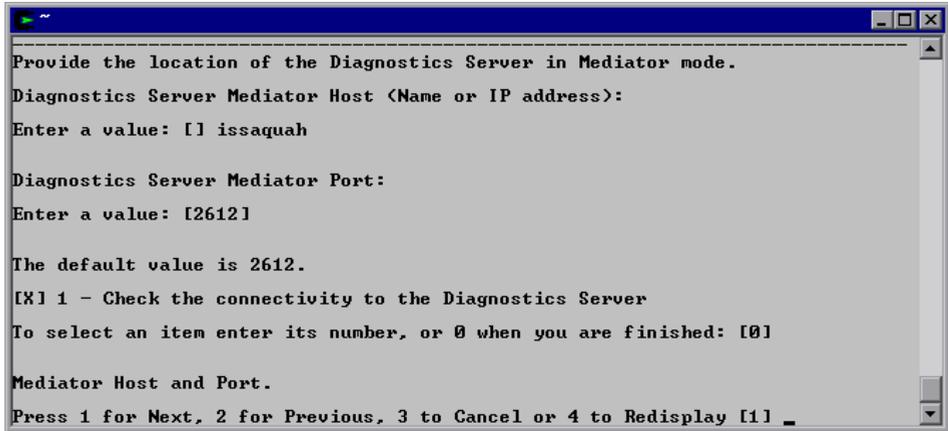
Select the active systems that apply to this Collector. You can select either one or both options.

- To collect data on a SAP R3 system, select **SAP R3**.
- To collect data on an Oracle 10g database server, select **Oracle**.

Select **Next** to continue with the installation.

6 Provide the details for the Diagnostics Server in Mediator mode.

Provide the information that enables communication with the Diagnostics Server in Mediator mode.



If there is only one Diagnostics Server in the Diagnostics deployment where the Collector will run, enter the host name of the Diagnostics Server and its event port information.

If there is more than one Diagnostics Server in the deployment, enter the information for the Diagnostics Server in Mediator mode that is to receive the events from the Collector.

- a Enter the host name or IP address of the host of the Diagnostics Server in Mediator mode.

Note: You should specify the fully qualified host name. In a mixed OS environment, where UNIX is one of the systems, this is essential for proper network routing.

- b Enter the number of the port where the Diagnostics Server is listening for Collector communication. The default port is **2612**. If you have changed the port since the Diagnostics Server was installed, specify that port number instead of the default.

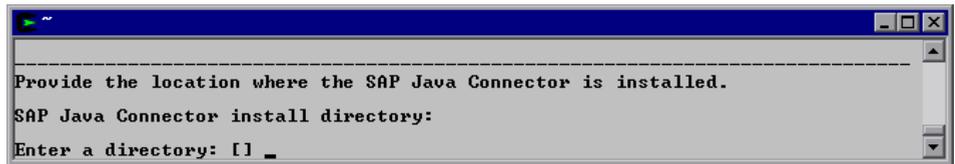
- c If you want to make sure that the Diagnostics Server is running and accessible from the installation host, select **Check the connectivity to the Diagnostics Server**.

Select **Next** to continue with the installation.

If you instructed the installer to perform the test for connectivity, it tests the connectivity at this point. If there are negative results, it reports these before proceeding with the next installation step.

If you do not want to address these problems at this stage, clear the **Check the connectivity to the Diagnostics Server** option so that the installation can proceed.

7 Provide the location of the SAP Java Connector.



At the prompt, type the name of the directory where the SAP Java Connector is installed.

This directory must contain the following files:

- sapjco.jar
- librfc.dll
- sapjcorfc.dll

Notes:

- If you do not know the SAP Java Connector directory name, contact your SAP representative.
 - If any of the files is missing from this directory, contact your SAP representative.
-

Select **Next** to continue with the installation.

8 Review the pre-installation summary.

The installation settings that you selected are displayed. Review the information to make sure that you are satisfied.

To change your settings, select **Previous** to return to the previous prompts.

To start the installation of the Collector, select **Next**.

9 Close the installation wizard.

When the installation has completed, review the post-installation summary information to make sure that the installation completed successfully.

Select **Finish** to exit the installation.

10 Configure the XML files for your active systems.

In step 5 you selected the active systems that will be monitored. For each of these active systems, you need to configure properties that enable the Collector host and the active system host to communicate.

For instructions on configuring the active system properties, see “Configuring the Active System Property Files” on page 84.

11 Verify the Collector installation.

Using the System Health Monitor, you can verify that the Collector is running as it should be. For details, see “Verifying the Diagnostics Collector Installation” on page 87.

Silent Installation of the Diagnostics Collector

A *silent installation* is an installation that is performed automatically, without the need for user interaction. In place of user input, the silent installation accepts input from a response file for each step of the installation.

For example, a system administrator who needs to deploy a component on multiple machines can create a response file that contains all the prerequisite configuration information, and then perform a silent installation on multiple machines. This eliminates the need to provide any manual input during the installation procedure.

Before you perform silent installations on multiple machines, you need to generate a response file that will provide input during the installation procedure. This response file can be used in all silent installations that require the same input during installation.

To generate a response file:

Perform a regular installation with the following command-line option:

```
<installer> -options-record <responseFileName>
```

This creates a response file that includes all the information submitted during the installation.

To perform a silent installation:

Perform a silent installation using the relevant response file.

You perform the silent installation with the **-silent** command-line option as follows:

```
<installer> -silent -options <responseFileName>
```

Configuring the Active System Property Files

When you install the Collector, you are asked to indicate the active systems that the Collector will monitor. After installation, you define instances of the active systems to be monitored. These instances are manually defined in the XML files provided with the installation. An instance definition in the XML file behaves like a probe on the instance of the active system.

Configuring SAP R3 Probes

SAP R3 system deployment can include one or more SAP R3 application instances. These instances together form an SAP R3 system.

Depending on user permissions, access to the system or application instances on the system may be direct, or may require connection through the SAP Message Server. For each SAP R3 Probe that you define, you need to know what connection option is used.

You define and configure SAP R3 Probes in the `<collector_install_dir>\etc\r3config.xml` file. The layout, elements, and attributes of the xml file are described in the `<collector_install_dir>\etc\r3config.xsd`.

To configure an SAP R3 Probe:

- 1 Open `<collector_install_dir>\etc\r3config.xml`.
- 2 If you are defining an SAP R3 Probe where access to the SAP R3 instance is through the SAP Message Server, locate the section of code preceded by the following comment:

```
<!--  
  Template to be used with the message server connection option.  
-->
```

If you are defining an SAP R3 Probe where access to the SAP R3 instance is direct, locate the section of code preceded by the following comment:

```
<!--  
  Template to be used with the direct connection option.  
-->
```

- 3 Make a copy of the comment together with the template code below the comment, and paste it at the end of the file.
- 4 Comment out the original template code by typing `<!--` in an empty line above the template code, and `-->` in an empty line thereafter.
- 5 In the copied code at the end of the file, alter the value of each property, as described in the following table and then save the file:

Property	Description	Value
r3system name	A logical name for the probe group under which this SAP Probe appears in the Diagnostics UI.	User-defined.
systemId	The ID of the SAP R3 system. Comprises 3 characters only.	Format: [XXX] Obtainable from the SAP system administrator.
client	The client name for the SAP R3 system.	Obtainable from the SAP system administrator.
user	The name of the user connecting to the SAP R3 system.	Obtainable from the SAP system administrator.
password	The password of the user connecting to the SAP R3 system.	Obtainable from the SAP system administrator.
messageServerHost (Message Server connection only)	The name of the SAP Message Server host machine.	Obtainable from the SAP system administrator.
r3Name (Message Server connection only)	Comprises 3 characters only.	Format: [XXX] Obtainable from the SAP system administrator.
group (Message Server connection only)		Obtainable from the SAP system administrator.

Configuring Oracle Probes

You define and configure the Oracle Probes in `<collector_install_dir>\etc\oracle-config.xml`. The layout, elements and attributes of the xml file are described in the `<collector_install_dir>\etc\oracle-config.xsd`.

To configure an Oracle Probe:

- 1 Open `<collector_install_dir>\etc\oracle-config.xml`.
- 2 Copy the template code, and paste it at the end of the file.
- 3 Comment out the template code by typing `<!--` in an empty line above the template code, and `-->` in an empty line thereafter.
- 4 In the copied code, alter the value of each property, as described in the following table and then save the file:

Properties	Description	Value
hostName	The name of the Oracle database server host machine.	Obtainable from the Oracle administrator.
portNumber	The number of the port where the Oracle database server listens for requests.	Default value: 1521
instanceName	The name given to the Oracle instance during installation of the Oracle database server.	Default value: Orcl Obtainable from the Oracle administrator.
userId	The ID of the user connecting to the Oracle database server. Note: The user need at least CREATE SESSION and SELECT ANY DICTIONARY to collect performance metrics.	Obtainable from the Oracle administrator.
password	The password of the user connecting to the Oracle database server.	Obtainable from the Oracle administrator.

Properties	Description	Value
probeName	The logical name to represent this Oracle instance in the Diagnostics UI. This name must be unique.	User-defined. If this value is not defined, the same value given for instanceName will be used here.
probeGroupName	The logical name of the Probe group under which this Probe appears in the Diagnostics UI. This can be an existing Probe group, or you can define a new one.	User-defined. For example: Existing: Default New: Oracle

Verifying the Diagnostics Collector Installation

The Collector starts running automatically when the installation is complete. You can verify the Collector installation using the System Health Monitor. For detailed instructions about how to use the System Health Monitor, see Appendix C, “Using the System Health Monitor.”

If you have been following the recommended installation sequence, after installing the Collector you can verify the following:

- The Diagnostics Server in Commander mode was successfully installed.
- Where relevant, additional Diagnostics Servers in Mediator mode were successfully installed and are communicating with the Diagnostics Server in Commander mode.
- The Collector was successfully installed and has established connectivity with the relevant Diagnostics Server.

Depending on your deployment, the Collector is shown on the System Health Monitor as a child of either the Diagnostics Server in Commander mode or a Diagnostics Server in Mediator mode.



When the Collector has been stopped, the Collector node in the System Health monitor is grey. When the Collector has been started the node is green.

To change the name of the Collector that appears on the System Health Monitor:

- 1 Open `<collector_install_dir>\etc\collector.properties` using a text editor.
- 2 Locate the following line in the code:

```
id = CollectorName
```

- 3 Replace **CollectorName** with a new name for the Collector.
- 4 Save the file. In order to see the new name on the System Health monitor, you must stop the Collector, and restart it. For details, see "Starting and Stopping the Diagnostics Collector," below.

Starting and Stopping the Diagnostics Collector

Instructions for a Windows Machine

To start the Collector on a Windows machine:

Select **Start > Programs > Mercury Diagnostics Collector > Start Mercury Diagnostics Collector**.

To stop the Collector on a Windows machine:

Select **Start > Programs > Mercury Diagnostics Collector > Stop Mercury Diagnostics Collector**.

Instructions for a Solaris Machine

The *nanny* is a process that runs as a daemon to ensure that the Collector is always running. The following procedures start and stop the Collector using the nanny.

To start the Collector on a Solaris machine:

- 1 Ensure that the **M_LROOT** environment variable is defined as the root directory of the Collector.

For example, in *tsh*, you could enter the following:

```
$> setenv M_LROOT <collector_install_dir>/nanny/solaris
```

If the **M_LROOT** environment variable is not defined as the root directory, you will encounter the following error:

```
Warning : MDRV: cannot find lrun root directory . Please check your
M_LROOT
Unable to format message id [-10791]
m_agent_daemon ( is down )
```

- 2 Change directories to **<collector_install_dir>/nanny/solaris/bin**.
- 3 Run **m_daemon_setup** with the **-install** option, as in the following example:

```
$> ./m_daemon_setup -install
```

To stop the Collector on a Solaris machine:

- 1 Change directories to `<collector_install_dir>/nanny/solaris/bin`.
- 2 Run `m_daemon_setup` with the `-remove` option, as in the following example:

```
$>./m_daemon_setup -remove
```

Instructions for a UNIX Machine

The following procedures start and stop the Collector without using the nanny.

To start the Collector on a UNIX machine:

Run `<collector_install_dir>/bin/collector.sh`.

To stop the Collector on a UNIX machine:

Terminate the process using a utility such as “kill.”

Determining the Version of the Diagnostics Collector

When you request support, it is useful to know the version of the Diagnostics Collector.

The version number of the Collector can be found:

- ▶ in the `<collector_install_dir>\version.txt` file
- ▶ among the details of the Collector on the System Health Monitor

Uninstalling the Diagnostics Collector

To uninstall the Collector:

- ▶ On a Windows machine, run `uninstaller.exe` which is located in the `<collector_install_dir>_uninst` directory.
- ▶ On a UNIX machine, run `uninstall*` which is located in the `<collector_install_dir>_uninst` directory.

5

Installing the Mercury Diagnostics Probe for .NET

This chapter describes how to install a Mercury Diagnostics Probe for .NET.

This chapter describes:	On page:
About the Mercury Diagnostics Probe for .NET	91
Installing the .NET Probe	92
Verifying the .NET Probe Installation	106
Configuring the .NET Probe	108
Determining the Version of the .NET Probe	109
Uninstalling the .NET Probe	109

About the Mercury Diagnostics Probe for .NET

The Mercury Diagnostics Probe for .NET (.NET Probe) provides a low-overhead capture solution that works with Mercury's Application Delivery and Application Management products. The .NET Probe captures events from a .NET application and sends the event metrics to the Diagnostics Server.

The .NET Probe uses runtime instrumentation to capture method latency information from specified applications. By default, the .NET Probe captures methods from the ASP.NET and ADO tiers, and from MSMQ for ASP.NET applications. The probe can be configured to capture custom business logic methods by adding instrumentation points to your application's **.points** file.

An *instrumented method* is any method defined in any of the points files, for which the probe was configured. *Initialization* refers to the registration of the probe with the Diagnostics Server. *.NET Probe initialization* occurs only when an instrumented method is encountered for the first time. For ASP.NET applications, the probe is started the first time that a page in the application is requested after the Web publishing service has been started.

Installing the .NET Probe

The following section provides detailed instructions for installing the .NET Probe.

About the Installation

You install the .NET Probe on the System Under Test (SUT).

For information about the recommended system configurations for hosting the .NET Probe, see “Requirements for the Mercury Diagnostics Probe for .NET Host” on page 17.

The default configuration of the .NET Probe has been set to monitor your application effectively, with very little impact on the performance of your application. For instructions on advanced .NET Probe configuration, see Chapter 15, “Advanced .NET Probe Configuration.”

Starting the Installation

You start the installer from the product CD or from the Diagnostics Downloads page in Mercury Business Availability Center.

To start the installation from the product CD:

Insert the Mercury Diagnostics installation CD for Windows into a CD-ROM drive. If the installer does not start automatically, double-click **setup.exe** in the root folder of the CD-ROM.

The installer displays the Mercury Diagnostics main installation menu.

Click **Mercury Diagnostics Probe for .NET** to initiate the installer for the .NET Probe.

To install the .NET Probe from a different location, locate the file **Mercury .NET Probe.msi** in the **Diagnostics_Probes** directory on the CD, copy it to the new location, and then run it.

Continue with “Running the Installation,” below.

To start the installation from the downloads page (for Mercury Business Availability Center users):

- 1** Select **Admin > Diagnostics** from the top menu in Mercury Business Availability Center, and click the **Downloads** tab.
- 2** On the Downloads page, click **.NET Probe For Windows NT/2000/XP** to download the .NET Probe installer.

Note: The probe installers are available in Mercury Business Availability Center only if you provided the path to the probe installers directory during the installation of the Diagnostics Server in Commander mode.

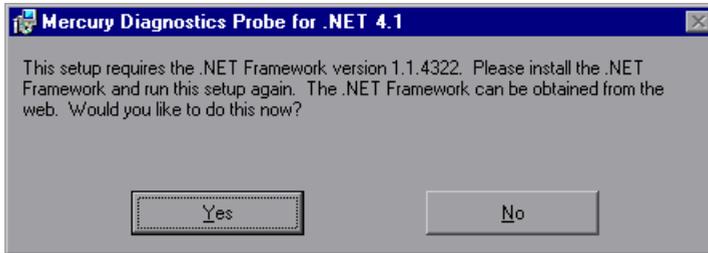
Continue with “Running the Installation,” below.

Running the Installation

After you have launched the installer, you are ready to begin the main installation procedure.

To install the .NET Probe on a Windows machine:

- 1 If the .NET Framework has not been installed on the probe host, the installer displays the following dialog box:



Click **Yes** to obtain the .NET Framework from the Web.

Clicking **No** ends the installation.

- 2 If you clicked **Yes**, the installer opens Microsoft's download page, where you can select the .NET Framework download.
- 3 After the .NET Framework installation restarts your computer, run the .NET Probe installer again.
- 4 The License Agreement is displayed.

Read the agreement, and select **I accept the terms of the license agreement** to accept the agreement. Click **Next** to continue.

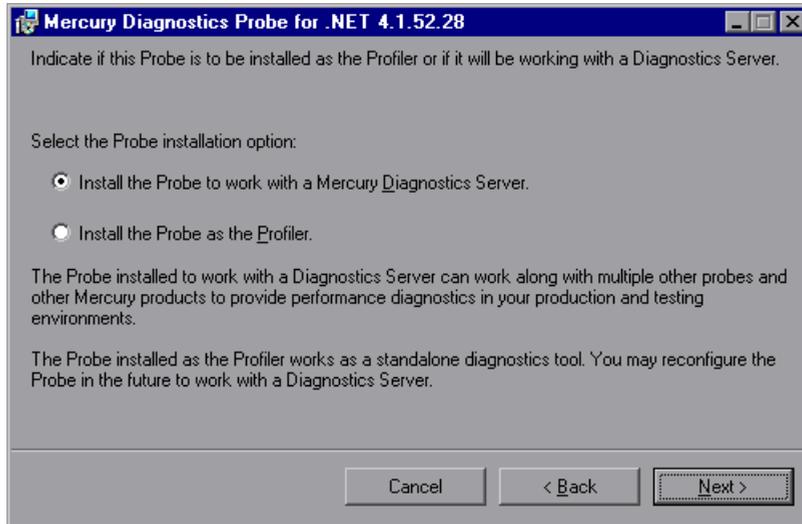
- 5 Provide the location where you want the probe installed.

Accept the default directory or select a different location either by typing in the path to the installation directory into the **Installation Directory Name** box, or by clicking **Browse** to navigate to the installation directory.

If you want to check the amount of available disk space on the drives of the host, click **Disk Cost**. Use this functionality to make sure that there is enough room for the probe installation.

Click **Next** to continue.

- 6 L:\Load_testing\Diagnostics\Install_Guide\dotnet_probe_install.fmSelect whether you want to install the probe as the Profiler only, or with a Mercury Diagnostics Server.



Make the selection that is appropriate for the environment where you will be using the probe:

- To use the probe with other probes and other Mercury products to provide performance metrics in your production or testing environment, select **Install the Probe to work with a Mercury Diagnostics Server**.
- To use the probe as a stand-alone diagnostics tool for a development environment, select **Install the Probe as the Profiler**.

Note: When the .NET Probe is installed to work with a Mercury Diagnostics Server, the Mercury Diagnostics Profiler for .NET is disabled by default. To enable it, see “Enabling the Mercury Diagnostics Probe for .NET” on page 271.

When the .NET Diagnostics Profiler is enabled, you can drill down from the Mercury Diagnostics screens into the screens of the .NET Diagnostics Profiler or you can access the Profiler in stand-alone mode.

When the probe is installed as the Profiler only, you can reconfigure the probe in the future to work with the Diagnostics Server.

Click **Next** to continue.

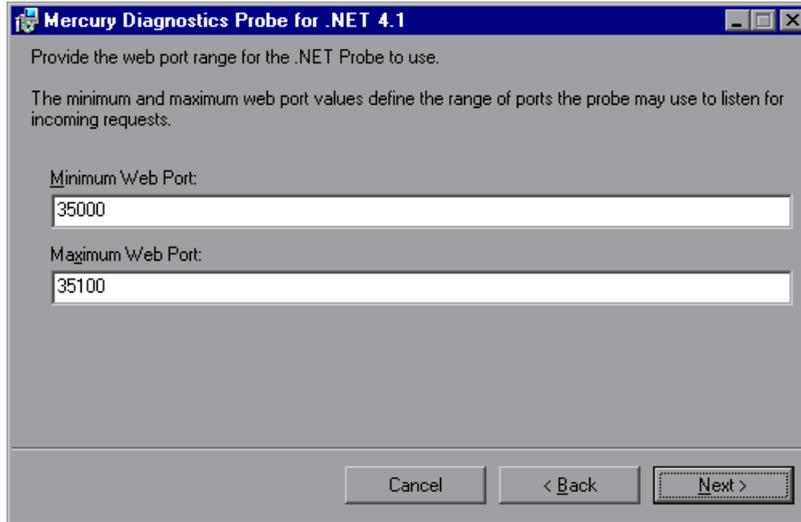
Next Step: At this stage, the installation procedure differs, depending on the environment in which you are installing the probe.

- ▶ If you are installing the probe as the Profiler only, continue with “Installing the Probe as a Profiler Only,” below.
 - ▶ If you are installing the probe to work with a Mercury Diagnostics Server, continue with “Installing the Probe to Work with a Mercury Diagnostics Server” on page 100.
-

Installing the Probe as a Profiler Only

If you are installing the probe to work as a Profiler only, continue with the following procedure.

- 1 Provide the Web port range for the .NET Probe to use.



Mercury Diagnostics Probe for .NET 4.1

Provide the web port range for the .NET Probe to use.

The minimum and maximum web port values define the range of ports the probe may use to listen for incoming requests.

Minimum Web Port:
35000

Maximum Web Port:
35100

Cancel < Back Next >

- **Web Port Min.** Type the lowest port number, in a range of ports on the probe host, that you want to assign to the probe.
- **Web Port Max.** Type the highest port number, in a range of ports on the probe host, that you want to assign to the probe.

Note: The default range is 100 ports, 35000 - 35100.

The upper and lower limits of the Web Port Range are defined by the **Web Port Min** and **Web Port Max** fields. The Web Port Range contains the ports that the probe can use.

When a probe is started, it attempts to find an unused port from within this range; starting from the lowest port number in the range and working its way up to the highest. Ports within the range may already be in use if another probe or application has previously claimed them.

The minimum size for the port range is equal to the maximum number of probes that will be concurrently running on the probe's host.

Considerations when setting the Web Port Range:

- If the probes are working with ASP.NET applications, it is recommended that you double the number of ports to account for ASP.NET's appdomain recycling.
- A port within the range does not have to be available, as long as there are enough available ports within the range described above.
- If you have a firewall between the probe and a component that will be communicating with the probe, you must open the firewall for the ports within the range. For this reason you may want to adjust the range to be just big enough.

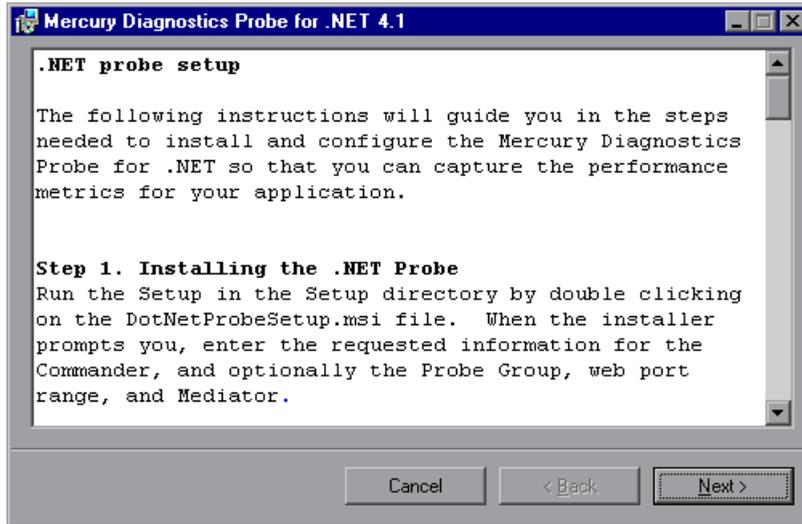
Click **Next** to continue.

- 2** The pre-installation summary screen opens. Click **Back** to make any changes. Click **Next** to start the .NET Probe installation.

The installation process begins, and a progress bar indicates how the installation is proceeding.

- 3 When the .NET Probe installation completes, instructions for configuring the probe are displayed in the installer window.

The instructions for configuring the probe can also be viewed in `<probe_install_dir>readme.rtf`. The .NET Probe configuration is discussed in greater detail in Chapter 15, “Advanced .NET Probe Configuration.”



Click **Next** when you are ready to proceed.

- 4 A confirmation message displays the results of the probe installation.

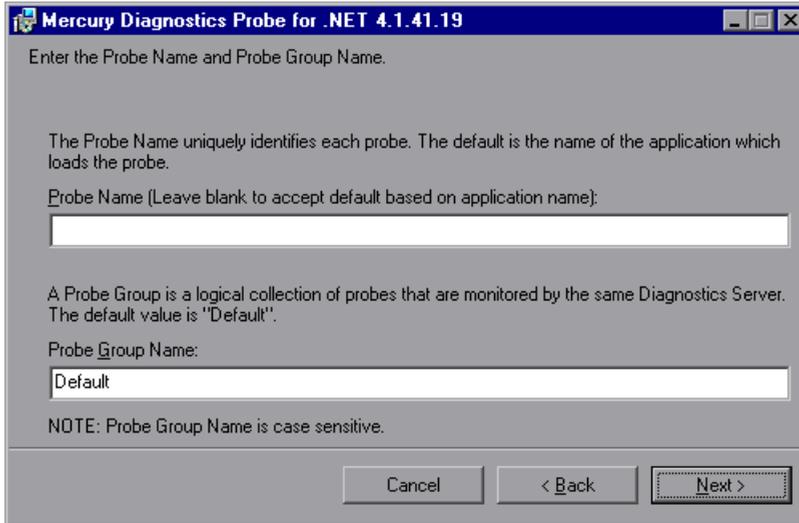
Click **Close** to exit the installer.

- 5 After you exit the installer, you must restart either the IIS or the Web publishing service before you can use the .NET Probe with ASP.NET applications.

Installing the Probe to Work with a Mercury Diagnostics Server

If you are installing the probe to work with a Mercury Diagnostics Server, continue with the following procedure.

- 1 Enter the probe name and probe group name.



The screenshot shows a Windows-style dialog box titled "Mercury Diagnostics Probe for .NET 4.1.41.19". The main text reads "Enter the Probe Name and Probe Group Name." Below this, there is explanatory text: "The Probe Name uniquely identifies each probe. The default is the name of the application which loads the probe." This is followed by a text input field labeled "Probe Name (Leave blank to accept default based on application name):". Below that, another explanatory text reads: "A Probe Group is a logical collection of probes that are monitored by the same Diagnostics Server. The default value is 'Default'." This is followed by a text input field labeled "Probe Group Name:" containing the text "Default". At the bottom, a note states: "NOTE: Probe Group Name is case sensitive." The dialog box has three buttons at the bottom: "Cancel", "< Back", and "Next >".

- **Probe Name.** The name that identifies the probe within Mercury Diagnostics. If you leave this field blank, the .NET Probe will auto-generate a probe name based on the application's domain name.

Note: It is recommended that you leave **Probe Name** blank and allow the probe to auto-generate the probe name. Read the following information carefully if you decide to enter your own probe name.

Considerations when entering a probe name:

- Valid characters that can appear in the probe name are: letters, digits, dashes, underscores, and periods.
- Assign a probe name that will help you recognize the application that the probe is monitoring, and the type of probe it is.

For example, the probe name for the first .NET Probe installed to monitor the application named PetWorld can be:

PetWorld_Dotnet_Probe1

- When you specify a probe name, all of the probes on the host are forced to use the same probe name. If you want to override the default names, use the following substitution macros to enhance the name at run time:

\$(MACHINENAME): Machine's host name

\$(APPDOMAIN): Application's domain name

\$(PID): Application's process ID

The default probe name auto-generated by the probe when the probe name field is left blank is equivalent to specifying “\$(APPDOMAIN).NET”.

To obtain behavior compatible with version 3.3.x of the probe, specify “\$(MACHINENAME).\$(APPDOMAIN).NET”.

- **Probe Group:** Probe groups are logical groupings of probes that report to the same Diagnostics Server. The performance metrics for a probe group are tracked, and can be displayed on many of the Diagnostics screens.

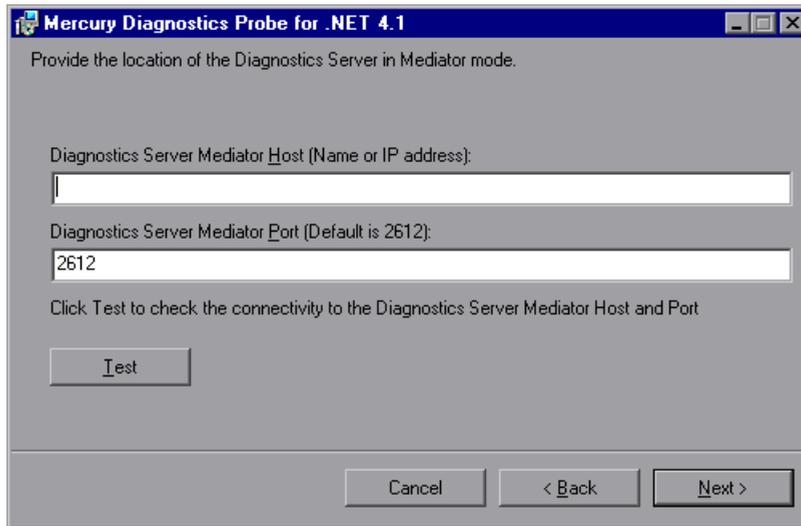
You may want to assign all of the probes for a particular enterprise application to a single probe group so that you can monitor the performance of the group as well as the individual probes.

The default value for the probe group name is Default.

Note: The probe group name is case-sensitive.

Click **Next** to continue.

- 2 Provide the information needed to enable the .NET Probe to communicate with the Diagnostics Server in Mediator mode.



- a In the **Diagnostics Server Mediator Host** box, type the host name or IP address of the host for the Diagnostics Server in Mediator mode.

You should specify the fully qualified host name, not just the simple host name. In a mixed OS environment, where UNIX is one of the systems, this is essential for proper network routing.

- b In the **Diagnostics Server Mediator Port** box, type the port number where the Diagnostics Server Mediator port is listening for probe communication. The default port number is 2612. If you have changed the port since the Diagnostics Server was installed, you should specify that port number here instead of the default.
- c If you want to perform a connectivity check to make sure that the Diagnostics Server is running and accessible from the installation host, click **Test**.

The connectivity check lets you know right away if you have made an error in the information that you provided about the Diagnostics Server in Mediator mode, or if there is a connection problem between the

Diagnostics Server's host and the probe's host. If the connection to the Diagnostics Server Mediator host cannot be resolved, an error message is displayed.

- d Click **Next** to proceed.
- 3 Provide the Web port range for the .NET Probe to use.

Mercury Diagnostics Probe for .NET 4.1

Provide the web port range for the .NET Probe to use.

The minimum and maximum web port values define the range of ports the probe may use to listen for incoming requests.

Minimum Web Port:
35000

Maximum Web Port:
35100

Cancel < Back Next >

- **Web Port Min.** Type the lowest port number, in a range of ports on the probe host, that you want to assign to the probe.
- **Web Port Max.** Type the highest port number, in a range of ports on the probe host, that you want to assign to the probe.

Note: The default range is 100 ports, 35000 - 35100.

The upper and lower limits of the Web Port Range are defined by the **Web Port Min** and **Web Port Max** fields. The Web Port Range contains the ports that the probe can use.

When a probe is started, it attempts to find an unused port from within this range; starting from the lowest port number in the range and working its way up to the highest. Ports within the range may already be in use if another probe or application has previously claimed them.

The minimum size for the port range is equal to the maximum number of probes that will be concurrently running on the probe's host.

Considerations when setting the Web Port Range:

- If the probes are working with ASP.NET applications, it is recommended that you double the number of ports to account for ASP.NET's appdomain recycling.
- A port within the range does not have to be available, as long as there are enough available ports within the range described above.
- If you have a firewall between the probe and a component that will be communicating with the probe, you must open the firewall for the ports within the range. For this reason you may want to adjust the range to be just big enough.

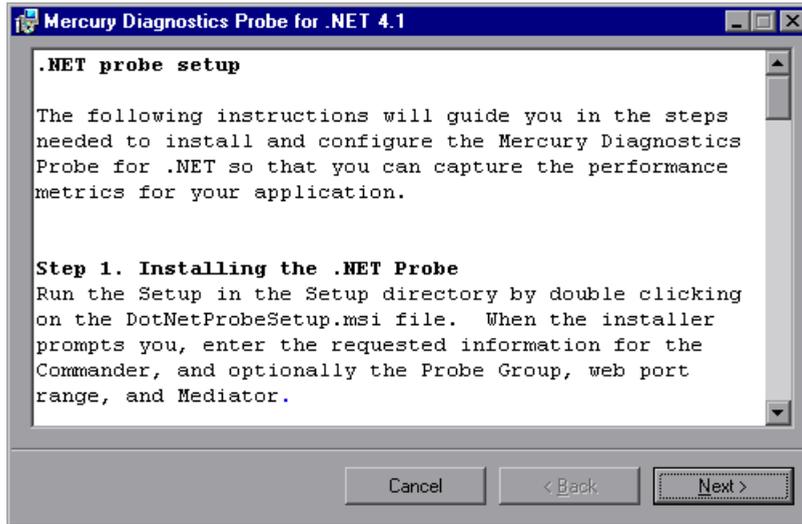
Click **Next** to continue.

- 4** The pre-installation summary screen opens. Click **Back** to make any changes. Click **Next** to start the .NET Probe installation.

The installation process begins, and a progress bar indicates how the installation is proceeding.

- 5 When the .NET Probe installation completes, instructions for configuring the probe are displayed in the installer window.

The instructions for configuring the probe can also be viewed in `<probe_install_dir>readme.rtf`. The .NET Probe configuration is discussed in greater detail in Chapter 15, “Advanced .NET Probe Configuration.”



Click **Next** when you are ready to proceed.

- 6 A confirmation message displays the results of the probe installation. Click **Close** to exit the installer.
- 7 After you exit the installer, you must restart either the IIS or the Web publishing service before you can use the .NET Probe with ASP.NET applications.
- 8 You can verify the installation as described in “Verifying the .NET Probe Installation,” below.

Note: The default configuration of the .NET Probe has been set to monitor your application effectively, with very little impact on the performance of your application. For more information, see “Configuring the .NET Probe” on page 108.

Verifying the .NET Probe Installation

Important! This section applies only if you installed the probe to work with a Mercury Diagnostics Server.

Use the System Health Monitor to verify the installation of the .NET Probe. For instructions on how to use the System Health Monitor, see the *Mercury Diagnostics User's Guide*.

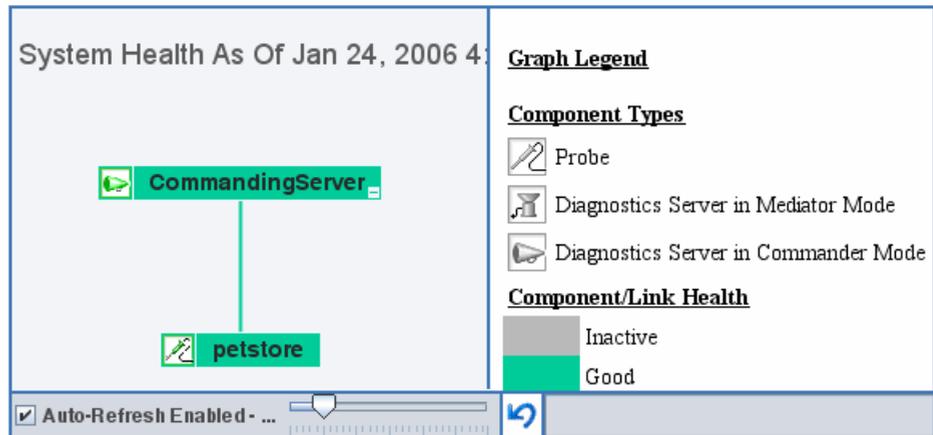
If you have been following the recommended installation sequence, after you have installed the .NET Probe and restarted the instrumented application, you will be able to verify the following:

- ▶ The Diagnostics Server in Commander mode was successfully installed.
- ▶ Where relevant, additional Diagnostics Servers in Mediator mode were successfully installed and are communicating with the Diagnostics Server in Commander mode.
- ▶ The .NET Probe was successfully installed and is communicating with the Diagnostics Server.

Note: The .NET Probe will not be displayed in System Health when first installed because the probe does not register with the Diagnostics Server until it is started. The probe is started and is registered with the Diagnostics Server when the instrumented application is run. For ASP.NET applications, this happens the first time that a page is requested for the instrumented application.

The default logging level for the .NET Probe has been set to **info** so that if the probe is not displayed in System Health when you believe that it should be, you can check the log file to help troubleshoot the problem. If you do not see a log file you can check the Windows Event Viewer.

Depending on your deployment, the new .NET Probe is shown on the System Health Monitor to be reporting to either the Diagnostics Server in Commander mode or the Diagnostics Server in Mediator mode.



Configuring the .NET Probe

The default configuration of the .NET Probe allows you to begin capturing performance metrics for your applications. Moreover, the probe's configuration can be customized to suit the configuration of your environment and the performance issues that you would like to diagnose.

The installer configures your ASP.NET applications and the .NET Probe so that they can work together to capture the basic workload of your applications. It is possible that one or more of your ASP.NET applications has been deployed in a manner that prevents the installer from detecting it, or you may want to enhance the standard instrumentation to capture the performance metrics for the custom classes in your application.

The .NET Probe is installed with a tool that helps you discover the methods and classes in an application. For information about the Reflector tool, see “Discovering the Classes and Methods in an Application” on page 265.

For instructions on advanced .NET Probe configuration, see Chapter 15, “Advanced .NET Probe Configuration.”

Note: When the .NET Probe is installed to work with a Mercury Diagnostics Server, the Mercury Diagnostics Profiler for .NET is disabled by default. To enable it, see “Enabling the Mercury Diagnostics Probe for .NET” on page 271.

Determining the Version of the .NET Probe

When you request support, it is useful to know the version of the Diagnostics components that you have installed.

To determine the version of the .NET Probe:

Right-click the file <probe_install_dir>\bin\Mercury.Profiler.dll, and display the component version information by selecting **Properties** from the menu.

Uninstalling the .NET Probe

To uninstall the .NET Probe, use the **Add/Remove Programs** utility in Windows.

6

Installing the Mercury Diagnostics Probe for J2EE

This chapter provides instructions for installing a Mercury Diagnostics Probe for J2EE (J2EE Probe) on Windows machines, UNIX machines, and z/OS mainframe machines. Instructions have also been provided for using a generic installer to install the J2EE Probe on other platforms.

This chapter describes:	On page:
About Installing and Configuring the J2EE Probe	112
Installing the J2EE Probe on a Windows Machine	113
Installing the J2EE Probe on a UNIX Machine	130
Installing the J2EE Probe on a z/OS Mainframe	146
Installing the J2EE Probe Using the Generic Installer	149
Silent Installation of the J2EE Probe	150

About Installing and Configuring the J2EE Probe

Before you can use the J2EE Probe to monitor your application, you need to perform the following operations:

1 Install the J2EE Probe.

The J2EE Probe is installed on the machine hosting the application that you wish to monitor. For more information about installing the J2EE Probe, see one of the following sections, depending on your operating system and method of installation.

- ▶ “Installing the J2EE Probe on a Windows Machine” on page 113
- ▶ “Installing the J2EE Probe on a UNIX Machine” on page 130
- ▶ “Installing the J2EE Probe on a z/OS Mainframe” on page 146
- ▶ “Installing the J2EE Probe Using the Generic Installer” on page 149
- ▶ “Silent Installation of the J2EE Probe” on page 150

2 Configure the J2EE Probe.

The probe is automatically configured during the installation process. In certain situations however, you may be required to change these default settings or manually configure the probe. For more information, see “Configuring the Probe Manually” on page 154.

3 Configure the application server.

In order to allow the probe to monitor your application, you need to instrument the JRE and manually modify the application startup script. For more information, see “Configuring the Application Server” on page 157.

Note: For information about the recommended system configurations for hosting the J2EE Probe, see “Requirements for the Mercury Diagnostics Probe for J2EE Host” on page 16.

Installing the J2EE Probe on a Windows Machine

The following steps provide detailed instructions for installing the J2EE Probe on a Windows machine. These instructions also apply when you are installing the probe on a UNIX machine using the graphical installer.

Launching the Installer

The installer can be launched from the product CD or from the Diagnostics Downloads page in Mercury Business Availability Center.

To launch the installer from the product CD:

Insert the Mercury Diagnostics installation CD for Windows into a CD-ROM drive. If the installer does not launch automatically, double-click **setup.exe** in the root folder of the CD.

The installer displays the Mercury Diagnostics main installation menu.

Click **Mercury Diagnostics Probe for J2EE** to initiate the installer for the J2EE Probe.

To install the probe from a different location, locate the executable file **J2EEProbeSetup_win.exe** file in the **Diagnostics_Probes** directory on the CD, copy it to the new location, and then run it.

Continue with “Running the Installation,” below.

To launch the installer from the downloads page, for Mercury Business Availability Center users:

- 1** Select **Admin > Diagnostics** from the top menu in Mercury Business Availability Center, and click the **Downloads** tab.
- 2** On the Downloads page, click **J2EE Probe For Windows NT/2000/XP**.

Note: The probe installers are available in Mercury Business Availability Center only if you provided the path to the probe installers directory during the installation of the Diagnostics Server in Commander mode.

Continue with “Running the Installation,” below.

Running the Installation

After you have launched the installer, the software license agreement opens and you are ready to run the installation.

To install the J2EE Probe on a Windows machine:

1 Accept the software license agreement.

Read the agreement and select **I accept the terms of the license agreement**.

Click **Next** to proceed.

2 Specify the location where you want to install the probe.

Accept the default directory or select a different location either by typing the path to the installation directory into the **Installation Directory Name** box, or by clicking **Browse** to navigate to the installation directory.

If there is a pre-existing installation of the probe on the host machine, you can upgrade the existing probe or install the probe to a different installation directory.

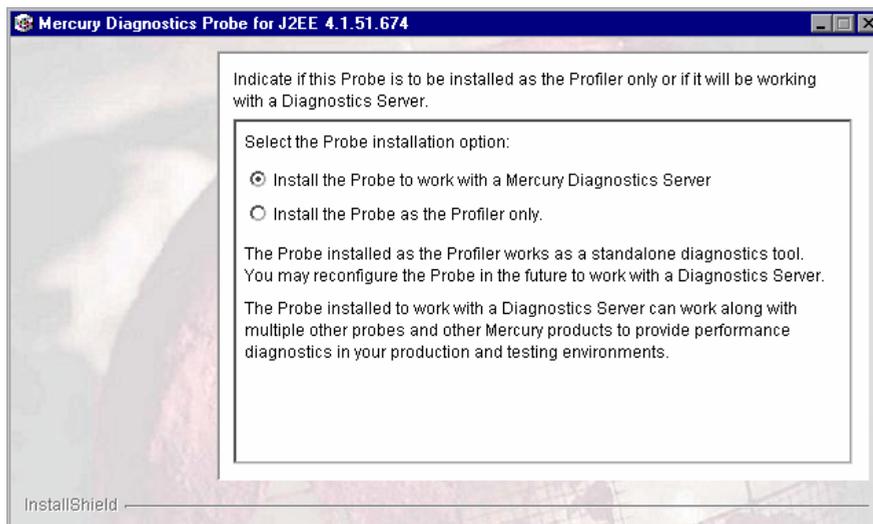
- To upgrade the existing probe installation, ensure that the installation directory of the existing probe has been specified in the **Installation Directory Name** box.

When the installer runs, it will back up any property settings of the existing J2EE Probe installation to a directory called **etc.old**.

- To install the probe in an installation directory that is different from the existing probe's directory, ensure that the installation directory of the existing probe has not been specified in the **Installation Directory Name** box.

Click **Next** to proceed.

3 Select whether you want to install the probe with a Mercury Diagnostics Server, or as the Profiler only.



- ▶ If you want to use the probe with other probes and other Mercury products to provide performance diagnostics in your production or testing environment, select **Install the Probe to work with a Mercury Diagnostics Server**.
- ▶ If you want to use the probe as a stand-alone diagnostics tool for a development environment, select **Install the Probe as the Profiler only**.

Note: When the probe is installed to work with a Mercury Diagnostics Server, the configuration is set such that the J2EE Diagnostics Profiler is available as well. You can drill down from the Mercury Diagnostics screens into the screens of the J2EE Diagnostics Profiler or you can access the Profiler in stand-alone mode.

When the probe is installed as a Profiler only, you can reconfigure the probe in the future to work with the Diagnostics Server.

Click **Next** to proceed.

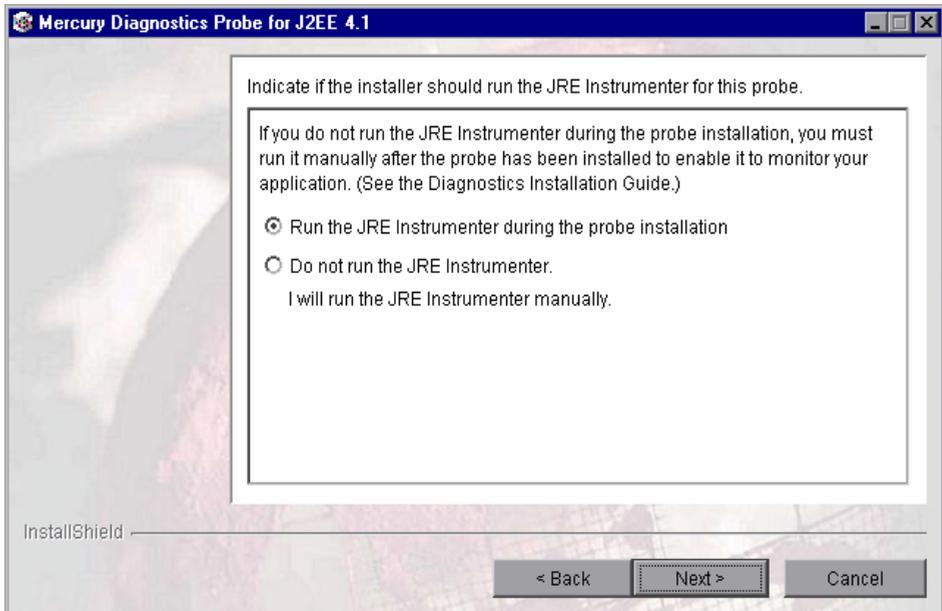
Next Step: At this stage, the installation procedure differs, depending on the environment in which you are installing the probe.

- ▶ If you are installing the probe as a Profiler only, continue with “Installing the Probe as a Profiler Only,” below.
 - ▶ If you are installing the probe to work with a Mercury Diagnostics Server, continue with “Installing the Probe to Work with a Mercury Diagnostics Server” on page 120.
-

Installing the Probe as a Profiler Only

If you are installing the probe to work as a Profiler only, continue with the following procedure.

- 1 Indicate whether or not you want to run the JRE Instrumenter during the installation of the Probe.**



You can run the instrumenter now, or you can skip this step by selecting **Do not run the JRE Instrumenter**. If you choose to skip this step now, you must run the JRE Instrumenter manually before you can use the Probe. For more information, see “Running the JRE Instrumenter” on page 165.

Note: If you are installing this Probe to work with SAP NetWeaver, select **Do not run the JRE Instrumenter**. You will configure SAP NetWeaver manually. See “Configuring the SAP NetWeaver Application Server” on page 204.

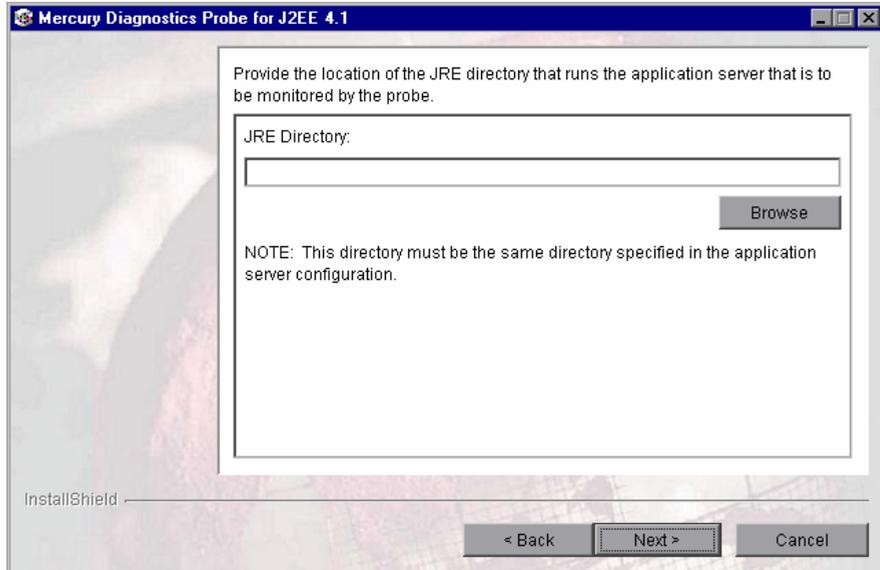
If you are installing this Probe to work with WebSphere in an Integrated Development Environment (IDE), select **Do not run the JRE Instrumenter**. You will run the JRE Instrumenter manually. See “Running the JRE Instrumenter for WebSphere IDE” on page 187.

Click **Next** to proceed.

Next Step: If you chose to run the JRE instrumenter now, continue below. If you chose not to run the JRE instrumenter now, skip to step 3.

2 Enter the JRE directory.

If you indicated that you want to run the JRE Instrumenter during the installation of the Probe, enter the JRE directory.



In the **JRE Directory** box, type the path to the JRE bin folder where the JRE used by the monitored application server can be found. For example, if you have installed WebLogic 8.1 in your **D:\bea** directory, the **java.exe** file can be found in the bin folder of the JRE directory, **D:\bea\jdk142_05\JRE**.

Click **Next** to proceed.

3 Confirm the pre-installation summary information.

The pre-installation summary screen opens. Review the summary information. Click **Back** to make any changes. Click **Next** to proceed.

The installation process begins, and a progress bar indicates how the installation is proceeding.

4 After installation, a status message notifies you that the installation was successfully completed.

Click **Finish** to close the installation wizard.

5 Configure the Application Server

Before you can use the probe to monitor your application, you need to instrument the JRE and manually modify the application startup script.

- a** If you chose to skip running the JRE Instrumenter during this installation, you need to run it manually before you can use the probe. For more information, see “Running the JRE Instrumenter” on page 165.
- b** Once you have installed the probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the probe is started together with the monitored application. For detailed instructions, see Chapter 9, “Configuring the Application Servers.”

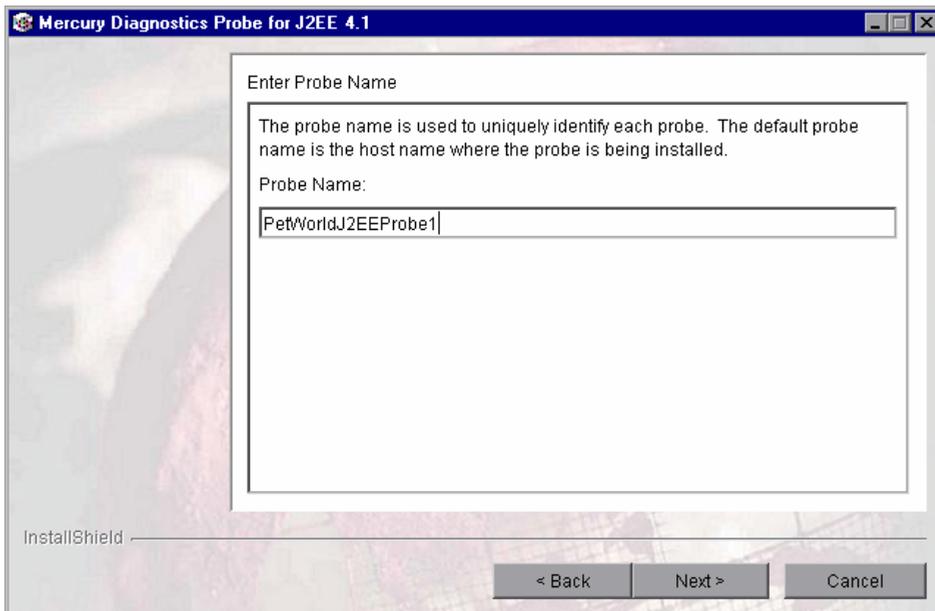
Note: The probe is automatically configured during the installation process. In certain situations however, you may be required to change these default settings or to manually configure the probe. For more information, see “Configuring the Probe Manually” on page 154.

For situations that require more advanced configurations of the probe and application server, see Chapter 16, “Advanced J2EE Probe and Application Server Configuration.”

Installing the Probe to Work with a Mercury Diagnostics Server

If you are installing the probe to work with a Mercury Diagnostics Server, continue with the following procedure.

1 Assign a name for the probe.



Type a unique name that will be used to identify the probe within Mercury Diagnostics. The following characters can be used in the name: - , _ , and all alphanumeric characters.

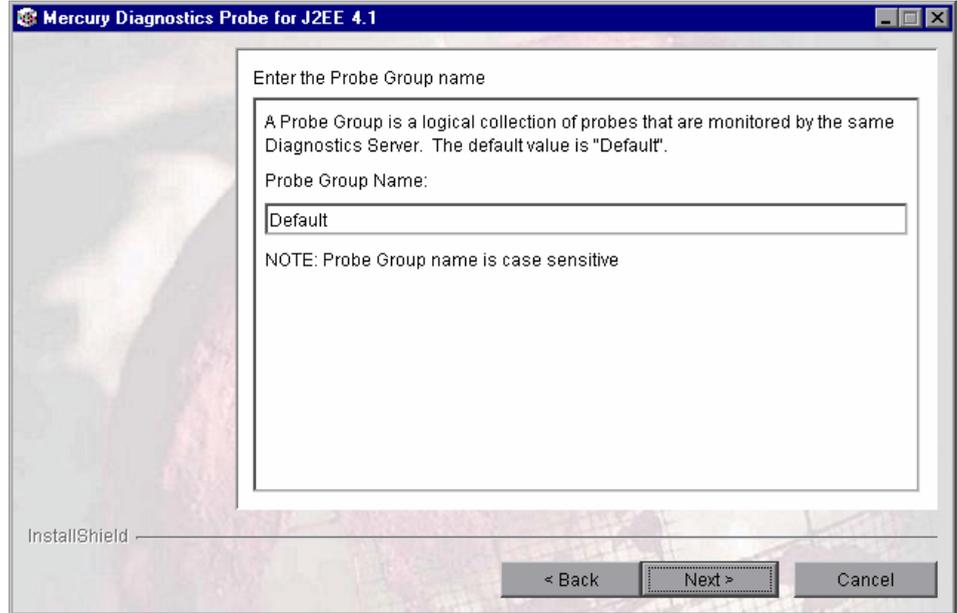
When assigning a name to a probe, choose a name that will help you to recognize the application that the probe is monitoring, and the type of the probe (in this case, a J2EE Probe).

For example, the probe name for the first J2EE Probe installed, that will be monitoring an application named PetWorld, could be:

PetWorldJ2EEProbe1

Click **Next** to proceed.

2 Assign a name for the probe group.



Probe groups are logical collections of probes that are monitored by the same Diagnostics Server. The performance metrics for a probe group are tracked and can be displayed on many of the Diagnostics screens.

For example, you may want to assign all of the probes for a particular enterprise application to a single probe group so that you can monitor the performance of the group as well as the individual probes.

Type a name for the probe group, or accept the default.

Note: The probe group name is case-sensitive.

Click **Next** to proceed.

3 Enter the details for the Diagnostics Server in Mediator mode.

Enter the necessary information to enable communication with the Diagnostics Server in Mediator mode.

Mercury Diagnostics Probe for J2EE 4.1

Provide the location of the Diagnostics Server in Mediator mode.

Diagnostics Server Mediator Host (Name or IP address):
pen

Diagnostics Server Mediator Port:
2612

The default value is 2612.

Check the connectivity to the Diagnostics Server Mediator Host and Port.

InstallShield

< Back Next > Cancel

Note: If there is only one Diagnostics Server in the Diagnostics deployment where the probe will run, enter the Diagnostics Server's host name and event port information here.

When there is more than one Diagnostics Server in your deployment, enter the information for the Diagnostics Server in Mediator mode that is to receive the events from the probe.

- a** In the **Diagnostics Server Mediator Host** box, type the host name or IP address of the host for the Diagnostics Server in Mediator mode.

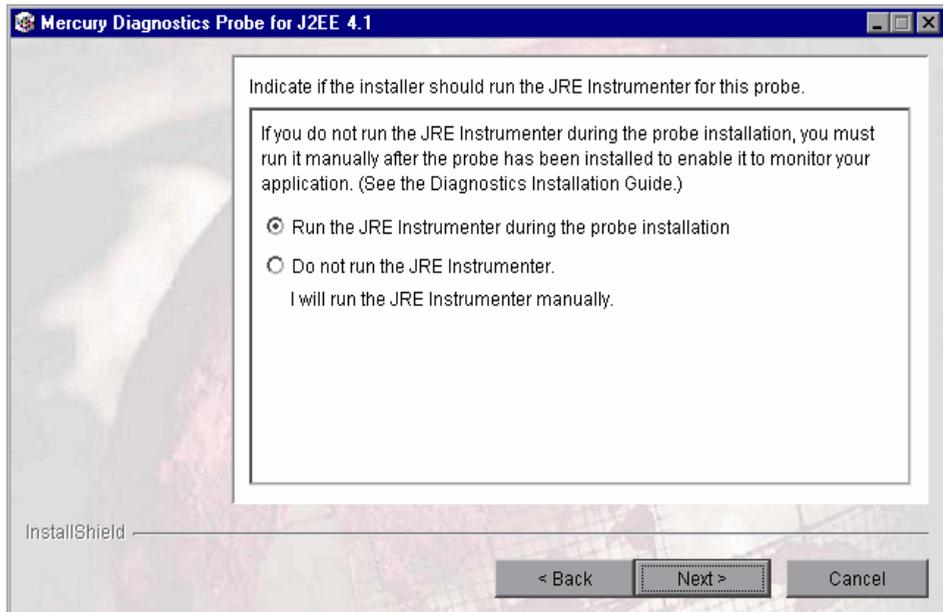
You should specify the fully qualified host name, not just the simple host name. In a mixed OS environment, where UNIX is one of the systems, this is essential for proper network routing.

- b** In the **Diagnostics Server Mediator Port** box, type the port number where the Diagnostics Server Mediator is listening for probe communication. The default port number is 2612. If you have changed the port since the Diagnostics Server was installed, you should specify that port number here instead of the default.
- c** If you want to make sure that the specified Diagnostics Server is running and accessible from the installation host, select **Check the connectivity to the Diagnostics Server Mediator Host and Port**.

Click **Next** to proceed.

If you selected **Check the connectivity to the Diagnostics Server Mediator Host and Port** and connectivity problems were encountered, the installer provides you with the results of the connectivity check. If you do not want to address these problems at this stage, clear the check box, proceed with the installation, and address the problem later.

4 Indicate whether or not you want to run the JRE Instrumenter during the installation of the Probe.



You can run the instrumenter now, or you can skip this step by selecting **Do not run the JRE Instrumenter**. If you choose to skip this step now, you must run the JRE Instrumenter manually before you can use the Probe. For more information, see “Running the JRE Instrumenter” on page 165.

Note: If you are installing this Probe to work with SAP NetWeaver, select **Do not run the JRE Instrumenter**. You will configure SAP NetWeaver manually. See “Configuring the SAP NetWeaver Application Server” on page 204.

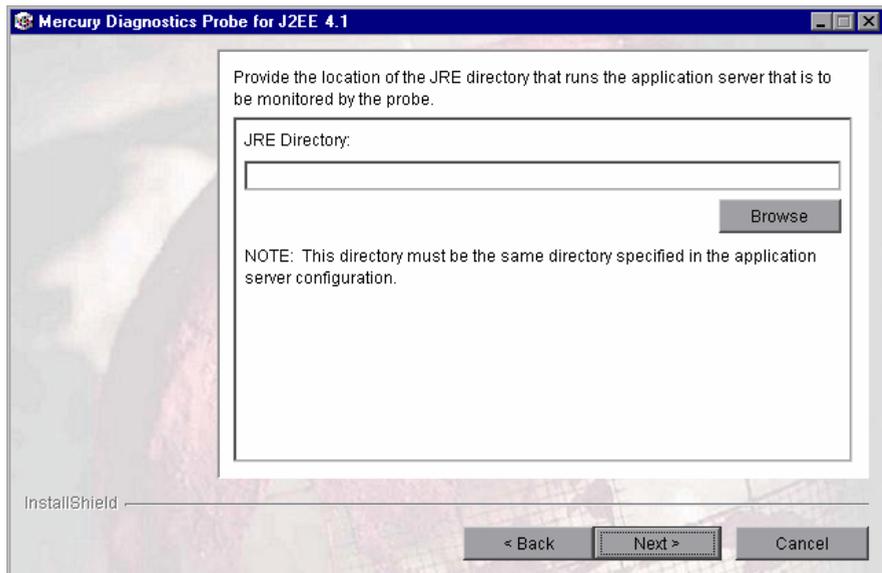
If you are installing this Probe to work with WebSphere in an Integrated Development Environment (IDE), select **Do not run the JRE Instrumenter**. You will run the JRE Instrumenter manually. See “Running the JRE Instrumenter for WebSphere IDE” on page 187.

Click **Next** to proceed.

Next Step: If you chose to run the JRE instrumenter now, continue below. If you chose not to run the JRE instrumenter now, skip to step 6.

5 Enter the JRE directory.

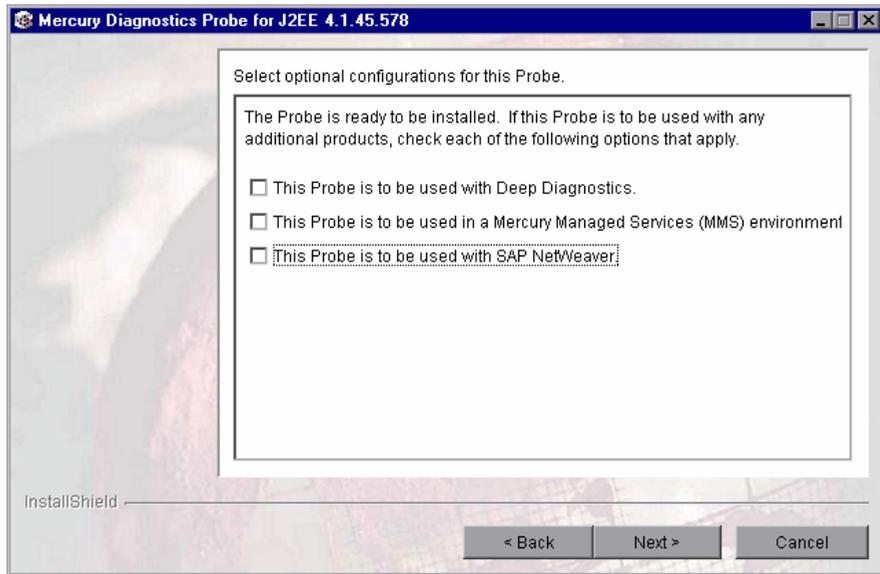
If you indicated that you want to run the JRE Instrumenter during the installation of the Probe, enter the JRE directory.



In the **JRE Directory** box, type the path to the JRE bin folder where the JRE used by the monitored application server can be found. For example, if you have installed WebLogic 8.1 in your **D:\bea** directory, the **java.exe** file can be found in the bin folder of the JRE directory, **D:\bea\jdk142_05\JRE**.

Click **Next** to proceed.

6 Select the configuration options if the probe is to be used with additional products.

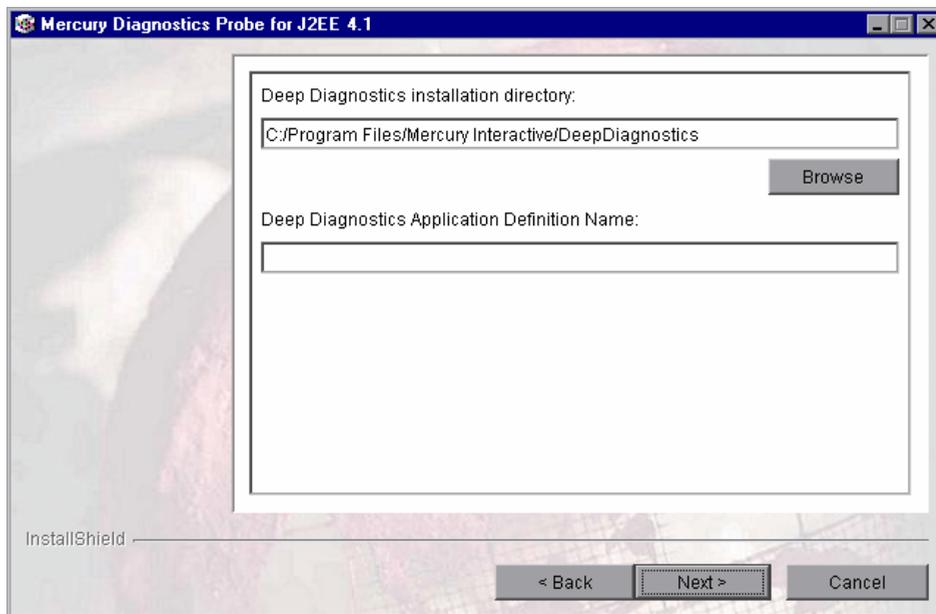


Indicate whether you will be using the probe with Deep Diagnostics, whether you will be using the probe within a Mercury Managed Services (MMS) environment, and whether you will be using the probe with a SAP NetWeaver Web application server.

Click **Next** to proceed.

7 Enter the Deep Diagnostics information.

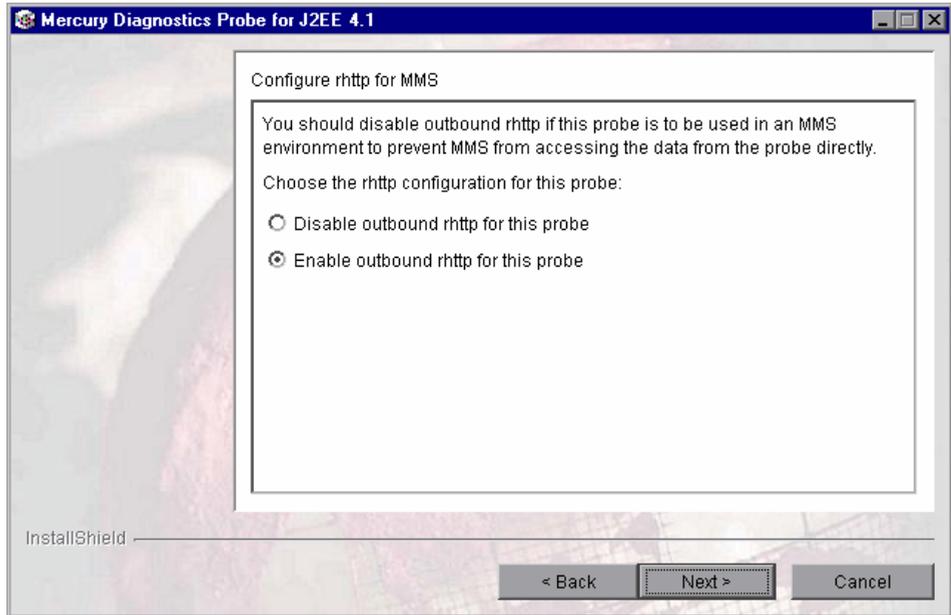
If you indicated that you want to use this probe with Deep Diagnostics, enter the Deep Diagnostics information.



- a** In the **Deep Diagnostics installation directory** box, type the location where Deep Diagnostics was installed, or click **Browse** to navigate to the installation location.
- b** Type the Deep Diagnostics Application Definition name. The Deep Diagnostics Application Definition name must exactly match the name specified for the Application Definition in Deep Diagnostics. For more information, refer to the *Mercury Deep Diagnostics for J2EE Installation Guide*.
- c** Click **Next** to proceed.

8 Configure the Mercury Managed Services environment.

If you indicated that the probe is to be used in a Mercury Managed Services environment (see step 6), you are given the option to disable outbound rhttp (reverse http).



The Diagnostics Server uses rhttp to query metrics from the J2EE Probe. You should disable the probe's outgoing rhttp to prevent Mercury Managed Services from accessing the data directly from the probe.

Select the appropriate option and click **Next**.

9 Confirm the pre-installation summary information.

The pre-installation summary screen opens. Review the summary information. Click **Back** to make any changes.

Click **Next** to proceed.

The installation process begins, and a progress bar indicates how the installation is proceeding.

10 Close the installation wizard.

After installation, a status message notifies you that the installation was successfully completed.

Click **Finish** to close the installation wizard.

11 Configure the Application Server

Before you can use the probe to monitor your application, you need to instrument the JRE and manually modify the application startup script.

- a** If you chose to skip running the JRE Instrumenter during this installation, you need to run it manually before you can use the probe. For more information, see “Running the JRE Instrumenter” on page 165.
- b** Once you have installed the probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the probe is started together with the monitored application. For detailed instructions, see Chapter 9, “Configuring the Application Servers.”

Note: The probe is automatically configured during the installation process. In certain situations however, you may be required to change these default settings or to manually configure the probe. For more information, see “Configuring the Probe Manually” on page 154

For situations that require more advanced configurations of the probe and application server, see Chapter 16, “Advanced J2EE Probe and Application Server Configuration.”

12 Verify the probe installation.

If you are installing this probe to work with a Mercury Diagnostics Server, you can verify the probe installation as described in “Verifying the J2EE Probe Installation” on page 158.

Installing the J2EE Probe on a UNIX Machine

J2EE Probe installers have been provided for several UNIX platforms. The following instructions provide you with the steps necessary to install the J2EE Probe in most UNIX environments using either a graphical installation or a console mode installation.

The installer screens that are displayed in a graphical installation are the same as those documented for the Windows installer in “Installing the J2EE Probe on a Windows Machine” on page 113.

In some instances, you may not be able to use the regular UNIX installers. In these cases, you should use the Generic installer as described in “Installing the J2EE Probe Using the Generic Installer” on page 149.

The instructions and screen shots that follow are for a probe installation on a Solaris machine. These same instructions should apply for the other certified UNIX platforms.

Launching the Installer

The Installer can be launched from the product CD or from the Diagnostics Downloads page in Mercury Business Availability Center.

To launch the installer from the product CD:

Copy the installer from the **Mercury Diagnostics installation CD for <UNIX version>\Diagnostics_Probes** directory to the machine where the probe is to be installed.

Continue with “Running the Installation,” below.

To launch the installer from the downloads page (for Mercury Business Availability Center users):

- 1 Select **Admin > Diagnostics** from the top menu in Mercury Business Availability Center, and click the **Downloads** tab.
- 2 On the Downloads page, click the link to the installer that is appropriate for your environment and save it to the machine where the probe is to be installed.

Note: The probe installers are available in Mercury Business Availability Center only if you provided the path to the probe installers directory during the installation of the Diagnostics Server in Commander mode.

Continue with “Running the Installation,” below.

Running the Installation

After you have copied the installer to the machine where the probe is to be installed, you are ready to run the installation.

Note: The following instructions assume an understanding of UNIX console screens and commands. For more information about UNIX screens and commands, see Appendix H, “Using UNIX Commands.”

To install the J2EE Probe on a Unix machine:

1 Run the installer.

Where necessary, change the mode of the installer file to make it executable.

- ▶ To run the installer in console mode, enter the following at the UNIX command prompt:

```
./<installer> -console
```

The installer displays the installation prompts in console mode as shown in the following step.

- ▶ To run the installer in the graphical mode enter the following at the UNIX command prompt:

```
./<installer>
```

The installer displays the same screens that are displayed for the Windows installer, as shown in “Installing the J2EE Probe on a Windows Machine” on page 113.

2 Accept the software agreement.

The software license agreement is displayed.

Read the agreement. As you read, you can press **ENTER** to move to the next page of text, or type **q** to jump to the end of the license agreement.

Accept the terms of the agreement, and select **Next** to continue with the installation.

3 Specify the location where you want to install the probe.

At the **Directory Name** prompt, accept the default installation location shown in brackets, or enter the path to a different location.

If there is a pre-existing installation of the probe on the host machine, you can upgrade the existing probe or install the probe to a different installation directory.

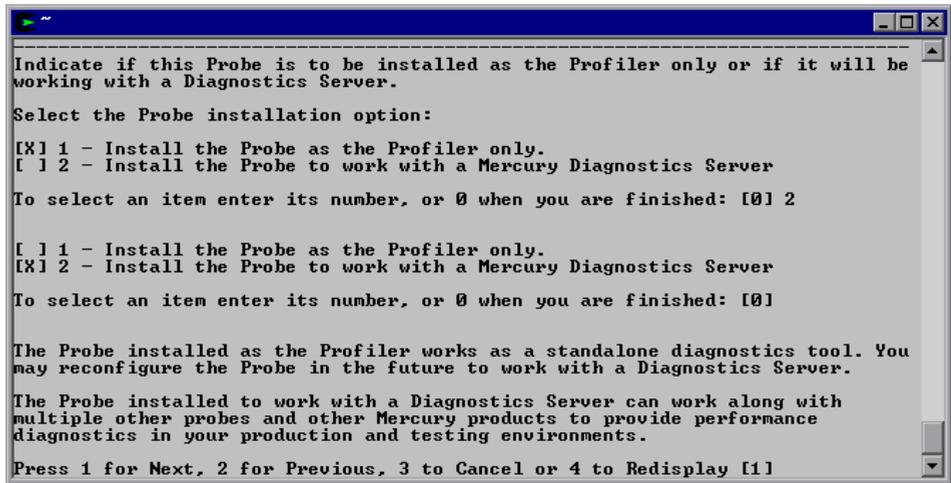
- ▶ To upgrade the existing probe installation, specify the installation directory of the existing probe at the **Directory Name** prompt.

When the installer runs, it backs up any property settings for the existing J2EE Probe installation to a directory called **etc.old**.

- ▶ To install the probe in an installation directory that is different from the existing probe’s directory, specify an installation directory other than that of the existing probe.

Select **Next** to continue with the installation.

4 Specify whether you want to install the probe as the Profiler only, or with a Mercury Diagnostics Server.



The probe can be configured to work in one of the following ways:

- To use the probe with other probes and other Mercury products to provide performance diagnostics in your production or testing environment, select **Install the Probe to work with a Mercury Diagnostics Server**.
- To use the probe as a stand-alone diagnostics tool for a development environment, select **Install the Probe as the Profiler only**.

Select **Next** to continue with the installation.

Note: When the probe is installed to work with a Mercury Diagnostics Server, it is configured so that the J2EE Diagnostics Profiler is available as well. You can drill down from the Mercury Diagnostics screens into the screens of the J2EE Diagnostics Profiler or you can access the Profiler in stand-alone mode.

A probe that was installed in Profiler-only mode can be reconfigured to work with the Diagnostics Server if this becomes necessary.

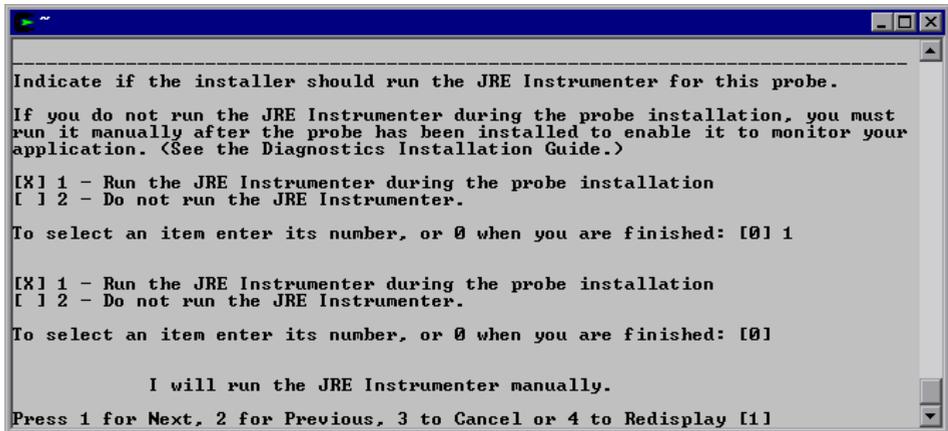
Next Step: At this stage, the installation procedure branches, depending on the option that you selected for installing the probe.

- ▶ If you are installing the probe as a Profiler only, continue with “Installing the Probe as a Profiler Only,” below.
 - ▶ If you are installing the probe to work with a Mercury Diagnostics Server, continue with “Installing the Probe to Work with a Mercury Diagnostics Server” on page 137.
-

Installing the Probe as a Profiler Only

If you are installing the probe to work as a Profiler only, continue with the following procedure.

- 1** Indicate whether or not you want to run the JRE Instrumenter during the installation of the probe.



- ▶ To run the instrumenter now, select **Run the JRE Instrumenter during the probe installation**.

- ▶ To skip this step, select **Do not run the JRE Instrumenter**. If you choose to skip this step now, you must run the JRE Instrumenter manually before you can use the probe. For instructions, see Chapter 8, “Running the JRE Instrumenter.”

Note: If you are installing this probe to work with SAP NetWeaver, skip the JRE Instrumenter. You will configure SAP NetWeaver manually. See “Configuring the SAP NetWeaver Application Server” on page 204.

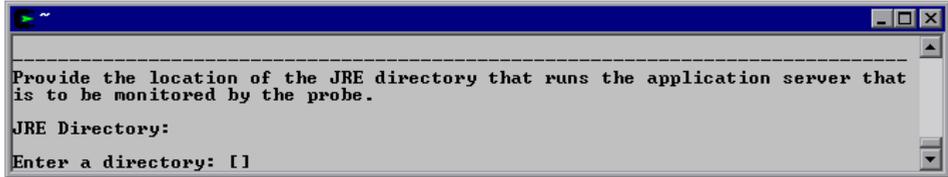
If you are installing this probe to work with WebSphere in an Integrated Development Environment (IDE), skip the JRE Instrumenter. You will run the JRE Instrumenter manually. See “Running the JRE Instrumenter for WebSphere IDE” on page 187.

Select **Next** to continue with the installation.

Next Step: At this stage, the installation procedure branches, depending on the option that you selected for running the JRE Instrumenter.

- ▶ If the installer is running the JRE Instrumenter, the installation process continues as documented in the next step.
 - ▶ Otherwise, the installation process skips the next step, and proceeds with the step following the next.
-

2 Enter the JRE directory.



Type the path to the JRE bin folder where the JRE used by the monitored application server can be found. For example, if you have installed WebLogic 8.1 in your `/opt/bea` directory, the JRE directory is `/opt/bea/jdk142_05/JRE`.

Select **Next** to continue with the installation.

3 Confirm the pre-installation summary information.

The installation settings that you selected are displayed. Review the information to make sure that you are satisfied.

To change your settings, select **Previous** to return to the previous prompts.

Select **Next** to start the installation.

4 Close the installation wizard.

When the installation has completed, review the post-installation summary information to make sure that the installation completed successfully.

Select **Finish** to exit the installation.

5 Configure the application server.

Before you can use the probe to monitor your application, you need to instrument the JRE and manually modify the application startup script.

- a** If you chose to skip running the JRE Instrumenter during this installation, you need to run it manually before you can use the probe. For more information, see “Running the JRE Instrumenter” on page 165.
- b** Once you have installed the probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the probe is started together with the monitored application. For detailed instructions, see Chapter 9, “Configuring the Application Servers.”

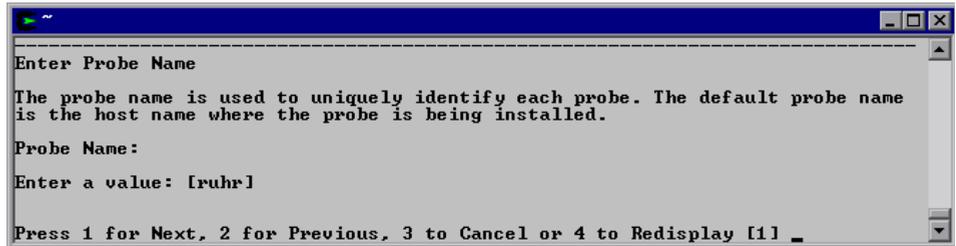
Note: The probe is automatically configured during the installation process. In certain situations however, you may be required to change these default settings or to manually configure the probe. For more information, see “Configuring the Probe Manually” on page 154.

For situations that require more advanced configurations of the probe and application server, see Chapter 16, “Advanced J2EE Probe and Application Server Configuration.”

Installing the Probe to Work with a Mercury Diagnostics Server

If you are installing the probe to work with a Mercury Diagnostics Server, continue with the following procedure.

1 Assign a name for the probe.



Enter a unique name that is to be used to identify the probe within Mercury Diagnostics. The following characters can be used in the name: - , _ , and all alphanumeric characters.

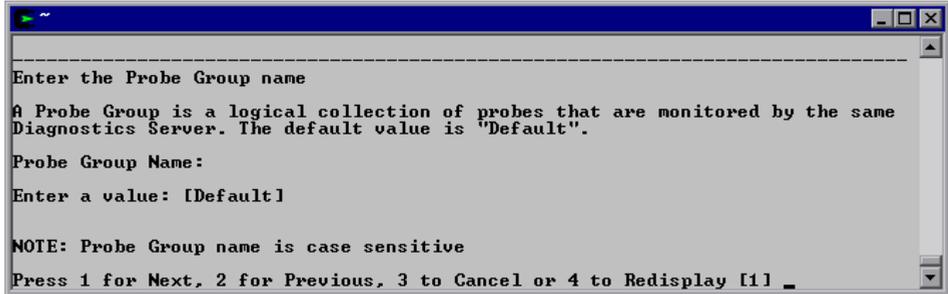
When assigning a name to a probe, choose a name that will help you recognize the application that the probe is monitoring, and the type of probe (in this case, a J2EE Probe).

For example, the probe name for the first J2EE Probe installed, that will be monitoring an application named PetWorld, could be:

PetWorldJ2EEProbe1

Select **Next** to continue with the installation.

2 Assign a name for the probe group.



Enter a name for the probe group for this probe.

Note: The probe group is case-sensitive.

Probe groups are logical groupings of probes that report to the same Diagnostics Server. The performance metrics for a probe group are tracked and can be displayed on many of the Diagnostics screens.

For example, you may want to assign all of the probes for a particular enterprise application to a probe group so that you can monitor both the performance of the group and the performance of the individual probes.

Select **Next** to continue with the installation.

3 Enter the details for the Diagnostics Server in Mediator mode.

Enter the necessary information to enable communication with the Diagnostics Server in Mediator mode.

```

Provide the location of the Diagnostics Server in Mediator mode.
Diagnostics Server Mediator Host <Name or IP address>:
Enter a value: [lissaquah]

Diagnostics Server Mediator Port:
Enter a value: [2612]

The default value is 2612.
[X] 1 - Check the connectivity to the Diagnostics Server
To select an item enter its number, or 0 when you are finished: [0]

Mediator Host and Port.
Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1] _

```

Note: If there is only one Diagnostics Server in the Diagnostics deployment where the probe will run, enter the Diagnostics Server's host name and event port information here.

When there is more than one Diagnostics Server in your deployment, enter the information for the Diagnostics Server in Mediator mode that is to receive the events from the probe.

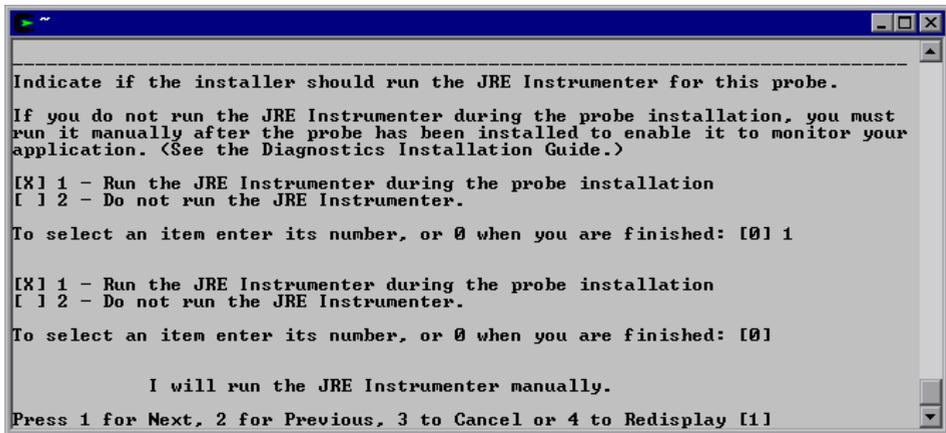
- a** Enter the host name or IP address of the host of the Diagnostics Server.
You should specify the fully qualified host name; not just the simple host name. In a mixed OS environment, where UNIX is one of the systems, this is essential for proper network routing.
- b** Enter the port number of the Diagnostics Server.
The default port for the Diagnostics Server is **2612**. If you have changed the port since the Diagnostics Server was installed, you should specify that port number here instead of the default.

- If you want to make sure that the specified Diagnostics Server is running and accessible from the installation host, select **Check the connectivity to the Diagnostics Server**.

Select **Next** to continue with the installation.

If you selected **Check the connectivity to the Diagnostics Server Mediator Host and Port** and connectivity problems were encountered, the installer provides you with the results of the connectivity check. If you do not want to address these problems at this stage, clear this option, proceed with the installation, and address the problem later.

4 Indicate if you want to run the JRE Instrumenter.



- To run the instrumenter now, select **Run the JRE Instrumenter during the probe installation**.
- To skip this step, select **Do not run the JRE Instrumenter**. If you choose to skip this step now, you must run the JRE Instrumenter manually before you can use the probe. For instructions, see Chapter 8, “Running the JRE Instrumenter.”

Note: If you are installing this probe to work with SAP NetWeaver, skip the JRE Instrumenter. You will configure SAP NetWeaver manually. See “Configuring the SAP NetWeaver Application Server” on page 204.

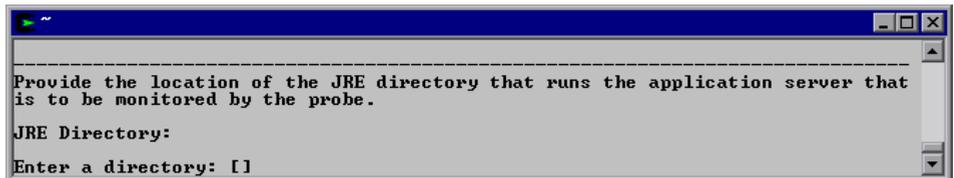
If you are installing this probe to work with WebSphere in an Integrated Development Environment (IDE), skip the JRE Instrumenter. You will run the JRE Instrumenter manually. See “Running the JRE Instrumenter for WebSphere IDE” on page 187.

Select **Next** to continue with the installation.

Next Step: At this stage, the installation procedure branches, depending on the option that you selected for running the JRE Instrumenter.

- ▶ If the installer is running the JRE Instrumenter, the installation process continues as documented in the next step.
 - ▶ Otherwise, the installation process skips the next step, and proceeds with the step following the next.
-

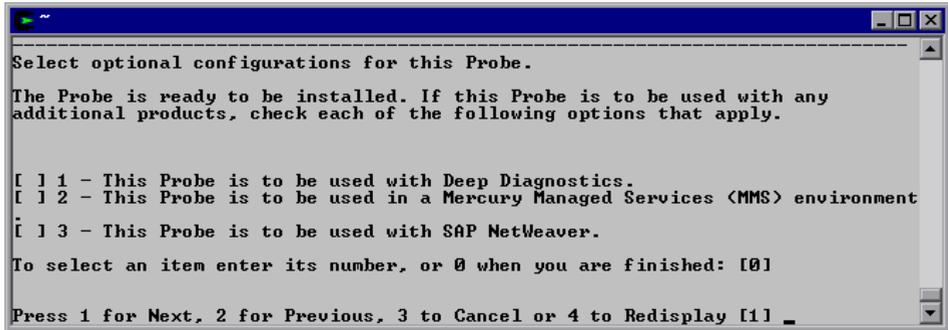
5 Enter the JRE directory.



Type the path to the JRE bin folder where the JRE used by the monitored application server can be found. For example, if you have installed WebLogic 8.1 in your `/opt/bea` directory, the JRE directory is `/opt/bea/jdk142_05/JRE`.

Select **Next** to continue with the installation.

6 Select the configuration options if the probe is to be used with additional products.



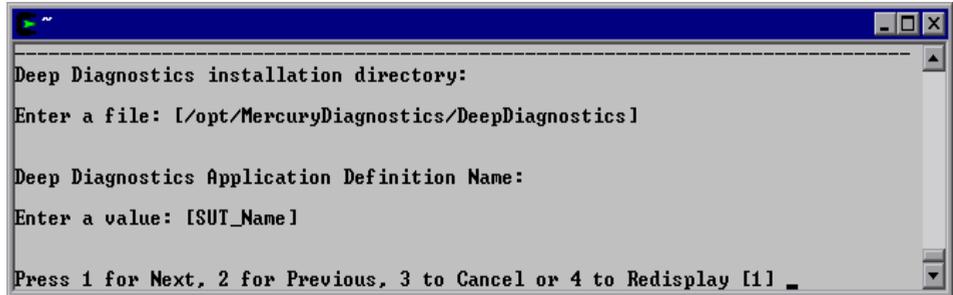
Indicate whether you will be using the probe with other products.

- To use the probe to collect data for Deep Diagnostics, select **This Probe is to be used with Deep Diagnostics.**
- To use the probe within a Mercury Managed Services (MMS) environment, select **This Probe is to be used in a Mercury Managed Services (MMS) environment.**
- To use the probe with a SAP NetWeaver Web application server, select **This Probe is to be used with SAP NetWeaver.**

Select **Next** to continue with the installation.

7 Enter the Deep Diagnostics information.

If you indicated that you want to use this probe with Deep Diagnostics, enter the Deep Diagnostics information.



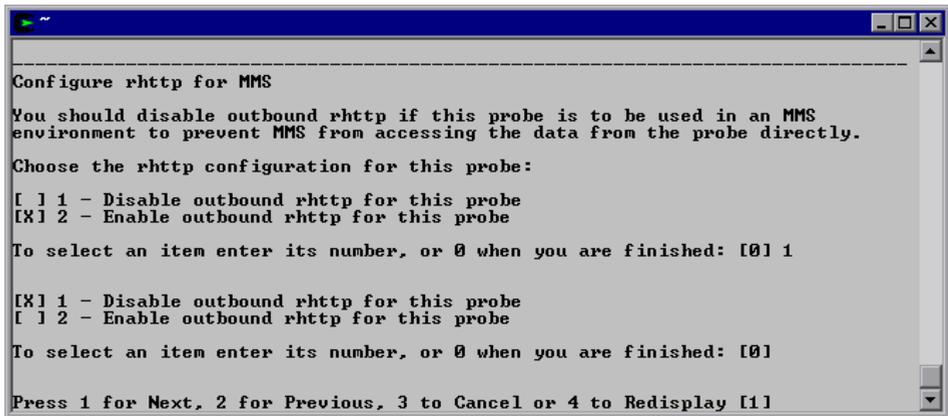
- a** Enter the Deep Diagnostics installation directory.
- b** Enter the Deep Diagnostics Application Definition name. The Deep Diagnostics Application Definition name must exactly match the name specified for the Application Definition in Deep Diagnostics.

Select **Next** to continue with the installation.

8 Configure the Mercury Managed Services environment.

If you indicated that the probe is to be used in a Mercury Managed Services environment (see step 6 above), it is recommended that you disable outbound rhttp (reverse http).

The Diagnostics Server uses rhttp to query metrics from the J2EE Probe. You should disable the probe's outgoing rhttp to prevent Mercury Managed Services from accessing the data directly from the probe.



- To disable outbound rhttp, select **Disable outbound rhttp for this probe**.
- To enable outbound rhttp, select **Enable outbound rhttp for this probe**.

Select **Next** to continue with the installation.

9 Confirm the pre-installation summary information.

The installation settings that you selected are displayed. Review the information to make sure that you are satisfied.

To change your settings, select **Previous** to return to the previous prompts.

To start the installation of the J2EE Probe, select **Next**. A progress bar and status message is displayed.

10 Close the installation wizard.

When the installation has completed, review the post-installation summary information to make sure that the installation completed successfully.

Select **Finish** to exit the installation.

11 Configure the application server.

Before you can use the probe to monitor your application, you need to instrument the JRE and manually modify the application startup script.

- a** If you chose to skip running the JRE Instrumenter during this installation, you need to run it manually before you can use the probe. For more information, see “Running the JRE Instrumenter” on page 165.
- b** Once you have installed the probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the probe is started together with the monitored application. For detailed instructions, see Chapter 9, “Configuring the Application Servers.”

Note: The probe is automatically configured during the installation process. In certain situations however, you may be required to change these default settings or to manually configure the probe. For more information, see “Configuring the Probe Manually” on page 154.

For situations that require more advanced configurations of the probe and application server, see Chapter 16, “Advanced J2EE Probe and Application Server Configuration.”

12 Verify the probe installation.

If you are installing this probe to work with a Mercury Diagnostics Server, you can verify the probe installation as described in “Verifying the J2EE Probe Installation” on page 158.

Installing the J2EE Probe on a z/OS Mainframe

This section provides instructions for installing the J2EE Probe from the .tar file that is included on the product CD.

Important Considerations when Installing the J2EE Probe on z/OS

Consider the following before you install a probe in a z/OS environment:

Editing Property Files on a z/OS Mainframe

When installed in a z/OS environment, the Probe expects the Diagnostics property files to be in EBCDIC format. Use an EBCDIC editor to update the property files and make sure to store the updates in the same format.

Viewing the System Logs in z/OS

You can view the system log by accessing the primary operator's console in SDSF.

Capturing Metrics on the z/OS Mainframe

The J2EE Probe has been configured to capture Deep Diagnostics Java metrics on the z/OS mainframe. Load test metrics are not captured for z/OS. The Probe can be configured to capture a limited number of system level metrics.

For more information on capturing system metrics in z/OS, see “Enabling z/OS System Metrics Capture” on page 423.

Installing the J2EE Probe on z/OS from the Product CD

A tar file containing the probe files has been included on the product CD.

To install the J2EE Probe on a z/OS mainframe:

- 1 Upload **ProbeInstall_zOS.tar.gz** from the **Diagnostics_Probes** folder on the Mercury Diagnostics installation CD (or directly from the Probe installation cd) to the directory on the z/OS machine where you want to unzip the installer.
- 2 Unzip **ProbeInstall_zOS.tar.gz** using gzip as shown in the following example:

```
gzip -d ProbeInstall_zOS.tar.gz
```

This command creates the unzipped file, **ProbeInstall_zOS.tar**.

- 3 To unpack the tar file, run the tar command as shown in the following example:

```
tar -xvf ProbeInstall_zOS.tar
```

This command creates the unpacked directory, **ProbeInstall**.

- 4 Prepare to run **jreinstrumenter.sh** by setting the permissions on the script so that it can be run (o+x).

- 5 Run the JRE Instrumenter. For z/OS, use the following command:

```
<path_to_java> -jar <path_to_jre_instrumenter> -b <jvm_directory>
```

where:

- ▶ <path_to_java> is the path to the java executable.
 - ▶ <path_to_jre_instrumenter> is
 <probe_install_dir>/lib/jreinstrumenter.jar
 - ▶ <jvm_directory> is the path to the JRE for the java installation that your application is using.
- 6 Configure the probe properties to enable the probe to work with the other Diagnostics components as described in “Configuring the Probe Manually” on page 154.
 - 7 Once you have installed the probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the probe is started together with the monitored application. For detailed instructions, see Chapter 9, “Configuring the Application Servers.”
 - 8 Verify the probe installation as described in “Verifying the J2EE Probe Installation” on page 158.
 - 9 If you want to collect z/OS system metrics, see “Enabling z/OS System Metrics Capture” on page 423.

Installing J2EE Probes on Multiple z/OS Machines

If you will be installing Probes on more than one z/OS machine, you may want to create a pax archive of the Probe implementation on the first machine and then use the pax archive to install the probe onto the other machines. Contact your system administrator for more information.

Installing the J2EE Probe Using the Generic Installer

Important! This section only applies if you are going to install the Probe to run with a Diagnostics Server.

The installers for the J2EE Probe have been built to support installing the Probe on all of the platforms for which the component has been certified. However, the Probe may work on other platforms that have not yet been certified. A generic installer has been provided on the product CD to allow you to install the Probe on these uncertified platforms.

To get the Probe to work on the platforms that are not supported by the regular installer, you must run the generic installer, and manually configure the Probe so that it will be able to communicate with the other Diagnostics components and be able to monitor the processing of your application.

To install and configure the J2EE Probe on an uncertified platform:

- 1 Locate **ProbeInstall.tar.gz** from the **Diagnostics_Probes** folder on the Mercury Diagnostics installation CD (or directly from the Probe installation cd).
- 2 Unzip **ProbeInstall.tar.gz** using **gzip** as shown in the following example:

```
gzip -d ProbeInstall.tar.gz
```

When this command completes, the unzipped file is called **ProbeInstall.tar**.

- 3 To unpack the tar file, execute the tar command as shown in the following example:

```
tar -xvf ProbeInstall.tar
```

This command creates the unpacked **ProbeInstall** directory.

- 4 Configure the probe to enable the probe to work the other Diagnostics components as described in “Configuring the Probe Manually” on page 154.

- 5 Configure the Application Server to allow the probe to monitor your application.
 - a Run the JRE Instrumenter as described in “Running the JRE Instrumenter” on page 165.
 - b Once you have installed the probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the probe is started together with the monitored application. For detailed instructions, see Chapter 9, “Configuring the Application Servers.”
- 6 Verify the probe installation as described in “Verifying the J2EE Probe Installation” on page 158.

Silent Installation of the J2EE Probe

A *silent installation* is an installation that is performed automatically, without the need for user interaction. In place of user input, the silent installation accepts input from a response file for each step of the installation.

For example, a system administrator who needs to deploy a component on multiple machines can create a response file that contains all the prerequisite configuration information, and then perform a silent installation on multiple machines. This eliminates the need to provide any manual input during the installation procedure.

Before you perform silent installations on multiple machines, you need to generate a response file that will provide input during the installation procedure. This response file can be used in all silent installations that require the same input during installation.

To generate a response file:

Perform a regular installation with the following command-line option:

```
<installer> -options-record <responseFileName>
```

This creates a response file that includes all the information submitted during the installation.

To perform a silent installation:

Perform a silent installation using the relevant response file.

You perform the silent installation with the **-silent** command-line option as follows:

```
<installer> -silent -options <responseFileName>
```


7

Post-Installation Configuration of the Mercury Diagnostics Probe for J2EE

This chapter provides instructions for configuring and working with the performing the Mercury Diagnostics Probe for J2EE (J2EE Probe) after it has been installed. You should review these instructions to make sure that you have configured the Probe to effectively and efficiently monitor your application's performance.

This chapter describes:	On page:
Configuring the Probe Manually	154
Configuring the Application Server	157
Configuring the Probe for Advanced Situations	158
Deploying BEA JRockit Mission Control	158
Verifying the J2EE Probe Installation	158
Determining the Version of the J2EE Probe	160
Upgrading to a Newer Version of the J2EE Probe	161
Uninstalling the J2EE Probe	164

Configuring the Probe Manually

The Probe is automatically configured during the installation process. The default configuration of the J2EE Probe has been set to monitor your application effectively, with very little impact on the performance of your application. In the following situations however, you may be required to change these default settings or to configure the Probe manually.

- ▶ The Windows and UNIX installers gather the information needed to automatically configure the Probe. If you skipped some of the optional steps in the installer, or if you need to change some of the information that you provided to the installer after the Probe has been installed, you may need to perform some or all of the Probe configuration steps manually.
- ▶ After you have installed the Probe using the generic installation process or the z/OS installation process, you must configure the Probe properties so that the Probe can work with the other Diagnostics Components.

Note: For information about more advanced Probe configuration instructions, see Chapter 16, “Advanced J2EE Probe and Application Server Configuration.”

To configure the J2EE Probe manually:

1 Configure the Probe properties:

a In `<probe_install_dir>/etc/probe.properties`:

- Set the **active.products** property to indicate the products that the Probe will be working with. See the comments in the property file for the valid values.
- Set the **id** property to indicate the Probe id for the Probe. Type the unique name that is to be used to identify the Probe within Mercury Diagnostics. The following characters can be used in the name: - , _ , and all alphanumeric characters.
- Set the **group** property to assign the Probe to a Probe Group. Make sure that the Probe Group is exactly the same as that of the other Probes with which you want the Probe grouped.

b In `<probe_install_dir>/etc/dynamic.properties`:

- Set the **mediator.host.name** property to indicate the host name for the Diagnostics Server in Mediator mode.
- Set the **mediator.port.number** property to indicate the port number where the Diagnostics Server Mediator listens for Probe communication. The default port number is 2612. If you have changed the port since the Diagnostics Server was installed, you should specify that new port number here instead of the default.

c In `<probe_install_dir>/etc/dispatcher.properties`:

Set the **registrar.url** property value to the URL for the Diagnostics Server in Mediator mode to indicate the host name and port for the Diagnostics Server in Mediator mode.

```
registrar.url=http://<mediator>:2006/commander/registrar/
```

Note: Make sure that you update the **registrar.url** property in the property file that is appropriate for the configuration of your network. The property **registrar.url** exists in two different property files in **<probe_install_dir>\etc**.

The **registrar.url** property in **dispatcher.properties** is used to provide the URL for the registrar when there is not a proxy.

The **registrar.url** property in **webserver.properties** is used to provide the URL for the registrar when there is a proxy server between the Probe and the registrar.

- d** If the Probe will be used with Deep Diagnostics, update the **pfmt.godel.probedir** property in **\$DD/etc/pfmt.godel.expert.properties**.
- 2** Ensure that the permission bits on the **<probe_install_dir>/log** directory are set so that the user who starts the application server has write (o+w) permissions.
- 3** If you want to use this Probe with Mercury Deep Diagnostics for J2EE, configure the J2EE Probe to collect data for Deep Diagnostics. For more information, see “Configuring the Probe to Work with Mercury Products” on page 307.

Important! Before you can use the J2EE Probe to monitor your application, you need to instrument the JRE and manually modify the application startup script. For more information, see “Configuring the Application Server” below.

Configuring the Application Server

In order to allow the Probe to monitor your application, you need to instrument the JRE and manually modify the application server startup script.

Instrumenting the JRE

The instrumentation process defines for the Probe, the classes and methods that are to be monitored in your application. It also controls the layers that will be used to report the performance metrics that were gathered.

You use the JRE Instrumenter to run the instrumentation. The JRE Instrumenter automatically runs at the end of the Probe installation to prepare your application for instrumentation unless you indicate otherwise.

If you elect to skip running the JRE Instrumenter during the Probe installation, you must run it manually before you can use the Probe to monitor your application. For more information, see Chapter 8, “Running the JRE Instrumenter.”

Modifying the Application Startup Script

Once you have installed the Probe and instrumented the JRE, you need to manually modify the startup script for your application server so that the Probe is started together with the monitored application.

For detailed instructions about how to modify your application server’s startup script, see Chapter 9, “Configuring the Application Servers.”

Note: The process for configuring the J2EE Probe and your application servers when there are multiple JVMs on a single machine is described in “Configuring the Probes for Multiple Application Server JVM Instances” on page 312.

Configuring the Probe for Advanced Situations

You should review the topics in Chapter 16, “Advanced J2EE Probe and Application Server Configuration” to determine if there are Probe configuration settings that will enable the Probe to work more efficiently in your environment.

Deploying BEA JRockit Mission Control

The Mercury Diagnostics license allows you to use BEA JRockit Mission Control. If the application that you are monitoring is using BEA JRockit JVM and if the J2EE Probe has been configured to work with a Diagnostics Server, then you will be able to access a license for BEA JRockit Mission Control and use Mission Control to monitor your application’s performance.

For more information see “Using BEA JRockit Mission Control” on page 341 for instructions on how to access and deploy the license for BEA JRockit Mission Control.

Verifying the J2EE Probe Installation

You can verify the Probe installation using the System Health Monitor. For detailed instructions about how to use the System Health Monitor, see the *Mercury Diagnostics User’s Guide*.

Important! This section only applies if you installed the Probe to run with a Diagnostics Server.

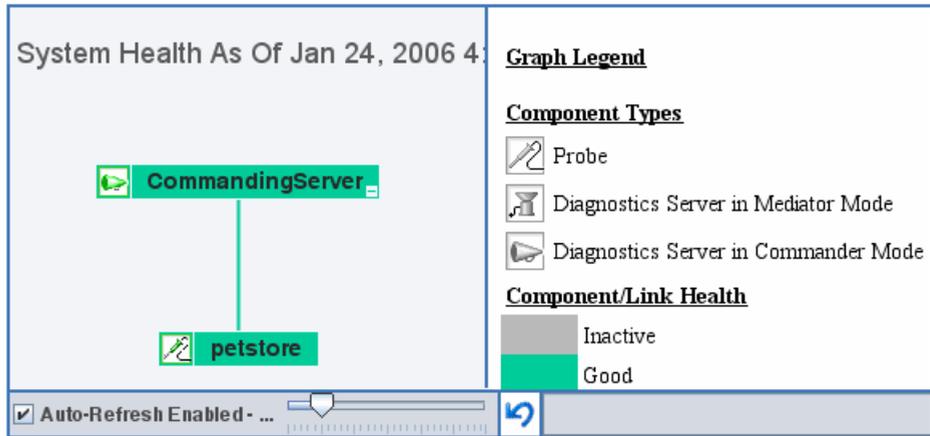
Note: The Probe does not register with the Diagnostics Server until it is started. The Probe is started when the instrumented application server is started. Therefore, you cannot verify the installation of the Probe using the System Health Monitor until you have configured the Probe and Application Server as described in “Introducing Installation and Configuration of Diagnostics Probes” on page 73.

If you have been following the recommended installation sequence, after you have installed the J2EE Probe you are able to verify the following:

- ▶ The Diagnostics Server in Commander mode was successfully installed.
 - ▶ Where relevant, additional Diagnostics Servers in Mediator mode were successfully installed and are communicating with the Diagnostics Server in Commander mode.
 - ▶ The J2EE Probe was successfully installed and has established connectivity with the Diagnostics Server.
-

Note: The J2EE Probe is not displayed in System Health when first installed because the Probe does not register with the Diagnostics Server until it is started. The Probe is started and is registered with the Diagnostics Server when you start the monitored application server.

Depending on your deployment, the new J2EE Probe is shown on the System Health Monitor as a child of either the Diagnostics Server in Commander mode or the Diagnostics Server in Mediator mode.



The J2EE Probe is colored grey in the System Health Monitor when the monitored application is closed. Once the application has been started the Probe is colored green.

Determining the Version of the J2EE Probe

When you request support, it is useful to know the version of the Diagnostics component that you have a question about.

You can find the version of the Probe in one of the following ways:

- ▶ Locate the version file `<Probe_install_dir>\version.txt`. The file contains the 4-digit version number, as well as the build number.
- ▶ The version number is available in the Probe log file (`<Probe_install_dir>/log/<probe_id>/probe.log`)
- ▶ The version number is available among the details of the Probe on System Health Monitor.

Upgrading to a Newer Version of the J2EE Probe

When you install the Probe, the installer detects if a Probe has already been installed in the installation directory that you specified. The installer upgrades the Probe installation to the current version and creates a backup of the property settings in the `<probe_install_dir>/etc` directory. The backup file is called `<probe_install_dir>/etc.old`.

Note: If you do not want the installer to upgrade the installation for an existing probe, you should specify a different installation directory when prompted by the installer.

The following instructions are instructions for upgrading from an older version of the Mercury Diagnostics Probe for J2EE.

Note: The current version of the J2EE Probe has been designed to work with the current versions of the Diagnostics components that are part of the Mercury LoadRunner 8.1, Mercury Performance Center 8.1, and Mercury Business Availability Center 6.x installations. The Probe may not work with earlier versions of these products. Likewise, previous versions of the Probe may not work with the current versions of these Mercury products.

To upgrade the J2EE Probe to a newer version of the Probe:

- 1** Shutdown the application servers that are being monitored by the Probes that you want to upgrade.
- 2** Create a backup copy of the startup script for the application server.

- 3 If you want to reuse the configuration of your current Probe with the new Probe, create a backup copy of the folder **<Probe_install_dir>/etc** called **old_etc_backup**.

Note: This optional step could help you save time by reusing the current settings for Probe configurations such as instrumentation changes, buffering changes, port numbers, Diagnostics Server locations, probe.id, and product mode.

- 4 Uninstall the Probe. For instructions, see “Uninstalling the J2EE Probe” on page 164.
- 5 Install the new version of the J2EE Probe as instructed in Chapter 6, “Installing the Mercury Diagnostics Probe for J2EE.”

Note: When you install the J2EE Probe, do not select the option to have the installer configure your application server, and do not select the option to allow the installer to instrument the JRE. When the Probe installation is complete, proceed with the steps below.

- 6 If you made a backup copy of the old **<Probe_install_dir>/etc** folder so that you could reuse the old Probe’s configuration in step 3 above, you should follow the backup information to configure the new Probe.
 - a Create a backup of the newly installed folder **<Probe_install_dir>/etc** called **new_etc_backup** so that you have a copy of the files as they were created by the installer.
 - b Rename the **<Probe_install_dir>/etc/auto_detect.points** file that was just installed to **new_auto_detect.points**.
 - c Copy the backup copy of **auto_detect.points** into **<Probe_install_dir>/etc**.

- d** The old **auto_detect.points** file will work with the new Probe. If you want to take advantage of the enhancements in the new version of the Probe, copy the following sections from **new_auto_detect.points** file:
- All sections that begin “BEA-”, including the preceding comments
 - All sections that begin “SAP_”, including the preceding comments
 - The “Synchronization” section, including the preceding comments
 - The “RMI” section, including the preceding comments
 - args_by_class
 - Any enhancements/optimizations that you made
- e** Copy the rest of the files, excluding **modules.properties** and **webserver.properties**, from **old_etc_backup**. Replace the files in the current **<Probe_install_dir>/etc/**.

Note: Do not replace **modules.properties** and **webserver.properties**. You must use the new version of these files.

- f** In **capture.properties**, add the following property:

```
ac.register.sink=true
```

- g** In **dispatcher.properties**, add the following properties:

```
dispatcher.properties.file.name=dispatcher.properties
registrar.server_ping.time=90s
ac.autostart=true
```

- h** In **inst.properties**, revise the value of the following property:

```
classes.to.exclude=!aik\.security\.*;!c8e\.*;!org\.jboss\.net\.protocol\.file\
Handler;!org\.jboss\.net\.protocol\.file\.URLConnection;!.*ByCGLIB.*
```

- i** In **inst.properties**, add the following property:

```
rmi=com.mercury.opal.capture.inst.RMIServerInstrumenter,com.mercury.o
pal.capture.inst.RMIProxyInstrumenter
```

- 7 Configure the J2EE Probe and application server following the instructions in Chapter 4, “Introducing Installation and Configuration of Diagnostics Probes.”

Note: If a firewall separates your Probe and the Diagnostics Servers from the other Diagnostics components, make sure that you have configured the firewall to allow communications across the firewall using the port numbers in the range that you specify. For more information, see “Configuring Diagnostics Components to Work with a Firewall” on page 389.

To avoid having to reconfigure your firewall, you may want to override the default port numbers for the new Probe so that they will match the port numbers that the older Probe used.

Uninstalling the J2EE Probe

To uninstall the J2EE Probe:

- ▶ On a Windows machine, run **uninstaller.exe** which is located in the `<Probe_install_dir>_uninst` directory.
- ▶ On a UNIX machine, run **uninstall*** which is located in the `<Probe_install_dir>_uninst` directory.

8

Running the JRE Instrumenter

This chapter explains how to run the JRE Instrumenter to enable the Mercury Diagnostics Probe for J2EE to monitor your application.

This chapter describes:	On page:
About the JRE Instrumenter	165
Running the JRE Instrumenter	166

About the JRE Instrumenter

When you install a J2EE Probe, you must run the JRE Instrumenter to prepare your application for the instrumentation that allows the Probe to monitor the processing. The JRE Instrumenter runs automatically during the Probe installation unless you elect not to run it at that time. If you elect not to run the JRE Instrumenter during the installation of the Probe, you can run it manually as described in this chapter.

Moreover, if the JDK (java.exe executable) used by your application server changes, you must manually run the JRE Instrumenter again in order to enable the Probe to monitor the processing.

Note: If a Probe is being used to monitor multiple JVMs, the JRE Instrumenter must be run once for each JVM so that the Probe can be prepared to instrument the applications that are running on each JVM. For details, see “Configuring the Probes for Multiple Application Server JVM Instances” on page 312.

In this chapter `<probe_install_dir>` is used to indicate the directory where the J2EE Probe was installed.

Running the JRE Instrumenter

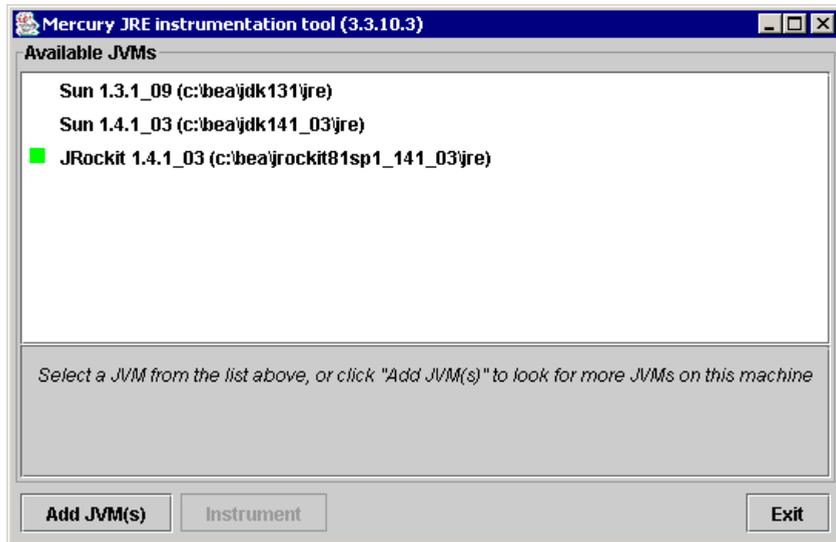
If you chose not to run the JRE Instrumenter during the Probe installation you need to run it manually.

To run the JRE Instrumenter:

1 Start the JRE Instrumenter.

Open `<probe installation directory>\bin` to locate the JRE Instrumenter executable. If the Probe is installed on a Windows machine, run the `jreinstrumenter.cmd` command. If the Probe is installed on a UNIX machine, run the `jreinstrumenter.sh` command.

When the Instrumenter starts, the Mercury JRE Instrumentation Tool dialog box opens.



The JVMs that are available for instrumentation are listed in the **Available JVMs** list.

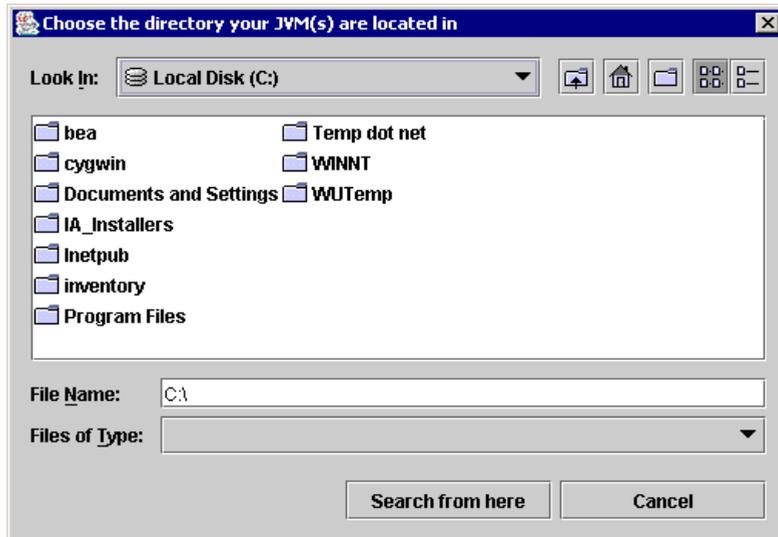
The JVMs that have already been instrumented are listed with green square preceding the name of the JVM.

Note: If the JVM that you want to instrument appears in the **Available JVMs** list, you can skip this step and continue with step 3.

2 Add other JVMs to the Available JVMs list.

- a In the Mercury JRE Instrumentation Tool dialog box, click **Add JVM(s)** to search for other JVMs to be added to the Available JVMs list.

The Choose the directory dialog box opens.

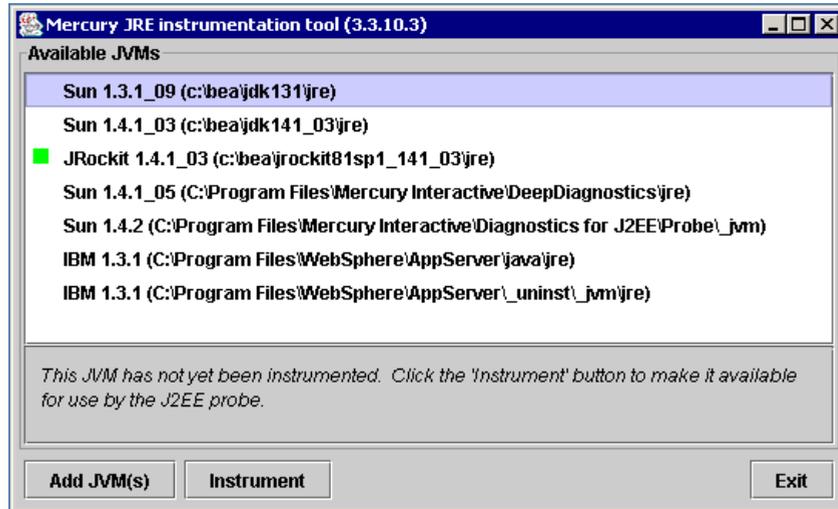


- b Navigate to the location where you would like to begin searching for JVMs using the standard Windows Explorer type navigation tools.
- c Select the file where you would like to begin the search so that its name appears in the **File Name** box.

- d** Click **Search from here** to start searching for JVMs.

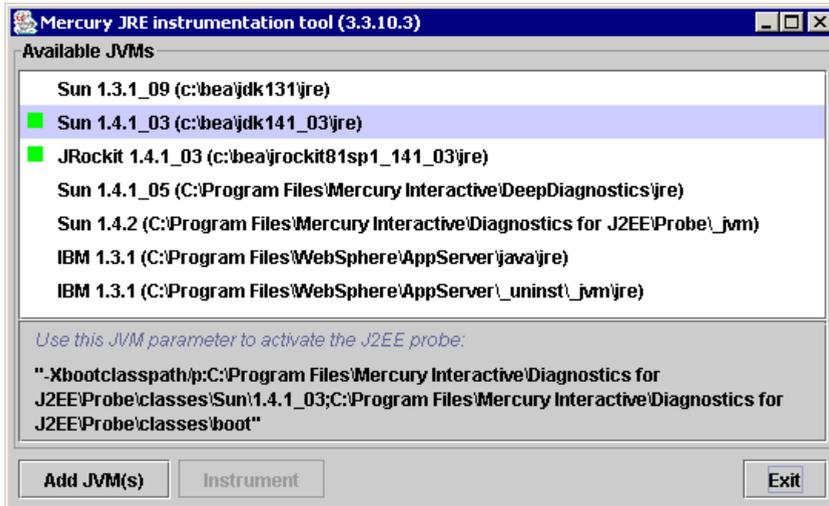
The dialog box closes and the Mercury JRE Instrumentation Tool dialog box opens again, with the command buttons disabled while the tool searches for JVMs.

As the tool locates JVMs, it lists them in the **Available JVMs** list.



3 Instrument the selected JVM.

Select a JVM from the **Available JVMs** list and click **Instrument**. The Instrumenter prepares the JVM to work with the Probe and displays the JVM parameter in the box above the command buttons.



4 Include the JVM parameter in the application server's startup script.

Copy the JVM parameter displayed in the box above the command buttons and include it in the application server's startup script.

5 Exit the JRE Instrumenter.

Click **Exit** to close the Mercury JRE Instrumentation Tool.

9

Configuring the Application Servers

The chapter provides instructions on how to configure your application server to allow the probe to monitor your application.

This chapter describes:	On page:
About Configuring the Application Server	172
Configuring WebSphere Application Servers	173
Configuring WebLogic Application Servers	188
Configuring the Oracle9i Application Server	198
Configuring the JBoss Application Server	201
Configuring the SAP NetWeaver Application Server	204
Configuring a Generic Application Server	206
Adjusting the Heap Size for the J2EE Probe in the Application Startup Script	208

About Configuring the Application Server

Once you have executed the JRE Instrumenter for the J2EE Probe, you must modify the startup script for your application so that the probe that is to monitor the application will be started when the application is started.

You can configure the application servers by updating the application server startup scripts manually. The following sections provide instructions for updating the certified application servers manually.

It is possible that your site administrator has site-specific methods for making configuration modifications. In this case, the generic procedure described in “Configuring a Generic Application Server” on page 206 should provide the information that the site administrator needs to implement the required configuration changes.

Note: If there are no instructions for your specific type of application server in the sections that follow, follow the procedure in “Configuring a Generic Application Server” on page 206.

The process for configuring the J2EE Probe and your application servers when there are multiple JVMs on a single machine is described in “Configuring the Probes for Multiple Application Server JVM Instances” on page 312.

Note: In this chapter “<probe_install_dir>” is used to indicate the directory where the J2EE Probe was installed.

When modifying the -Xbootclasspath parameter, be sure to use quotes if there are spaces in the path that you specify.

Configuring WebSphere Application Servers

Supported versions	Operating systems
3.5, 4.x, 5.x, 6.x	AIX 4.3 Windows 2000/2003 Solaris 8

WebSphere servers are controlled using the WebSphere Application Server Administrative Console. The Console has control over the JVM command line and allows you to add classpath elements, define runtime variables (-D variables), and configure the bootclasspath for WebSphere 3.5, 4.x, 5.x, and 6.x versions. You use the Administrative Console to add the Xbootclasspath property and any additional arguments that are needed for the JVM command line.

Note that the appearance of the Console may differ for different versions of WebSphere. The instructions that follow show different screens for each of the certified WebSphere Application Server versions. The way that changes are implemented in each version of WebSphere may vary slightly. As a result, the examples shown in this section may not correspond exactly to your WebSphere version. The examples should provide the information that you need to enter the required parameters in the appropriate location in the Console.

Note: WebSphere applies the changes that are made on each tab only when you click **Apply** on the tab. Your changes will not be applied until you click **Apply**.

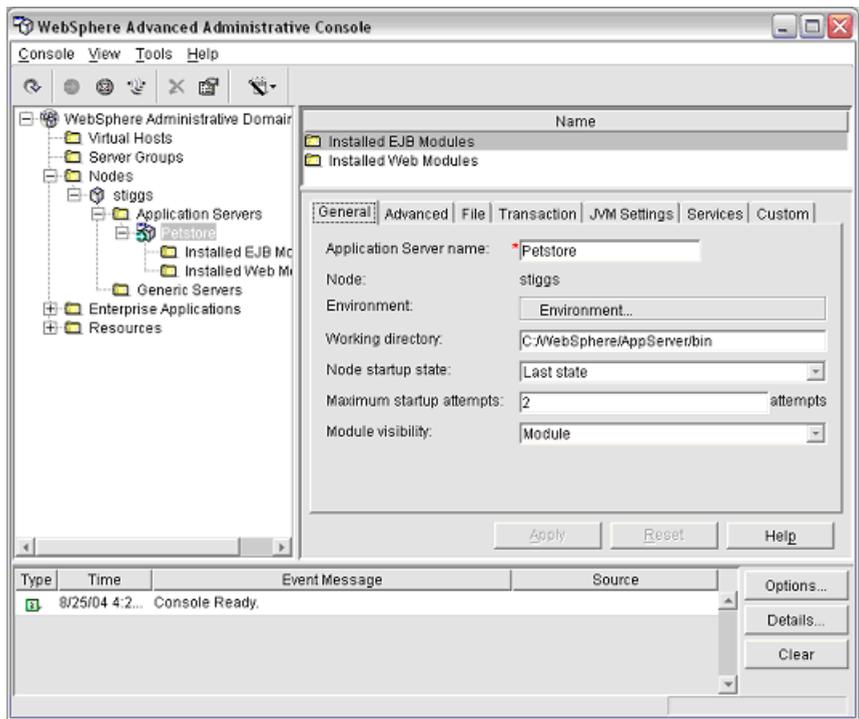
WebSphere 4.0

To configure a WebSphere 4.0 application server:

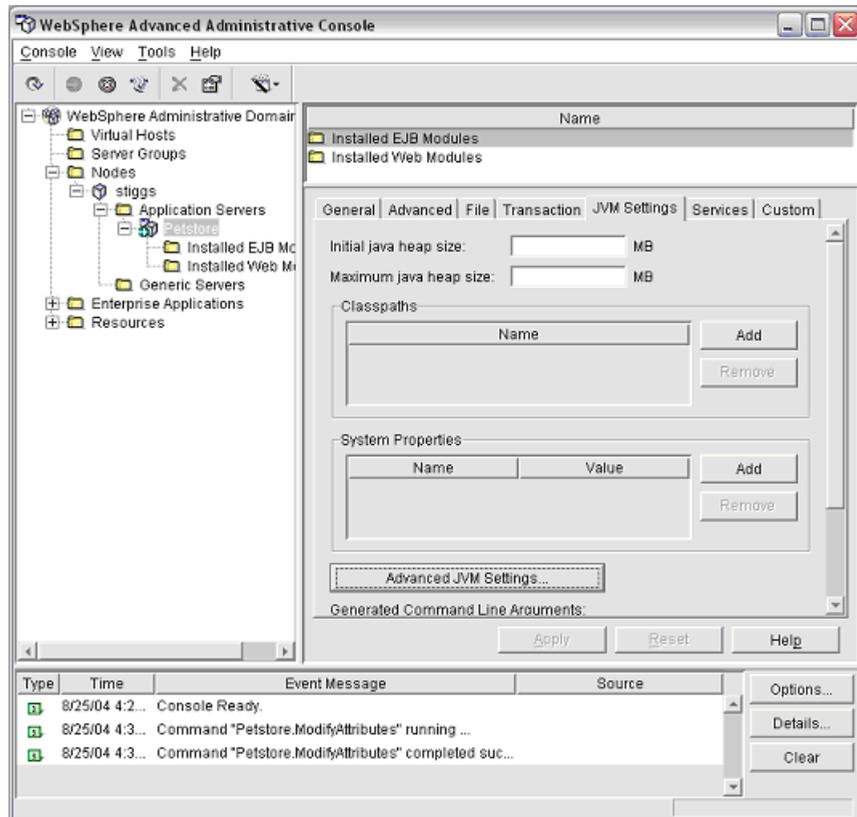
- 1 Use your Web browser to access the Web page for the WebSphere Administrative Console for the application server instance for which the probe was installed:

```
<WebSphere_Install_Dir>\AppServer\bin\adminclient.sh(bat)
```

- 2 From the tree view under the **Nodes** parent, select the application server for the application to be monitored. The General tab for the selected application server is displayed to the right of the tree view.



- 3 Click the **JVM Settings** tab. The JVM Settings information is displayed as shown in the following example:



- 4 Click **Advanced JVM Settings**. The Advanced JVM Settings dialog box opens.



- 5 In the **Boot classpath (prepend)** box, type the boot classpath property as follows:

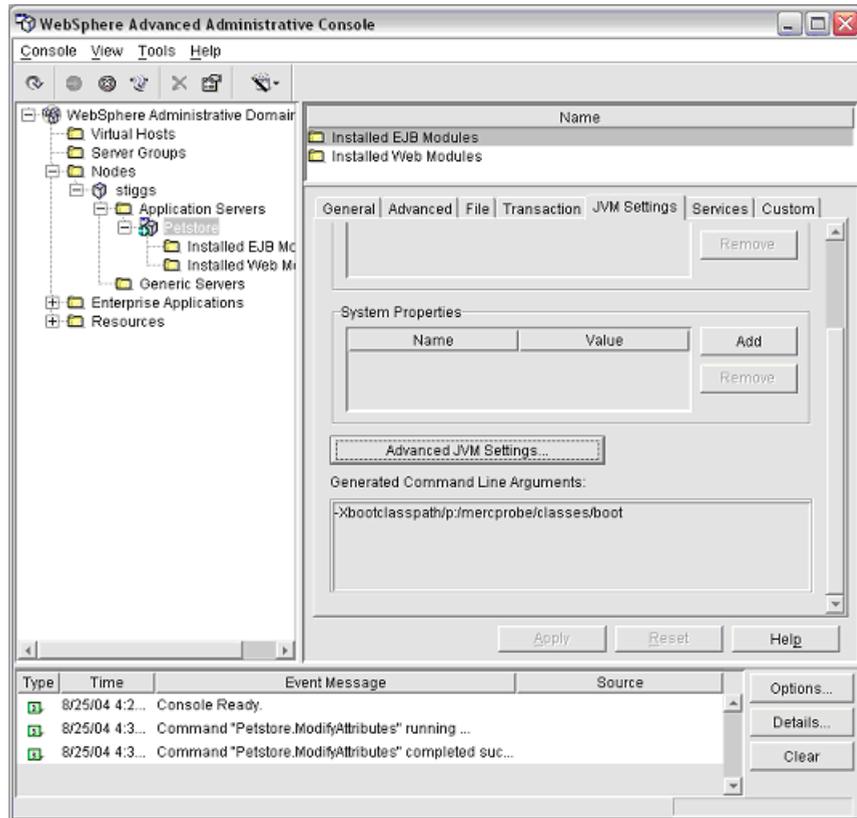
```
<probe_install_dir>\classes\IBM\1.3.1_04;<probe_install_dir>\classes\boot
```

where **<probe_install_dir>** is the path to the directory where the probe was installed.

Note: If your JVM uses a JIT option, such as `-hotspot` or `-classic`, make sure that the `-Xbootclasspath` option is entered following the JIT option in the Java Command Line arguments on the Console.

Click **OK** to close the Advanced JVM Settings dialog box.

- 6 In **Generated Command Line Arguments** area of the JVM Settings tab, verify that the Xbootclasspath property was created properly.



The property should be set as follows:

```
"-Xbootclass-
path/p:<probe_install_dir>\classes\IBM\1.3.1_04;<probe_install_dir>\classes\boot"
```

Note: When modifying the `-Xbootclasspath` parameter. Be sure to use quotes if there are spaces in the path that you specify.

- 7 Scroll down to display the command buttons, and click **Apply** to save your changes which will take effect when you restart the application server.
- 8 Restart the WebSphere application server from the Console menu. You do not need to restart the host for your application server.
- 9 To verify that the probe was configured correctly, check for entries in the `<probe_install_dir>\log\<probe.id>\probe.log` file. If there are no entries in the file, either you did not run the JRE Instrumenter or you did not enter the `Xbootclasspath` correctly. For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

You can also refer to the Websphere **stdout.log** and **stderr.log** files for troubleshooting any problems arising from this configuration.

WebSphere 5.x and 6.x

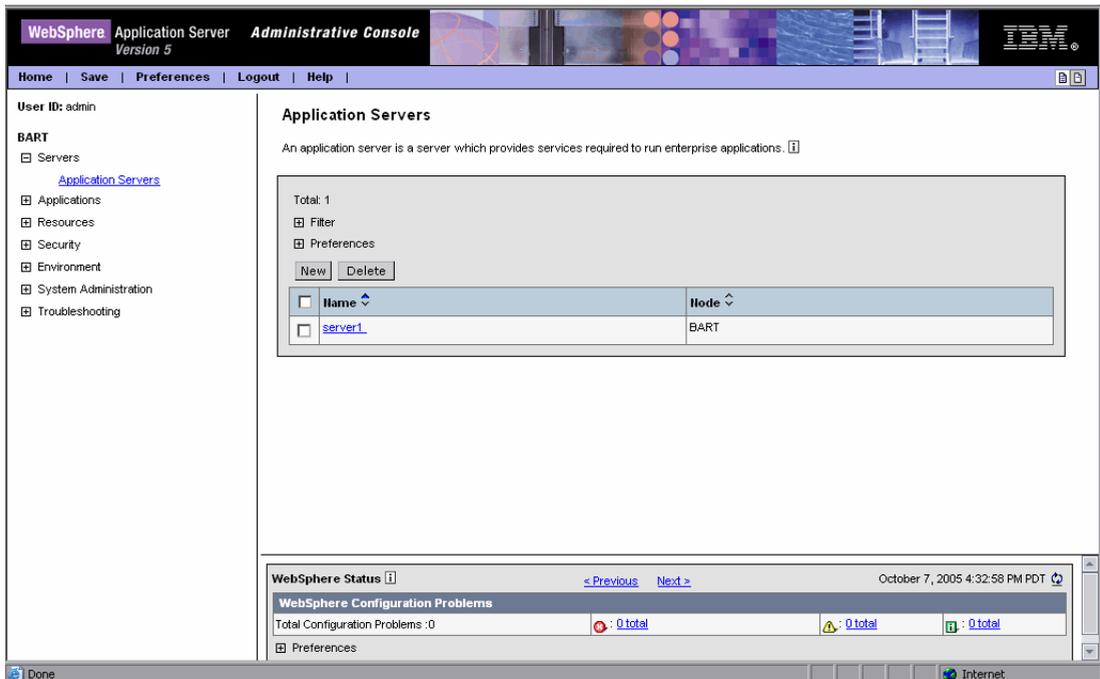
To configure a WebSphere 5.x and 6.x application server:

- 1 Use your Web browser to access the WebSphere Application Server Administrative Console for the application server instance for which the probe was installed:

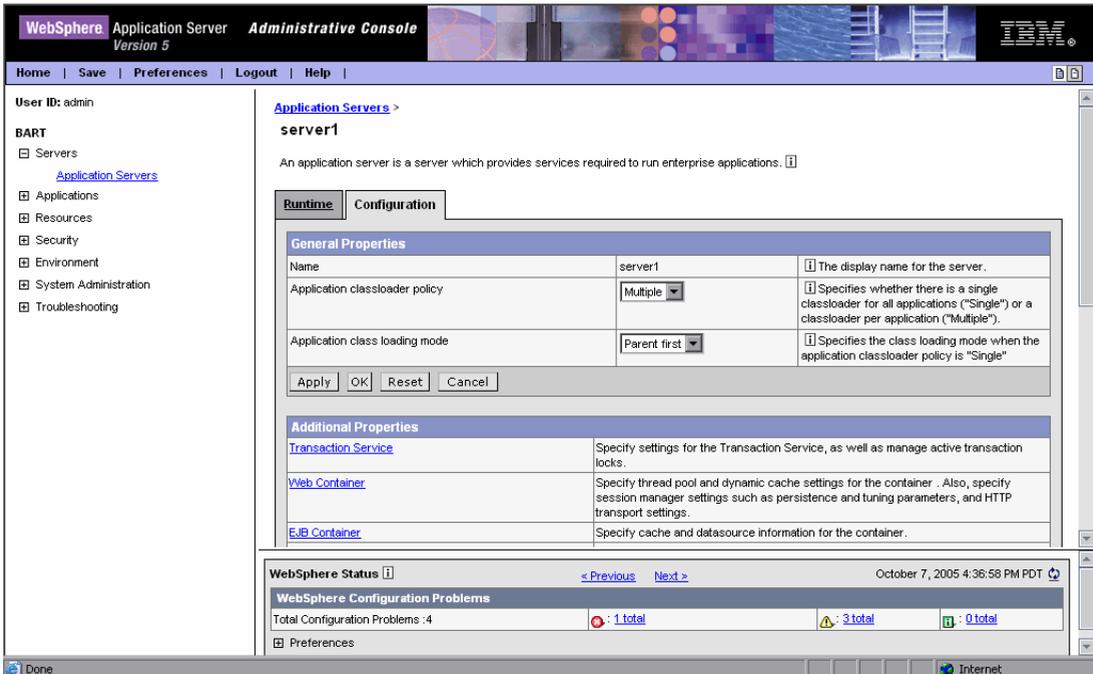
```
http://<App_Server_Host>:9090/admin
```

Replace **<App_Server_Host>** with the machine name for the application server host.

The Websphere Application Server Administrative Console opens.



- 2 In the left panel, select **Servers > Application Servers**.
 - 3 From the list of application servers in the right panel, select the name of the server that you want to configure so that it will be monitored by the probe.
- The Configuration tab for the selected application server is displayed.



4 Scroll down in the Configuration tab, and in the **General Properties** column, look for the **Process Definition** property.

The screenshot shows the WebSphere Administrative Console interface. The top navigation bar includes 'Home', 'Save', 'Preferences', 'Logout', and 'Help'. The left sidebar shows a tree view with 'Application Servers' selected. The main content area displays a list of configuration properties for the selected server. The 'Process Definition' property is highlighted with a red oval. Below the list is a 'WebSphere Status' section showing configuration problems.

Property Name	Description
EJB Container	Specify cache and datasource information for the container.
Dynamic Cache Service	Specify settings for the Dynamic Cache service of this server.
Logging and Tracing	Specify Logging and Trace settings for this server.
Message Listener Service	Configuration for the Message Listener Service. This service provides the Message Driven Bean (MDB) listening process, whereby MDBs are deployed against ListenerPorts that define the JMS destination to listen upon. These Listener Ports are defined within this service along with settings for its Thread Pool.
ORB Service	Specify settings for the Object Request Broker Service.
Custom Properties	Additional custom properties for this runtime component. Some components may make use of custom configuration properties which can be defined here.
Administration Services	Specify various settings for administration facility for this server, such as administrative communication protocol settings and timeouts.
Diagnostic Trace Service	View and modify the properties of the diagnostic trace service.
Debugging Service	Specify settings for the debugging service, to be used in conjunction with a workspace debugging client application.
IBM Service Logs	Configure the IBM service log, also known as the activity log.
Custom Services	Define custom service classes that will run within this server and their configuration properties.
Server Components	Additional runtime components which are configurable.
Process Definition	A process definition defines the command line information necessary to start/initialize a process.
Performance Monitoring Service	specify settings for performance monitoring, including enabling performance monitoring, selecting the PMI module and setting monitoring levels.

WebSphere Status | October 7, 2005 4:38:58 PM PDT

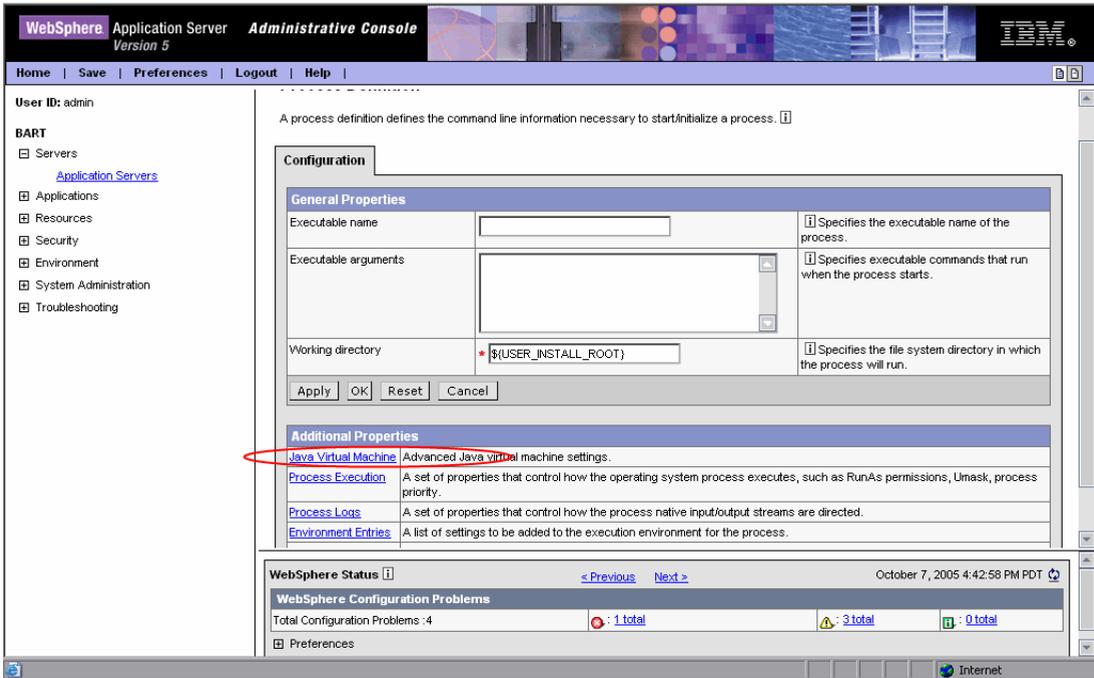
WebSphere Configuration Problems

Total Configuration Problems: 4	1 total	3 total	0 total
---------------------------------	---------	---------	---------

Done | Internet

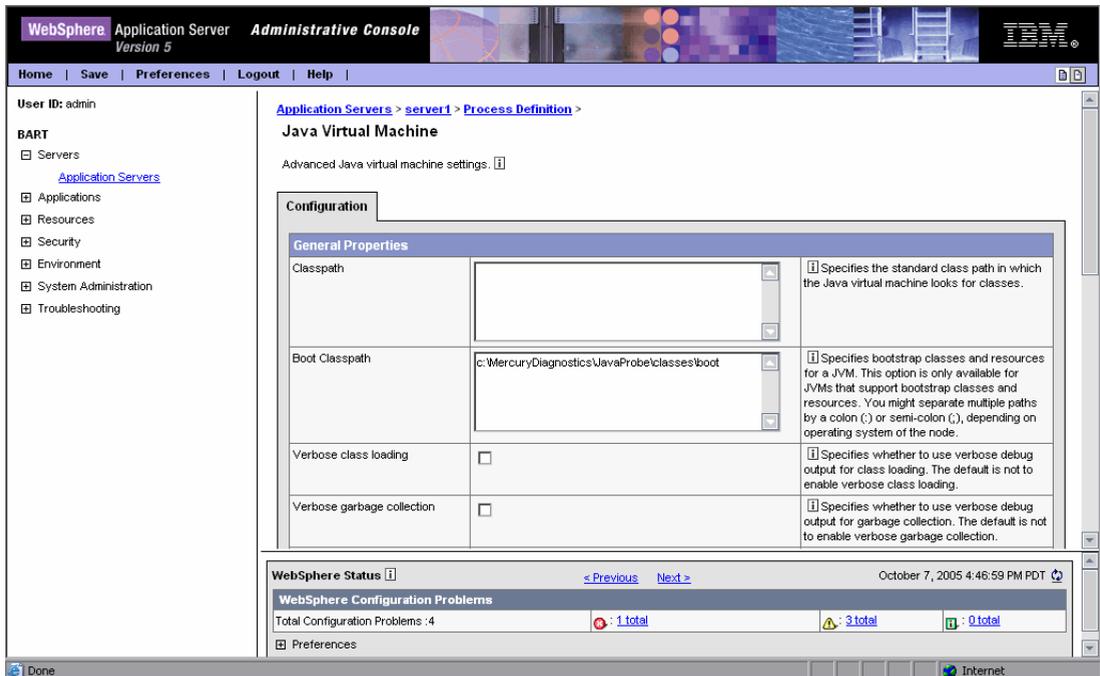
5 Click **Process Definition**.

6 Scroll down in the right panel, and look for **Java Virtual Machine**.



7 Click **Java Virtual Machine**.

8 The Configuration tab for the Java Virtual Machine is displayed.

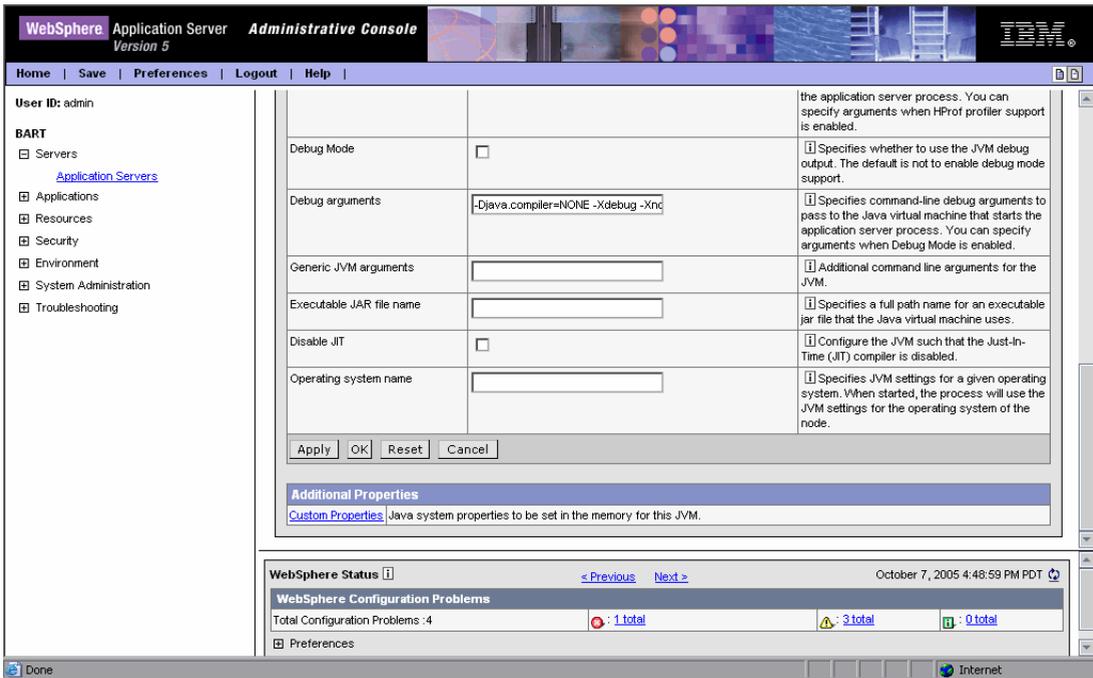


9 In the **Boot Classpath** box, type the path to the boot directory for the probe as follows:

```
<probe_install_dir>\classes\IBM\1.4.2_06;<probe_install_dir>\classes\boot
```

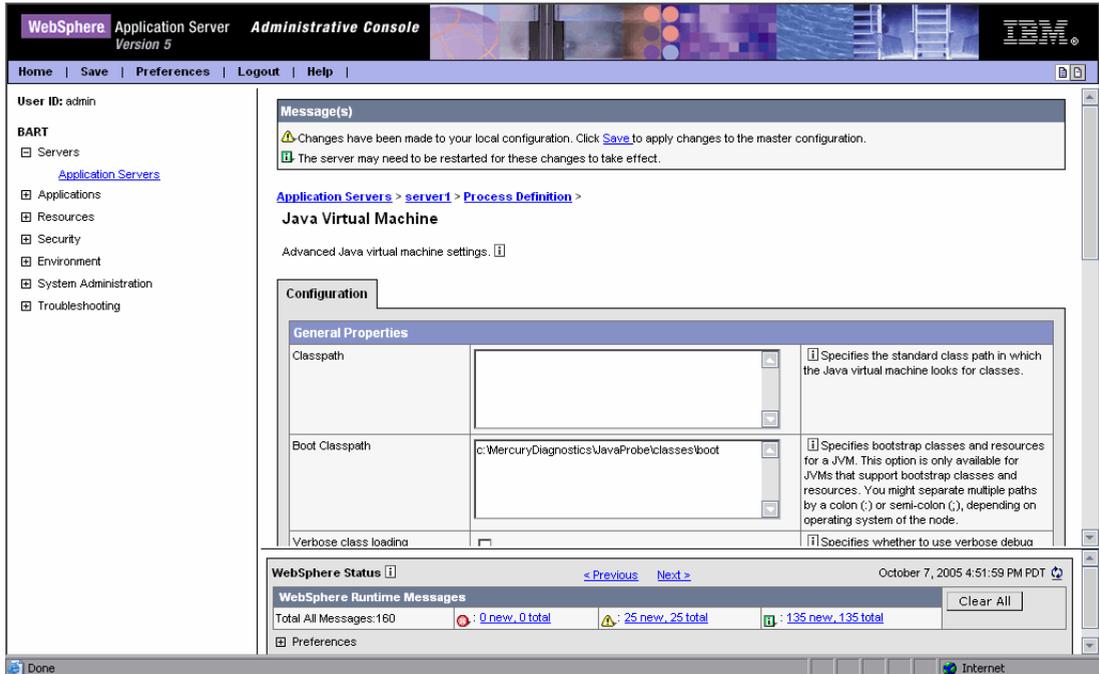
where **<probe_install_dir>** is the path to the location where the probe was installed.

10 Scroll to the bottom of the Configuration tab until the command buttons are visible.



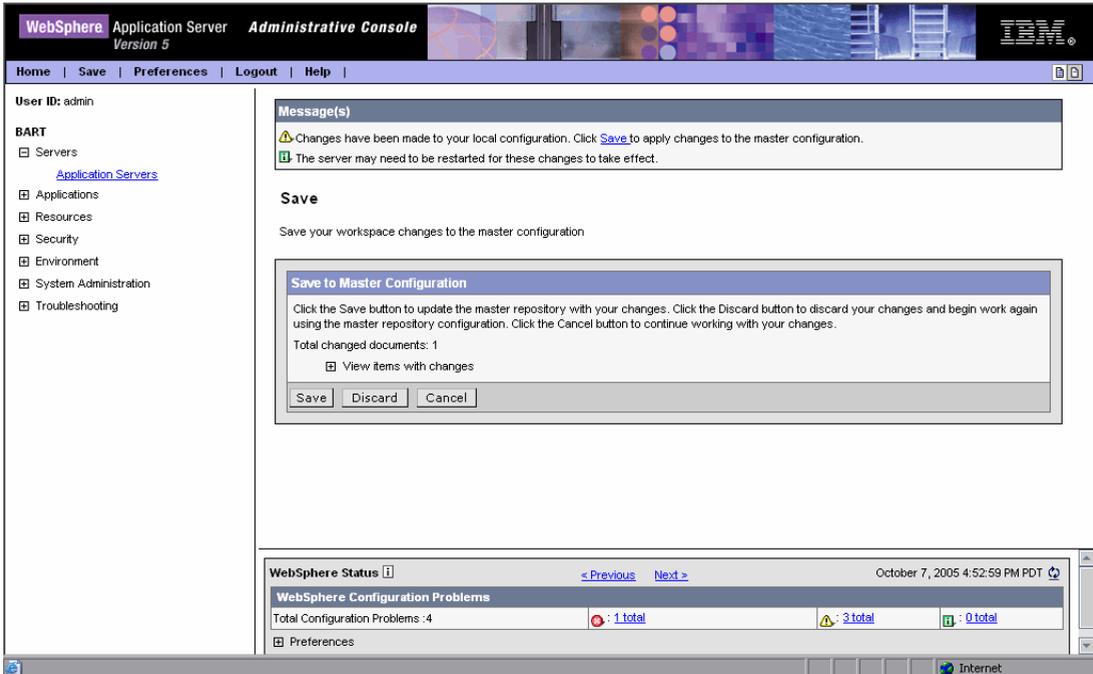
Click **Apply**.

11 A message confirms that your changes have been applied.



Click **Save** to apply the changes to the master configuration.

12 In the **Save to Master Configuration** area, click **Save**.



13 Restart the WebSphere application server. You do not need to restart the host for your application server.

14 To verify that the probe was configured correctly, check for entries in the `<probe_install_dir>\log\<probe_id>.log` file. If there are no entries in the file, this indicates that either you did not run the JRE Instrumenter or you did not enter the Xbootclasspath correctly. For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

Running the JRE Instrumenter for WebSphere IDE

If you are using WebSphere IDE, you must run the JRE Instrumenter manually in order to make sure that the correct JAVA executable for the WSAD IDE has been instrumented.

The WSAD IDE has 10 different java.exe executables to choose from. You must make sure to instrument the one that is used to run the IDE.

To instrument the correct java.exe:

- 1** Determine the version of WebSphere that you are using.
- 2** Determine the location of the appropriate java.exe. See the table below.
- 3** Run the JRE Instrumenter as described in “Running the JRE Instrumenter” on page 165.

Version	Executable
WAS 5.0	IDE INSTALL\runtimes\base_v5\java\bin\java.exe
WAS 5.1	IDE INSTALL\runtimes\base_v51\java\bin\java.exe
WAS 4.x	IDE INSTALL\runtimes\aes_v4\java\bin

Note: Before you can capture data, you must modify the Xbootclasspath for the application’s JVM startup parameters per your IDE’s instructions.

If you are NOT running your application inside the IDE, follow the instructions below to configure the JVM parameters using the WebSphere Administrative Console.

Configuring WebLogic Application Servers

Supported versions	Operating systems
6.1, 7.0, 8.1, 9.0	Windows 2000 SP6 Solaris 8

WebLogic application servers are configured by adding the **Xbootclasspath** property to the script that is used to start the application server. WebLogic is started by running shell scripts in a UNIX environment, or command scripts in a Windows environment. Because the startup scripts that WebLogic provides are frequently customized by a site administrator, it is not possible to provide detailed configuration instructions that apply to all situations. Instead, the following sections provide instructions for each of the certified versions of the WebLogic application server for a generic implementation. Your site administrator should be able to use these instructions to show you how to make these changes in your customized environment.

Note: Make sure you understand the structure of the startup scripts, how the property values are set, and how to use environment variables before you make any configuration changes for the Probe. Always create a backup copy of any file that you are going to update prior to making the changes.

WebLogic 6.1

To configure a WebLogic 6.1 application server:

- 1 Locate the startup script used to start WebLogic for your domain. This file is typically located in a path similar to this example:

```
D:\bea\wlserver6.1\config\\start<Dom_Name>.cmd
```

Replace **<Dom_Name>** with the name of the script that starts your application.

For example, if your domain name is `petstore`, the path looks like this:

```
D:\bea\wlserver6.1\config\petstore\startPetStore.cmd
```

- 2 Create a backup copy of the startup script prior to making any changes to the script.
- 3 Use your editor to open the startup script.
- 4 Add the **Xbootclasspath** parameter to the Java command line that starts the application server. The parameter must be placed at the beginning of the Java parameters following any JIT options such as **-hotspot** or **-classic**.

Following is an example of the **Xbootclasspath** parameter:

```
<probe_install_dir>\classes\Sun\1.3.1_04;<probe_install_dir>\classes\boot
```

where **<probe_install_dir>** is with the path to the directory where the Probe was installed.

Following is an example of a WebLogic startup script before adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -classpath "%CLASSPATH%"
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer
-Dbea.home="C:\bea" -Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1\lib\weblogic.policy" weblogic.Server
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

Following is an example of a WebLogic startup script after adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m
-Xbootclasspath/p:"C:\Program Files\Mercury Interactive\common\Java-
aProbe\classes\boot"
-classpath "%CLASSPATH%"
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer
-Dbea.home="C:\bea" -Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1\lib\weblogic.policy" weblogic.Server
```

- 5 Save the changes to the startup script.
- 6 Restart the WebLogic application server. You do not need to reboot the host machine, just the application server.

Note: To verify that the Probe was configured correctly, check for entries in the **<probe_install_dir>/log/<probe_id>/probe.log** file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the Xbootclasspath correctly.

Note: For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

WebLogic 7.0

Note: For WebLogic 7.0 sp06 you must also configure JRockit as described in the section on configuring a WebLogic 8.1 application server for the JRockit JVM, on page 195.

To configure a WebLogic 7.0 application server:

- 1 Locate the startup script used to start WebLogic for your domain. This file is typically located in a path similar to this example:

```
D:\bea\weblogic700\server\config\\start<Dom_Name>.cmd
```

Replace **<Dom_Name>** by the name of the script that starts your application. For example, if your domain name is `petstore`, the path looks like this:

```
D:\bea\ weblogic700\server\config\petstore\startPetStore.cmd
```

- 2 Create a backup copy of the startup script prior to making any changes to the script.
- 3 Use your editor to open the startup script.

- 4 To change the script, add the following immediately before the call to **startWLS.cmd** (at the end of the file):

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

For Windows:

```
set JAVA_OPTIONS="-Xbootclasspath/p:<probe_install_dir>\
classes\Sun\1.3.1_04;<probe_install_dir>\classes\boot" %JAVA_OPTIONS%
@rem Call WebLogic Server
call "C:\bea7\weblogic700\server\bin\startWLS.cmd"
```

For UNIX:

```
export JAVA_OPTIONS="-Xbootclass-
path/p:<probe_install_dir>\classes\Sun\1.3.1_04;<probe_install_dir>\classes\boot"
$JAVA_OPTIONS
# Call WebLogic Server
call /bea/weblogic700/server/bin/startWLS.sh"
```

- 5 Save the changes to the startup script.
- 6 Restart the WebLogic application server. You do not need to reboot the host machine for your application server.
- 7 To verify that the Probe was configured correctly, check for entries in the **<probe_install_dir>/log/<probe_id>/probe.log** file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the Xbootclasspath correctly.

Note: For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

WebLogic 8.1

To configure a WebLogic 8.1 application server for the Sun JVM:

- 1 Locate the startup script used to start WebLogic for your domain. This file is typically located in a path similar to this example:

```
D:\bea\weblogic81\config\\start<Dom_Name>.cmd
```

Replace **<Dom_Name>** by the name of the script that starts your application.

For example, if your domain name is **petstore**, the path looks like this:

```
D:\bea\weblogic81\config\petstore\startPetStore.cmd
```

- 2 Create a backup copy of the startup script prior to making any changes to the script.
- 3 Use your editor to open the startup script.
- 4 Add the **Xbootclasspath** parameter to the Java command line that starts the application server. The parameter must be placed at the beginning of the Java parameters following any JIT options, such as **-hotspot** or **-classic**.

Following is an example of the **Xbootclasspath** parameter:

```
JAVA_OPTIONS="-Xbootclasspath/p:<probe_install_dir>  
             \classes\Sun\1.4.2_04;<probe_install_dir>\classes\boot"
```

where **<probe_install_dir>** is the path to the directory where the probe was installed.

Following is an example of a WebLogic startup script before adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -Xbootclass-  
path/p:<probe_install_dir>\classes\Sun\1.3.1_04;<probe_install_dir>\classes\boot" -  
classpath "%CLASSPATH%"  
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer  
-Dbea.home="C:\bea"  
-Dweblogic.management.password=%WLS_PW%  
-Dweblogic.ProductionModeEnabled=%STARTMODE%  
-Dcloudscape.system.home=./samples/eval/cloudscape/data  
-Djava.security.policy=="C:\bea\weblogic81/lib/weblogic.policy"  
weblogic.Server
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

Following is an example of a WebLogic startup script after adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -Xbootclasspath/p:"C:\Program  
Files\Mercury Interactive\common\JavaProbe\classes\boot"-classpath "%CLASS-  
PATH%"  
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer -Dbea.home="C:\bea"  
-Dweblogic.management.password=%WLS_PW%  
-Dweblogic.ProductionModeEnabled=%STARTMODE%  
-Dcloudscape.system.home=./samples/eval/cloudscape/data  
-Djava.security.policy=="C:\bea\weblogic81/lib/weblogic.policy" weblogic.Server
```

- 5 Save the changes to the startup script.
- 6 Restart the WebLogic application server. You do not need to restart the application server host machine.
- 7 To verify that the Probe was configured correctly, check for entries in the `<probe_install_dir>\log\<probe_id>\probe.log` file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter, or did not enter the Xbootclasspath correctly.

Note: For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

To configure a WebLogic 8.1 application server for the JRockit JVM:

- 1** Locate the command file that invokes the WebLogic application server (for example, `startWLS.cmd`). This file is typically located in a path similar to the following example:

```
C:\bea\weblogic81\server\bin\ startWLS.cmd
```

- 2** Create a backup copy of the command file prior to making any changes to the script. You may want to give the new copy a name such as `startWLSWithJRockit.cmd`, and use this as the new version of the command file that will be manipulated in the following steps.
- 3** Use your editor to open the startup script.
- 4** Set the JAVA executable invoked by WebLogic to JRockit.
 - a** Locate the line in the command file where the value of the `JAVA_VENDOR` parameter is set.
 - b** Change the value of the `JAVA_VENDOR` variable to point to the JRockit folder as follows:

```
set JAVA_VENDOR=<BEA_HOME_DIR>jrockit
```

For example:

```
set JAVA_VENDOR=BEA
```

5 Modify the Java command line that starts the application server.

- a** Locate the line in the command file which begins as follows:

```
%JAVA_HOME%\bin\java %JAVA_VM% %JAVA_OPTIONS% .....
```

- b** Indicate the JRockit management URL by specifying the **Xmanagement:class** parameter immediately following the %JAVA_OPTIONS% variable.

The following is an example of the **Xmanagement:class** parameter:

```
-Xmanagement:class=com.mercury.opal.capture.proxy.JRockitManagement
```

- c** Allow the Probe to hook into the application server process by adding the **Xbootclasspath** parameter immediately following the %JAVA_OPTIONS% variable.

Following is an example of the **Xbootclasspath** parameter:

```
-Xbootclasspath/p:<probe_install_dir>\classes\boot
```

where **<probe_install_dir>** is the path to the directory where the probe was installed.

The following is an example of a WebLogic startup script before adding the **Xmanagement:class** and **Xbootclasspath** parameters:

```
"%JAVA_HOME%\bin\java" %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS%  
-Dweblogic.Name=%SERVER_NAME%  
-Dweblogic.management.username=%WLS_USER%  
-Dweblogic.management.password=%WLS_PW%  
-Dweblogic.management.server=%ADMIN_URL%  
-Dweblogic.ProductionModeEnabled=%PRODUCTION_MODE%  
-Djava.security.policy="%WL_HOME%\server\lib\weblogic.policy" weblogic.Server
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

The following is an example of a WebLogic startup script after adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" %JAVA_VM% %MEM_ARGS% %JAVA_OPTIONS%  
-Xmanagement:class=com.mercury.opal.capture.proxy.JRokitManagement  
-Xbootclasspath/p:"C:\Program Files\Mercury Interactive\common\JavaProbe\classes\boot"  
-Dweblogic.Name=%SERVER_NAME%  
-Dweblogic.management.username=%WLS_USER%  
-Dweblogic.management.password=%WLS_PW%  
-Dweblogic.management.server=%ADMIN_URL%  
-Dweblogic.ProductionModeEnabled=%PRODUCTION_MODE%  
-Djava.security.policy="%WL_HOME%\server\lib\weblogic.policy" weblogic.Server
```

- 6** Save the changes to the command file.
- 7** Restart the WebLogic application server (not the computer; just the application server).
- 8** To verify that the Probe was configured correctly, check for entries in the **<probe_install_dir>\log\<probe_id>\probe.log** file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the **Xbootclasspath** parameter correctly.

Note: For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

Configuring the Oracle9i Application Server

Supported version	Operating systems
9.0.3	Solaris 8

Oracle9i application servers are configured by adding the **Xbootclasspath** property to the XML file used to start the application server. Because the files that Oracle9i provides are frequently customized by the site administrator, it is not possible to provide detailed configuration instructions that will apply exactly for each situation. Instead, the following sections provide instructions configuring an Oracle9i application server for a generic implementation. Your site administrator should be able to use these instructions to guide you through making these changes in your customized environment.

Note: Make sure that you understand the structure of the startup scripts, how the property values are set, and the use of environment variables before you make any configuration changes for the probe. Always create a backup copy of any file that you are going to update, prior to making the changes.

To configure an Oracle9i application server:

- 1** Locate the XML file that is used to control the configuration of the application server when the server is started. This file is typically located at **<Oracle 9iAS_Install_Dir>/opmn/conf/opmn.xml**.
- 2** Create a backup copy of the **opmn.xml** file prior to making any changes.
- 3** Open the **opmn.xml** file to be edited using your editor.

4 Add the `Xbootclasspath` property. The property must be added to the “java-option value.”

The following is an example of the `Xbootclasspath` parameter:

```
-Xbootclasspath/p:<probe_install_dir>\classes\Sun\1.4.2_04;
<probe_install_dir>\classes\boot
```

where `<probe_install_dir>` is the path to the directory where the probe was installed.

Note: When modifying the `-Xbootclasspath` parameter. Be sure to use quotes if there are spaces in the path that you specify.

Following is an example of an Oracle 9iAS startup script before adding the `Xbootclasspath` parameter:

```
- <ias-instance xmlns="http://www.oracle.com/ias-instance">
- <notification-server>
  <port local="6100" remote="6200" request="6003" />
  <log-file path="/opt/oracle/ora9ias/opmn/logs/ons.log" level="3" />
</notification-server>
- <process-manager>
  <start-mode mode="ssl" />
</ohs>
- <oc4j instanceName="home" numProcs="1" maxRetry="3">
  <config-file path="/opt/oracle/ora9ias/j2ee/home/config/server.xml" />
  <oc4j-option value="-properties" />
  <port ajp="3000-3100" rmi="3101-3200" jms="3201-3300" />
  <environment>
    <prop name="LD_LIBRARY_PATH" value="/opt/oracle/ora9ias/lib" />
  </environment>
</oc4j>
- <oc4j instanceName="OC4J_Demos" gid="OC4J_Demos">
  <config-file path="/opt/oracle/ora9ias/j2ee/OC4J_Demos/config/server.xml" />
  <java-option value="-Xmx512M" />
  <oc4j-option value="-userThreads -properties" />
  <port ajp="3001-3100" rmi="3101-3200" jms="3201-3300" />
  <environment>
    <prop name="%LIB_PATH_ENV%" value="%LIB_PATH_VALUE%" />
  </environment>
</oc4j>
- <custom gid="dcm-daemon" numProcs="1" noGidWildcard="true">
  <start path="/opt/oracle/ora9ias/dcm/bin/dcmctl daemon -logdir /opt/oracle/ora9ias/dcm/logs/daemon_logs" />
  <stop path="/opt/oracle/ora9ias/dcm/bin/dcmctl shutdowndaemon" />
</custom>
  <log-file path="/opt/oracle/ora9ias/opmn/logs/lpm.log" level="3" />
</process-manager>
</ias-instance>
```

Following is an example of an Oracle 9iAS startup script after adding the **Xbootclasspath** parameter:

```

- <ias-instance xmlns="http://www.oracle.com/ias-instance">
- <notification-server>
  <port local="6100" remote="6200" request="6003" />
  <log-file path="/opt/oracle/ora9ias/opmn/logs/ons.log" level="3" />
</notification-server>
- <process-manager>
- <ohs gid="HTTP Server" maxRetry="3">
  <start-mode mode="ssl" />
</ohs>
- <oc4j instanceName="home" numProcs="1" maxRetry="3">
  <config-file path="/opt/oracle/ora9ias/j2ee/home/config/server.xml" />
  <java-option value="-Xmx512m -Xbootclasspath/p:C:\Program
  Files\MercuryInteractive\common\JavaProbe\classes\boot" />
  <oc4j-option value="-properties" />
  <port ajp="3000-3100" mi="3101-3200" jms="3201-3300" />
  <environment />
</oc4j>
- <oc4j instanceName="OC4J_Demos" gid="OC4J_Demos">
  <config-file path="/opt/oracle/ora9ias/j2ee/OC4J_Demos/config/server.xml" />
  <oc4j-option value="-userThreads -properties" />
  <port ajp="3001-3100" mi="3101-3200" jms="3201-3300" />
  <environment>
  <prop name="%LIB_PATH_ENV%" value="%LIB_PATH_VALUE%" />
</environment>
</oc4j>
- <custom gid="dcm-daemon" numProcs="1" noGidWildcard="true">
  <start path="/opt/oracle/ora9ias/dcm/bin/dcmctl daemon -logdir /opt/oracle/ora9ias/dcm/logs/daemon_logs" />
  <stop path="/opt/oracle/ora9ias/dcm/bin/dcmctl shutdowndaemon" />
</custom>
  <log-file path="/opt/oracle/ora9ias/opmn/logs/ipm.log" level="3" />
</process-manager>
</ias-instance>

```

- 5 Save the changes to the XML file.
- 6 Restart the Oracle application server. You do not need to reboot the host for the application server.
- 7 To verify that the probe was configured correctly, check for entries in the **<probe_install_dir>\log\<probe_id>probe.log** file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the **Xbootclasspath** parameter correctly. For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

Configuring the JBoss Application Server

This section explains how to configure the JBoss 3.2.1 application server.

Supported versions	Operating systems
3.0.4	Windows 2000 Server
3.2.1	Linux Debian

You configure JBoss application servers by adding the **Xbootclasspath** property to the script that is used to start the application server. JBoss is started by running shell scripts in a UNIX environment, or command scripts in a Windows environment. Because the startup scripts that JBoss provides are frequently customized by the site administrator, it is not possible to provide detailed configuration instructions that will apply exactly for each situation. Instead, the following sections provide instructions for each of the certified versions of the JBoss application server for a generic implementation. Your site administrator should be able to use these instructions to guide you to make these changes in your customized environment.

Note: Make sure that you understand the structure of the startup scripts, how the property values are set, and the use of environment variables before you make any configuration changes for the probe. Always create a backup copy of any file that you are going to update prior to making the changes.

To configure a JBoss application server:

- 1 Locate the startup script that is used to start JBoss for your application. This file is typically located in path similar to the following example:

```
D:\JBoss\bin\run.bat
```

Note: For UNIX the file extension is “.sh”.

- 2 Create a backup copy of the startup script prior to making any changes to the script.
- 3 Open the startup script to be edited using your editor.
- 4 Add the **Xbootclasspath** parameter to the Java command line that starts the application server. The parameter must be placed at the beginning of the Java parameters following any JIT options, such as **-hotspot** or **-classic**.

The following is an example of the **Xbootclasspath** parameter:

```
-Xbootclasspath/p:<probe_install_dir>\classes\Sun\1.4.2_04;  
  <probe_install_dir>\classes\boot
```

where **<probe_install_dir>** is the path to the directory where the probe was installed.

Note: When modifying the **-Xbootclasspath** parameter. Be sure to use quotes if there are spaces in the path that you specify.

The following is an example of a JBoss startup script before adding the **Xbootclasspath** parameter:

```
"%JAVA%" %JAVA_OPTS% -classpath "%CLASSPATH%" org.jboss.Main %ARGS%
```

The following is an example of a JBoss startup script after adding the **Xbootclasspath** parameter:

```
"%JAVA%" "-Xbootclasspath/p:<probe_install_dir>\classes\Sun\1.4.2_04;  
          <probe_install_dir>\classes\boot" %JAVA_OPTS%  
-classpath "%CLASSPATH%" org.jboss.Main %ARGS
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

- 5** Save the changes to the startup script.
- 6** Restart the JBoss application server with the probe by running the modified script. You do not need to restart the application server host.
- 7** To verify that the probe was configured correctly, check for entries in the **<probe_install_dir>/log/<probe_id>/probe.log** file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the **Xbootclasspath** parameter correctly. For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

Configuring the SAP NetWeaver Application Server

The following instructions describe how to configure the SAP NetWeaver application server so that your applications can be monitored by the Probe.

Supported versions	Operating systems
3.0.4	Windows 2000 Server
3.2.1	Linux Debian

Configuring the NetWeaver application server means instrumenting the JVM and adding the **Xbootclasspath** property to the script that is used to start the application server. The following sections provide instructions for a generic NetWeaver implementation. Your site administrator should be able to use these instructions to guide you in making the changes that are appropriate to your specific environment.

Note: Make sure that you understand the structure of the startup scripts, how the property values are set, and the use of environment variables before you make any changes to the configuration of your application server for the Probe. You should always create a backup of files before making any changes.

To configure a SAP NetWeaver application server:

- 1** Run the JRE Instrumenter as described in Chapter 8, “Running the JRE Instrumenter.”
- 2** Add the JVM that runs the NetWeaver application server.
- 3** Instrument the JVM. The JRE Instrumenter provides the **Xbootclass** parameter. This parameter must be added to the NetWeaver Configtool’s JMV parameters window as described in the next step.
- 4** Run the NetWeaver application server configuration tool. The configuration tool is called **configtool.bat** and it is located in the **usr\sap\j2e\jc00\j2ee\configtool** directory.

- 5 Add the **-Xbootclass** parameter into the Java parameters text window. This window is in the General tab when you select your server instance. For example, cluster-data | instance_ID70323 | server_ID7032350.
- 6 Save your changes and exit the configuration tool.
- 7 Restart the NetWeaver application server.
- 8 To verify that the Probe was configured correctly, check for entries in the `<probe_install_dir>/log/<probe_id>/probe.log` file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the **Xbootclasspath** parameter correctly.

Note: If you detect capturing problems, assign the following values to these properties in the `etc/capture.properties` file:

minimum.buffer.size = 250000

initial.private.buffer.count = 50

maximum.private.buffer.count = 200

gentle.reserve.buffer.count = 50

hard.reserve.buffer.count = 50

Configuring a Generic Application Server

Note: You should only use the instructions for a generic application server when you do not find configuration instructions for your specific application server in this document.

Your site administrator may configure the application server using an alternative, site-specific method. In this case, the generic procedure may be sufficient for the administrator to understand what changes must be made.

Important: Before making any changes, back up all startup scripts.

To update the application server configuration:

- 1** Locate the application server startup script or the file where the JVM parameters are set.
- 2** Create a backup copy of the application server startup script before you make any changes to the script.
- 3** Use an editor or the application server console to open the startup script.
- 4** Add the **Xbootclasspath** parameter to the Java command line that starts the application server, using the following syntax:

```
-Xbootclasspath/p:<probe_install_dir>\classes\Sun\1.4.2_04;  
  <probe_install_dir>\classes\boot
```

where **<probe_install_dir>** is the path to the directory where the Probe was installed.

This connects the Probe to the application. The parameter must be placed at the beginning of the Java parameters, following any JIT options such as **-hotspot** or **-classic**.

Following is an example of a WebLogic Java command line in a startup script before adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -classpath "%CLASSPATH%"
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer -Dbea.home="C:\bea"
-Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

The following is an example of a WebLogic Java command line in a startup script after adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m
-Xbootclasspath/p:<probe_install_dir>\classes\Sun\1.4.2_04;
<probe_install_dir>\classes\boot"
-classpath "%CLASSPATH%"
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer
-Dbea.home="C:\bea" -Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server
```

- 5** Save the changes to the startup script.
- 6** Restart the application server under test. You do not need to restart the application server host machine.
- 7** To verify that the Probe was configured correctly, check for entries in the `<probe_install_dir>/log/<probe_id>/probe.log` file. If there are no entries in the file, this indicates that you did not run the JRE Instrumenter or did not enter the **Xbootclasspath** parameter correctly.

Note: Alternatively, you can configure the JVM process definitions of the application server by going into the Administrative Console. However, this does not apply to WebLogic servers.

For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.

Adjusting the Heap Size for the J2EE Probe in the Application Startup Script

The size of the heap can impact the performance of the J2EE Probe.

The default value for the heap size is 64 MB. The heap size is set in the application’s start up script using the following VM argument:

```
-Xmx<size>
```

You can increase the heap size by updating the value specified in the **-Xmx<size>** option. Refer to your JVM documentation for more help on setting this parameter.

Part IV

Setting Up Integration with Other Mercury Products

10

Setting Up Mercury Business Availability Center 6.x to Use Diagnostics

To begin using Diagnostics in Mercury Business Availability Center, you enter the Diagnostics Server details and configure the relevant Mercury Business Availability Center components to view Diagnostics data.

This chapter describes:	On page:
About Setting Up Mercury Business Availability Center 6.x to use Diagnostics	212
Specifying the Diagnostics Server Details	214
Enabling the Diagnostics View	220
Enabling Diagnostics Monitoring in Business Process Monitor (BPM) Related Views	222
Assigning Permissions for Diagnostics Users	224
Accessing the Diagnostics Pages in Windows 2003	226

About Setting Up Mercury Business Availability Center 6.x to use Diagnostics

The process involved in setting up Mercury Business Availability Center to use Mercury Diagnostics differs according to the version of Mercury Business Availability Center that you are using.

To set up Diagnostics in Mercury Business Availability Center 6.1:

1 Specify the Diagnostics Server details.

Enter the Diagnostics Server details in Mercury Business Availability Center. For more information, see “Specifying the Diagnostics Server Details” on page 214.

2 Enable the Diagnostics View.

Enable the Diagnostics View in Dashboard and in the IT Universe tab. For more information, see “Enabling the Diagnostics View” on page 220

3 Enable Diagnostics monitoring in Business Process Monitor related views.

Change the Business Process Monitoring (BPM) source template to enable Diagnostics monitoring in BPM related views. For more information, see “Enabling Diagnostics Monitoring in Business Process Monitor (BPM) Related Views” on page 222.

4 Assign relevant permissions.

Grant different permissions to different Diagnostics users. (This procedure is optional.) For more information, see “Assigning Permissions for Diagnostics Users” on page 224.

To set up Diagnostics in Mercury Business Availability Center 6.2:

1 Specify the Diagnostics Server details.

Enter the Diagnostics Server details in Mercury Business Availability Center. For more information, see “Specifying the Diagnostics Server Details” on page 214.

2 Assign relevant permissions.

Grant different permissions to different Diagnostics users. (This procedure is optional.) For more information, see “Assigning Permissions for Diagnostics Users” on page 224.

Specifying the Diagnostics Server Details

To make Mercury Diagnostics accessible from Mercury Business Availability Center, you specify the Diagnostics Server details.

Note: If you are using Windows 2003, you need to configure your Internet browser settings to access the Diagnostics Configuration page. For more details, see “Accessing the Diagnostics Pages in Windows 2003” on page 226

To specify the Diagnostics Server details in Mercury Business Availability Center:

- 1 Log on to Mercury Business Availability Center.
- 2 Select **Admin > Diagnostics** to open the Diagnostics Configuration page.

MERCURY™ APPLICATIONS ▾ ADMIN ▾ MY BAC HELP ▾ SITE MAP L

Business Availability Center - Diagnostics Configuration User: administrator

Diagnostics Server Details

Make sure that the Diagnostics Server is accessible from the Business Availability Center machine and from users' Web browsers through the values you enter in the fields below.
Note: The values defined here are needed for application links and data connections.

Enter Diagnostics Server details:

Diagnostics Server host name:

Diagnostics Server port number:

Diagnostics Server protocol:

Note: If you try to access Mercury Diagnostics (by clicking **Diagnostics** on the Site Map page or by clicking **Applications > Diagnostics**) before you have configured the Diagnostics Server, you receive a message instructing you to register the Diagnostics Server.



Click the link to open the Diagnostics Configuration page.

3 Enter the details for the Diagnostics Server in Commander mode.

- ▶ **Diagnostics server host name.** Enter the name of the machine that is host to the Diagnostics Server in Commander mode.

Note: Even if the Diagnostics Server is installed on the same machine as Mercury Business Availability Center, you still need to enter the actual name of the host in the **Diagnostics server host name** box. It is not sufficient to type **localhost** instead of the host name.

- ▶ **Diagnostics server port number.** Enter the port number used by the Diagnostics Server in Commander mode. The default port number is 2006.

- ▶ **Diagnostics server protocol.** Select the communication protocol through which Mercury Business Availability Center connects to **Diagnostics** - either **HTTP** or **HTTPS**.

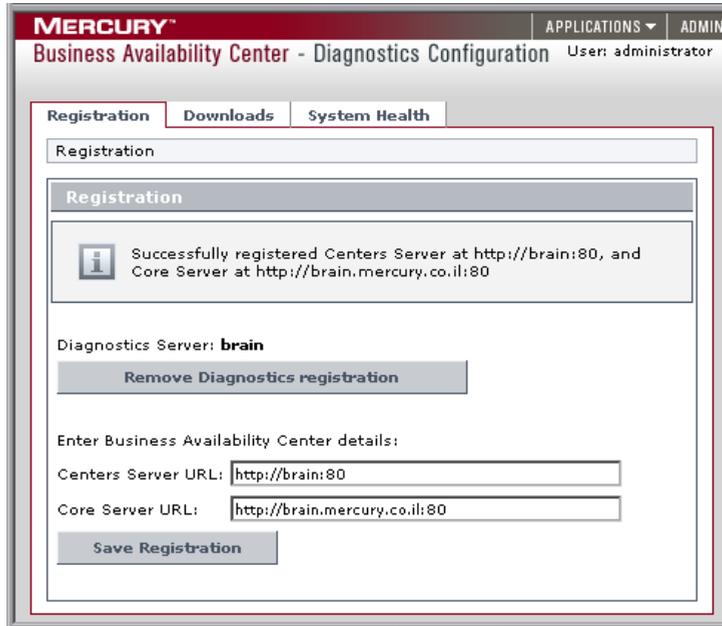
Note: If you select HTTPS as your communication protocol, additional configuration steps are required. For more information about the steps required, see Chapter 19, “Enabling HTTPS Between Diagnostics Components.”

- 4** After you have entered the Diagnostics Server details and ensured that these details are accurate, click **Submit** to complete the Diagnostics Server configuration process.

If the server name you have entered is incorrect or if the server is unavailable, an error message is displayed.

When you click **Submit**, the Mercury Diagnostics Server details are saved in Mercury Business Availability Center and the Mercury Business Availability Center server details are automatically registered on the Diagnostics Server machine.

The **Registration** tab in the Diagnostics Configuration page opens, displaying the Diagnostics Server details that you just specified and the Mercury Business Availability Center server details.



Where necessary, you can manually change the Mercury Business Availability Center server details in the **Enter Business Availability Center details** section of the Registration tab.

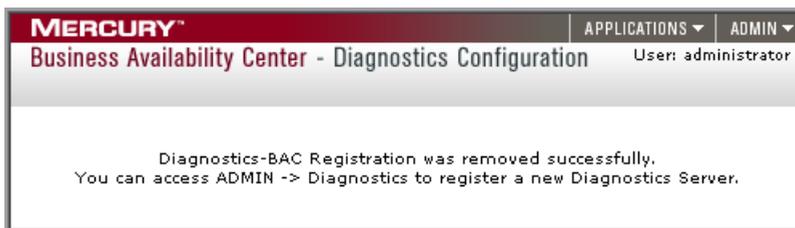
Note: Verify that the root URL in the **Centers Server URL** box, matches the root URL that you use to access Mercury Business Availability Center.

Changing the Diagnostics Server Details

You may want to change the Diagnostics Server details or remove the Diagnostics registration completely.

To remove the Diagnostics registration:

- 1** Select **Admin > Diagnostics**.
- 2** In the **Registration** tab, click **Remove Diagnostics Registration**.
- 3** In the message that opens, click **OK** to confirm that you want to remove the Diagnostics registration.
- 4** A message is displayed, confirming that you successfully removed the Diagnostics registration.



To register a new Diagnostics Server, select **Admin > Diagnostics** and follow the procedure explained in “Specifying the Diagnostics Server Details” on page 214.

Understanding the Diagnostics Configuration Page

You access the Diagnostics Configuration page by selecting **Admin > Diagnostics**. The Diagnostics Configuration page consists of the following three tabs:

Registration Tab

The Registration tab displays the following details:

- ▶ The Diagnostics Server details that you registered in Mercury Business Availability Center. To change these details, see “Changing the Diagnostics Server Details” on page 218.
- ▶ The Mercury Business Availability Center server details that were automatically registered on the Diagnostics Server machine. You can manually change the Mercury Business Availability Center server details in the **Enter BAC server details** section.

For information about registering Diagnostics in Mercury Business Availability Center for the first time, see “Specifying the Diagnostics Server Details” on page 214.

Downloads Tab

The Downloads tab provides links to the Mercury Diagnostics Probe installers, enabling you to download the Mercury Diagnostics Probe for your relevant platform.

If you did not specify the path to the Probe installers during the installation of the Mercury Diagnostics Server, the Downloads tab will not display any components. For more information, see Chapter 2, “Installing the Mercury Diagnostics Server.”

System Health Tab

Provides you with a map of all the components of your Mercury Diagnostics deployment and gives you real-time status and health information for each component. For detailed information about the System Health monitor, see Appendix C, “Using the System Health Monitor.”

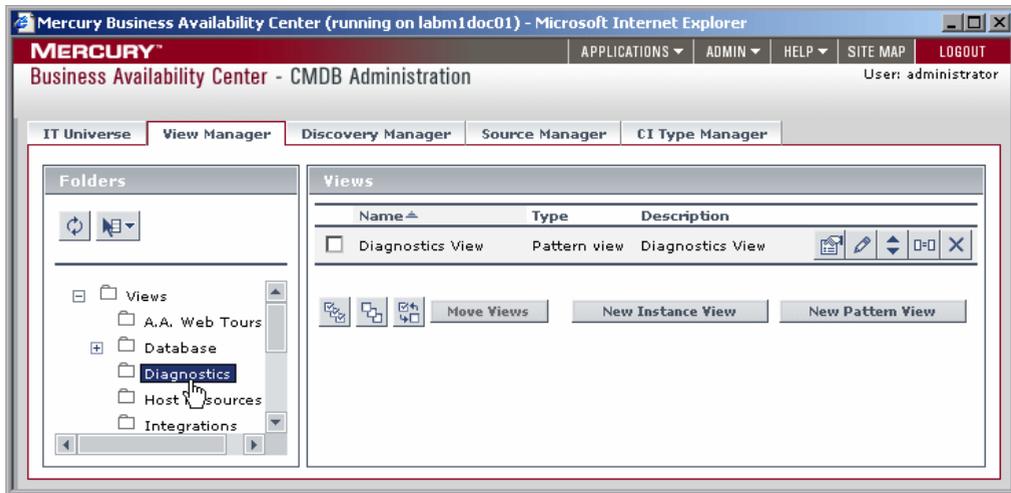
Enabling the Diagnostics View

Within Mercury Business Availability Center, both Dashboard and the IT Universe tab contain a *Diagnostics View* for working with Diagnostics data.

- ▶ In Mercury Business Availability Center 6.1, the Diagnostics View is not assigned by default and you therefore need to enable this view.
- ▶ In Mercury Business Availability Center 6.2, the Diagnostics View is enabled by default.

To enable the Diagnostics View in Mercury Business Availability Center:

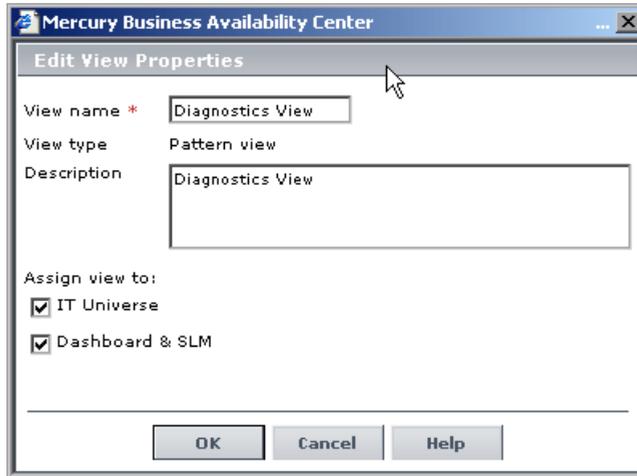
- 1 Select **Admin > CMDB** to open the CMDB Administration page.
- 2 Click the **View Manager** tab.
- 3 In the **Folders** pane (on the left), under the **Views** folder, click **Diagnostics**.





- 4 In the **Views** pane (on the right), click the **Edit view properties** button.

The Edit View Properties dialog box opens.



- 5 In the **Assign View to** section, select the **IT Universe** check box to enable the Diagnostics View in the IT Universe tab, and select the **Dashboard and SLM** check box to enable the Diagnostics View in Dashboard.

For more information about working with the Diagnostics View in Dashboard, see the *Mercury Diagnostics User's Guide*. For more information about working with Configuration Items in IT Universe, refer to *IT Universe Administration* in the *Mercury Business Availability Center Documentation Library*.

Enabling Diagnostics Monitoring in Business Process Monitor (BPM) Related Views

Important: This section does not apply to Mercury Business Availability Center 6.2, where Diagnostics monitoring in BPM related views is enabled by default.

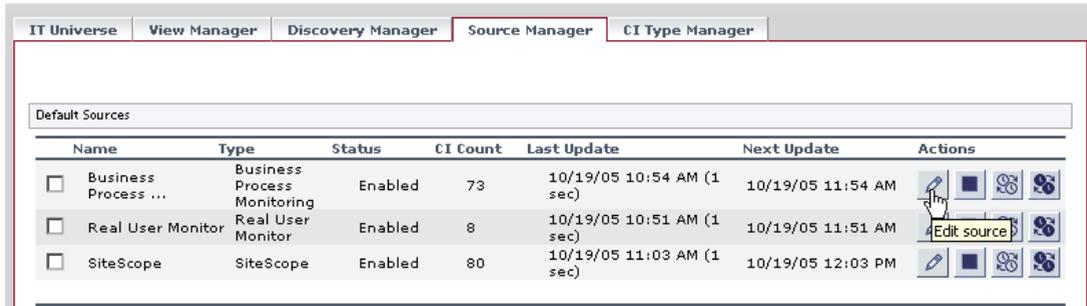
From Mercury Business Availability Center's Dashboard, you monitor the status of critical performance variables over time using Key Performance Indicators (KPIs). Diagnostics related KPIs are known as Application KPIs.

In Mercury Business Availability Center 6.1, you need to enable Application KPIs For BPM related views, before you can start using them to view the status of BPM transactions monitored by Diagnostics.

To enable Application KPIs in BPM related views, you replace the existing Business Process Monitoring source template with a modified template that has been provided with the Diagnostics product.

To enable Application KPIs in Business Process Monitor related views:

- 1 Select **Admin > CMDB** to open the CMDB Administration page.
- 2 Click the **Source Manager** tab to open the Source Manager.



Default Sources							
Name	Type	Status	CI Count	Last Update	Next Update	Actions	
<input type="checkbox"/> Business Process ...	Business Process Monitoring	Enabled	73	10/19/05 10:54 AM (1 sec)	10/19/05 11:54 AM		
<input type="checkbox"/> Real User Monitor	Real User Monitor	Enabled	8	10/19/05 10:51 AM (1 sec)	10/19/05 11:51 AM	Edit source	
<input type="checkbox"/> SiteScope	SiteScope	Enabled	80	10/19/05 11:03 AM (1 sec)	10/19/05 12:03 PM		



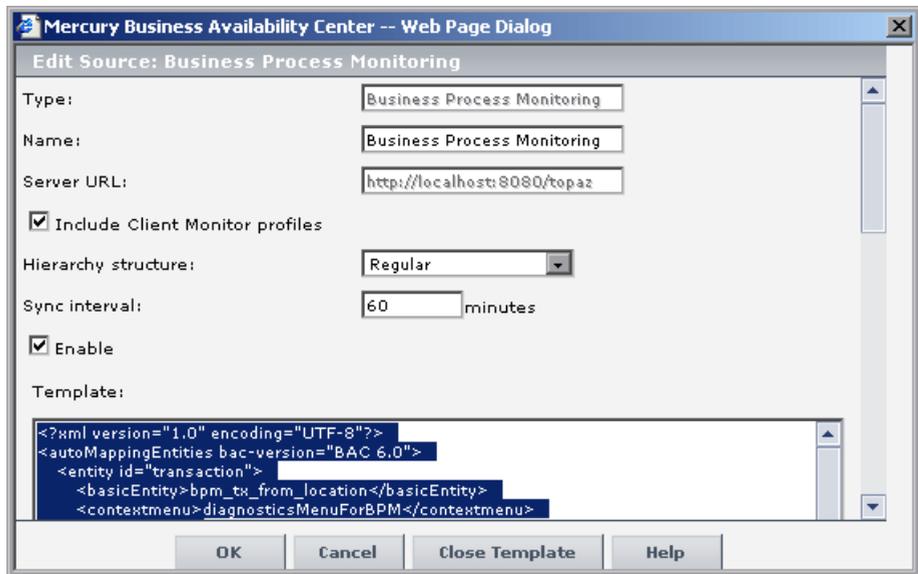
- 3 In the **Business Process Monitoring** row, click the **Edit Source** button to open the Edit Source dialog box.

- 4 Click the **Edit Template** button in the dialog box.

Mercury Business Availability Center displays a warning about making changes to the template. Click **OK**.

- 5 In the **Template** box, replace the existing template with the modified template as follows:

- a Using a text editor, open the **BPM_template_with_Application_KPI.xml** file located in the Mercury Diagnostics CD-ROM in the **BAC_Integration** directory, and copy the entire file to the clipboard.
- b In the **Template** box, delete all of the existing text and paste the new text that you copied from the modified template.



- c Click **OK**.

The new template is saved and Application KPIs will now be displayed, where relevant, in the BPM related views in Dashboard. This template modification procedure needs to be performed only once. Application KPIs are enabled for this session and for future sessions.

Assigning Permissions for Diagnostics Users

Mercury Business Availability Center enables you to apply permissions to users and user groups for specific resources that are defined in the system. There are specific types of Permission operations that administrators can grant Diagnostics users.

To assign permissions for Diagnostics Users in Mercury Business Availability Center:

- 1 Select **Admin > Platform** to open the Platform Administration page.
- 2 In the **Users and Permissions** tab, click **Permissions Management**.
- 3 In the left pane, in the **Select Context** list, select **Monitors**. The resource tree opens, displaying the resources included in the Monitors context.
- 4 In the resource tree, click **Diagnostics**.

The screenshot shows the Mercury Business Availability Center Platform Administration interface. The top navigation bar includes 'APPLICATIONS', 'ADMIN', 'HELP', 'SITE MAP', and 'LOGOUT'. The main title is 'Business Availability Center - Platform Administration' with the user 'administrator' logged in. The 'Users and Permissions' tab is active, showing the 'Resource: Diagnostics' configuration. The left pane shows the 'Select context' dropdown set to 'Monitors' and a tree view with 'Diagnostics' selected. The main area displays a list of users for the 'Diagnostics' resource, with 'fist_administrator_1' selected. Below this is a table for assigning permissions to 'fist_administrator_1'.

Operation	Grant	Granted from Group/Role/Parent	Inherit
Change	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
View	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Execute	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Full Control	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

An 'Apply Permissions' button is located at the bottom of the interface.

- 5 Select a user or user group in the user selection area on the upper right-hand area of the page.
- 6 In the **Operations** tab in the lower right-hand area of the page, select from the available operations for the highlighted resources.

Diagnostics administrators can grant users the following types of permission operations:

- ▶ **Change:** Enables viewing Diagnostics administration and configuring the Diagnostics settings.
 - ▶ **View:** Enables viewing the Diagnostics application when accessing Diagnostics from the Mercury Business Availability Center.
 - ▶ **Execute:** Enables changing the thresholds for Diagnostics measurements within the Mercury Diagnostics application.
 - ▶ **Full Control:** Enables performing all operations on Diagnostics, and granting and removing permissions for those operations.
- 7 Click **Apply Permissions** to complete the process.

For more detailed information about how to assign user permissions in Mercury Business Availability Center, refer to the section on “Configuring User Permissions” in *Platform Administration* in the *Mercury Business Availability Center Documentation Library*.

Accessing the Diagnostics Pages in Windows 2003

When your Internet browser is running in a Windows 2003 environment, you have to change your Internet browser settings in order to access the Diagnostics configuration and application pages in Mercury Business Availability Center.

To access the Diagnostics pages in a Windows 2003 environment:

- 1** In Internet Explorer, select **Tools > Internet Options** to open the Internet Options dialog box.
- 2** In the **Privacy** tab, under the **Web Sites** section, click **Edit** to open the Per Site Privacy Actions dialog box.
- 3** In the **Address of Web site** box, enter the name of the Diagnostics Server.
 - ▶ If you entered an IP address when you registered the Diagnostics Server in Mercury Business Availability Center, enter the IP address. If you entered a host name in Mercury Business Availability Center, enter the host name.
 - ▶ Include the `http://` or `https://` prefix, and the port number, as illustrated in the following example:
`http://<Diagnostics_server_host>:2006`
- 4** Click **Allow**. Click **OK** to close the Per Site Privacy Actions dialog box.
- 5** Click **OK** to close the Internet Options dialog box.

11

Installing LoadRunner 8.1 and the LoadRunner Diagnostics AddIn

To enable the Mercury Diagnostics functionality within LoadRunner, you must install the Mercury LoadRunner Diagnostics AddIn. You can only install the LoadRunner Diagnostics AddIn once you have already installed the Diagnostics Server in Commander mode and Mercury LoadRunner 8.1.

This chapter describes:	On page:
Understanding the LoadRunner Diagnostics AddIn	227
Installing Mercury LoadRunner	228
Installing the Mercury LoadRunner Diagnostics AddIn	228
Configuring LoadRunner for Mercury Diagnostics	230

Understanding the LoadRunner Diagnostics AddIn

The LoadRunner Diagnostics AddIn makes it possible for you to access Diagnostics functionality from within LoadRunner. Once the LoadRunner Diagnostics AddIn has been installed, you can configure LoadRunner to connect to the Diagnostics components, use the System Health Monitor to check on the status of the Diagnostics components, and use the Diagnostics components to gather performance metrics during your load tests.

Note: The version of the Diagnostics Server in Commander mode and the version of the LoadRunner Diagnostics AddIn must be exactly the same.

Installing Mercury LoadRunner

If you have not yet installed LoadRunner, you must install it before you can install the Mercury LoadRunner Diagnostics AddIn. For instructions on installing LoadRunner, refer to the *Mercury LoadRunner Installation Guide*.

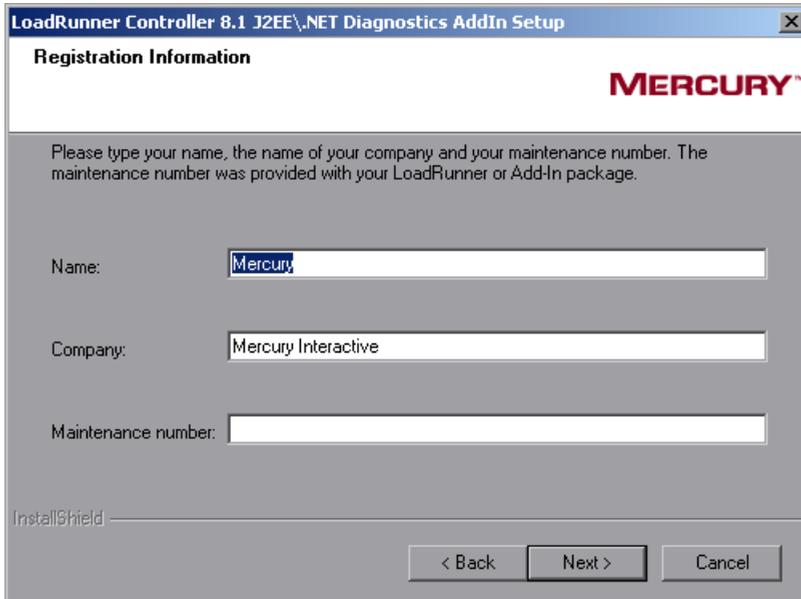
Installing the Mercury LoadRunner Diagnostics AddIn

The LoadRunner Diagnostics AddIn must be installed on the host machine for the LoadRunner Controller.

To install the LoadRunner Diagnostics AddIn:

Insert the Mercury Diagnostics Collectors CD into a CD-ROM drive. In the **LR_AddIn** directory, double-click **setup.exe** to start the InstallShield Wizard.

- 1 The software license agreement is displayed. Read the agreement and click **Yes** to accept it.
- 2 The Registration Information dialog box opens.

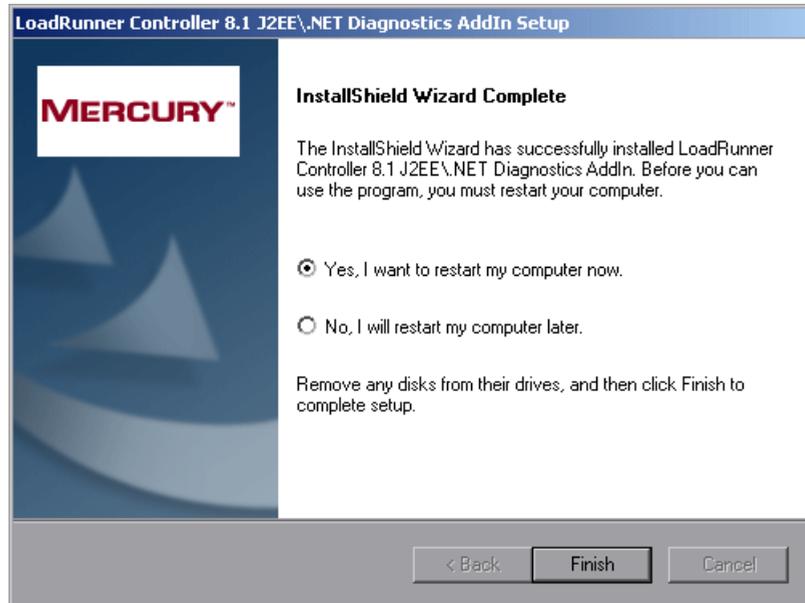


The screenshot shows a dialog box titled "LoadRunner Controller 8.1 J2EE\NET Diagnostics AddIn Setup" with a close button (X) in the top right corner. The dialog is titled "Registration Information" and features the Mercury logo in the top right. Below the title, there is a text prompt: "Please type your name, the name of your company and your maintenance number. The maintenance number was provided with your LoadRunner or Add-In package." There are three input fields: "Name:" with the text "Mercury" entered, "Company:" with the text "Mercury Interactive" entered, and "Maintenance number:" which is currently empty. At the bottom left, the text "InstallShield" is visible. At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

In the Registration Information dialog box, type your name, the name of your company, and your LoadRunner maintenance number. You can find the maintenance number in the maintenance pack shipped with LoadRunner.

Click **Next** to start the installation process. The installation process begins.

- 3 When the installation process is complete, the installation wizard displays a confirmation message.



To complete the installation process, choose **Yes, I want to restart my computer now** and click **Finish**.

Note: If you are installing the LoadRunner Diagnostics AddIn on a Windows XP machine with service pack 1 and Windows XP Hotfix Q328310 applied, you will receive an Application Error message for **ikernel.exe**. This message is issued because the Windows XP Hotfix Q328310 contains a Win32 API that does not execute as expected by the InstallShield engine. To resolve this problem, see the recommended solutions at the Java Technology Help web site, <http://java.com/en/download/help/ikernel.jsp>.

Configuring LoadRunner for Mercury Diagnostics

Before you can use the Diagnostics features from within LoadRunner you need to configure LoadRunner and provide the necessary information to enable communication with the Diagnostics components. See “Setting Up LoadRunner 8.1 to use Mercury Diagnostics” on page 231 for detailed instructions for configuring LoadRunner.

12

Setting Up LoadRunner 8.1 to use Mercury Diagnostics

This chapter describes how to configure LoadRunner 8.1 to enable Mercury Diagnostics for use in a load test.

This chapter describes:	On page:
About Setting Up LoadRunner to Use Mercury Diagnostics	231
Specifying the Diagnostics Server Details	232
Configuring LoadRunner Scenarios to use Mercury Diagnostics	234

About Setting Up LoadRunner to Use Mercury Diagnostics

LoadRunner and Mercury Diagnostics are integrated products that have been designed to work together to provide information to help you understand and improve the performance of your applications.

Before using Mercury Diagnostics with LoadRunner 8.1, you provide LoadRunner with the information that it needs in order to communicate with the Diagnostics components.

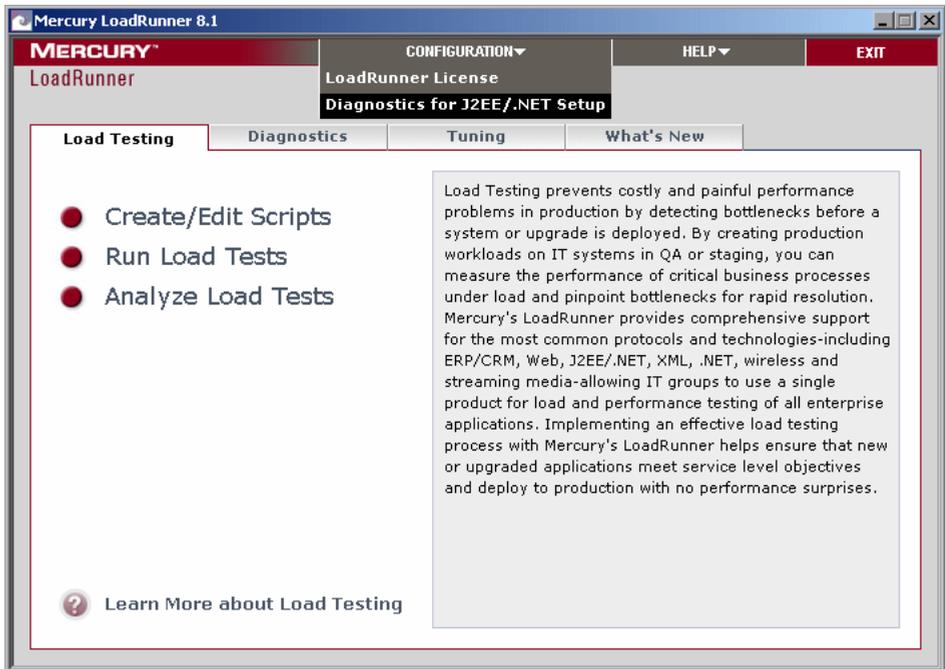
Specifying the Diagnostics Server Details

To make Mercury Diagnostics accessible from LoadRunner, you specify the Diagnostics Server details. You only need to specify these details the first time you use LoadRunner with Mercury Diagnostics.

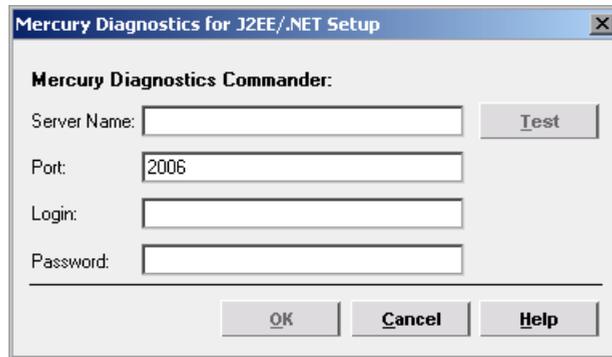
Note: Before specifying the Diagnostics Server details, ensure that the LoadRunner Controller is closed. When the Controller is open, you can view the Diagnostics configuration settings, but you cannot change them.

To specify the Diagnostics Server details in LoadRunner:

- 1 Select **Start > Programs > Mercury LoadRunner > LoadRunner** to open the Mercury LoadRunner launcher window.



- From the Mercury LoadRunner launcher window menu, select **Configuration > Diagnostics for J2EE/.NET Setup** to open the Diagnostics for J2EE/.NET Setup dialog box.



- Enter the details for the Diagnostics Server in Commander mode.
 - ▶ **Server Name.** Enter the name of the machine that is host to the Diagnostics Server in Commander mode.
 - ▶ **Port.** Enter the port number used by the Diagnostics Server in Commander mode. The default port number is 2006.

Note: LoadRunner does not support communication with the Diagnostics Server in Commander mode using HTTPS.

- ▶ **Login.** Enter the user name with which you log on to Mercury Diagnostics. The default user name is `admin`.
 - ▶ **Password.** Enter the password with which you log on to Mercury Diagnostics. The default password is `admin`.
- Click **Test** to verify that you entered the correct information for the Diagnostics Server in Commander mode and that there is connectivity between the Diagnostics Server in Commander mode and LoadRunner.
 - Click **OK** to complete the configuration process.

Configuring LoadRunner Scenarios to use Mercury Diagnostics

You configure your scenario for Diagnostics from the LoadRunner Controller for each load test scenario when you want to collect Diagnostics data.

To capture Diagnostics metrics in a load test scenario:

- 1** Configure the Diagnostics parameters for the scenario
- 2** Select the Probes that will be included in the scenario.

For more information, refer to the *Mercury Diagnostics User's Guide*.

13

Setting Up Performance Center 8.1 to Use Mercury Diagnostics

This chapter describes how to configure Performance Center 8.1 to enable Mercury Diagnostics for use in a load test.

This chapter describes:	On page:
About Setting Up Performance Center to Use Mercury Diagnostics	235
Specifying the Diagnostics Server Details	236
Configuring Performance Center Load Tests to Use Mercury Diagnostics	238

About Setting Up Performance Center to Use Mercury Diagnostics

Performance Center and Mercury Diagnostics are integrated products that have been designed to work together to provide information to help you understand and improve the performance of your applications.

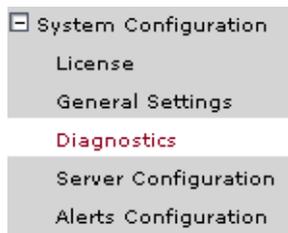
Before you can use Mercury Diagnostics with Performance Center 8.1, you provide Performance Center with the information that it needs in order to communicate with the Diagnostics components.

Specifying the Diagnostics Server Details

To make Mercury Diagnostics accessible from Performance Center, you specify the Diagnostics Server details. You only need to specify these details the first time you use Performance Center with Mercury Diagnostics. You provide this information on the Diagnostics page of the Performance Center Administration Site.

To specify the Diagnostics Server details in Performance Center:

- 1 Log on to the Performance Center Administration Site.
- 2 From the left navigation menu, choose **System Configuration > Diagnostics** to open the Diagnostics page.



- 3 In the **Setting J2EE/.NET Environment** section of the Diagnostics page, enter the details for the Diagnostics Server in Commander mode.

Setting J2EE/.Net Environment

J2EE/.Net Diagnostics Server Name

Server Name:

Port Number:

Login Name:

Password:

Use HTTPS:

- ▶ **Server Name.** Enter the name of the machine that is host to the Diagnostics Server in Commander mode.
- ▶ **Port Number.** Enter the port number used by the Diagnostics Server in Commander mode. The default port number is 2006.
- ▶ **Login Name.** Enter the user name with which you log on to Mercury Diagnostics. The default user name is **admin**.
- ▶ **Password.** Enter the password with which you log on to Mercury Diagnostics. The default password is **admin**.
- ▶ If Performance Center is communicating with Diagnostics through HTTPS, select the **HTTPS** box.

Note: If you select HTTPS as your communication protocol, additional configuration steps are required. For more information about the steps required, see Chapter 19, “Enabling HTTPS Between Diagnostics Components.”

- 4 Click **Save** to complete the configuration process.

Note: If you want to integrate Performance Center with a different Diagnostics Server in Commander mode, you need to update the above information.

Configuring Performance Center Load Tests to Use Mercury Diagnostics

You configure your load test for Diagnostics from the Performance Center User Site for each load test where you want to capture Diagnostics metrics.

To capture Diagnostics metrics in a load test :

- 1** Configure the Diagnostics parameters for the load test.
- 2** Select the Probes that will be included in the load test.

You provide this information on the Load Test Configuration page of the Performance Center User Site. For more information, refer to the *Mercury Diagnostics User's Guide*.

Part V

Advanced Component Configuration

14

Advanced Diagnostics Server Configuration

This chapter describes advanced configuration instructions for the Diagnostics Server that are intended for experienced users with in-depth knowledge of this product. Please use caution when modifying any of the Diagnostics component's properties.

This chapter describes:	On page:
Synchronizing Time Between Diagnostics Components	242
Configuring the Diagnostics Server for a Large Installation	247
Overriding the Default Diagnostics Server Host Name	250
Changing the Default Diagnostics Server Port	250
Configuring the Diagnostics Server for Multi-Homed Environments	252
Reducing Diagnostics Server Memory Usage	256
LoadRunner / Performance Center Diagnostics Server Assignments	257
Configuring Diagnostics Using the Diagnostics Server Configuration Pages	258
Configuring the Diagnostics Server for the LoadRunner Offline File	259

Synchronizing Time Between Diagnostics Components

In order for Diagnostics data to be stored and correlated properly, it is critical that time is synchronized between the Diagnostics components. To facilitate synchronization of data, the Diagnostics data is adjusted and saved to the synchronized GMT time of the Diagnostics Server in Commander mode. Synchronization makes it possible to display the data correctly for any local time in which the UI may be located.

The sections that follow describe how time synchronization works, and how to configure the components properly so that the time will be synchronized.

Understanding Time Synchronization

Time synchronization in Diagnostics begins with the Diagnostics Server in Commander mode determining the difference between its time and the GMT time provided by a designated **Time Source**. The **Time Source** to be used is set using the **timemanager.time_source** property in `<diagnostics_server_install_dir>/etc/server.properties`.

The valid values for the **timemanager.time_source** property are:

- **NTP**. Indicates that an NTP Server is to be used as the source of GMT time. This is the default value.

The NTP servers that are to be used are listed as values of the **timemanager.ntp.servers** property in `<diagnostics_server_install_dir>/etc/server.properties`.

Note: Ensure that one of the NTP servers in the list can be contacted from the Diagnostics Server, or add your local NTP server as the first server in the list.

- **BAC.** Indicates that the registered Mercury Business Availability Center core server is to be used as the source of GMT time.

Note: If Mercury Business Availability Center is configured to use Database time, you should also configure the Diagnostics Server in Commander mode to use this setting as the time source.

- **SERVER.** Indicates that the Diagnostics Server in Commander mode is to be used as the **Time Source**.

This should only be used when the Diagnostics Server is being used in Standalone mode.

The Diagnostics Servers that are in Mediator mode synchronize their time by establishing the time difference between the Diagnostics Server in Mediator mode and the Diagnostics Server in Commander mode.

If the Diagnostics Server in Commander mode has not yet synchronized with the **Time Source**, then the Diagnostics Servers in Mediator mode are considered to be “unsynched”. The Diagnostics Servers in Mediator mode that are unsynched will attempt to synchronize their time every 15 seconds until they succeed.

When a Diagnostics probe connects to a Diagnostics Server in Mediator mode or to a Diagnostics Server in Commander mode, the time difference is established between the Diagnostics Server and the probe.

If the probe attempts to connect to a Diagnostics Server that is still “unsynched,” then the probe connection is not allowed, and is dropped.

Since the data is stored based upon the GMT, differences in time zones or daylight savings times for the various components are not an issue. For example, the data that is displayed in the Diagnostics UI can be adjusted to display correctly for the time zone in which the UI is running.

Note: All data has been adjusted and saved to the synchronized GMT time of the Diagnostics Server in Commander mode. So, if the UI is running on a machine whose time has not been synchronized properly with the **Time Source**, the data displayed in the UI will appear shifted by the amount of time the UI machine is off from the synchronized GMT time.

Configuring the Time Synchronization on the Diagnostics Server

You synchronize the Diagnostics Server in Commanding mode by performing the procedure below.

Note: Time Synchronization settings for Diagnostics Servers in Mediator mode are ignored because their time is automatically synchronized with the Diagnostics Server in Commander mode.

To ensure that time on the Diagnostics Server in Commander mode can be synchronized:

- 1** The default configuration for the Diagnostics Server is set such that the value of the **timemanager.time_source** property in `<diagnostics_server_install_dir>/etc/server.properties` is **NTP**.

If the Diagnostics Server has an internet connection and the ability to connect to a server in the list of available NTP servers specified in the **timemanager.ntp.servers** property, then the default configuration will work and no changes are necessary.

Since Mercury Business Availability Center also uses NTP for time synchronization by default, this is the recommended setting.

- 2** If the Diagnostics Server does not have an internet connection or the ability to connect to the list of available NTP servers specified in **timemanager.ntp.servers** property, then you *must* do one of the following:
- ▶ Setup a local NTP server that can be contacted by the Diagnostics Server in Commander mode. List this local NTP server as the first entry in the **timemanager.ntp.servers** property in **<diagnostics_server_install_dir>/etc/server.properties**.

Note: It is recommended to have backup NTP servers in case the primary NTP server is not available.

- ▶ If you are using Diagnostics in a Mercury Business Availability Center or Mercury Managed Services environment, then you can set the **timemanager.time_source** property in **<diagnostics_server_install_dir>/etc/server.properties** to **BAC** to indicate Mercury Business Availability Center. This causes the Diagnostics Server to connect to the registered Mercury Business Availability Center core server to establish the time.

Note: To set up Mercury Business Availability Center to use Diagnostics, see Chapter 10, “Setting Up Mercury Business Availability Center 6.x to Use Diagnostics.”

- ▶ If the Diagnostics Server in Commander mode is to be used in Standalone mode, with no intention of using it with Mercury Business Availability Center, and you have no internet connection that will allow you to synchronize time with an NTP server, then you can set the **timemanager.time_source** property in `<diagnostics_server_install_dir>/etc/server.properties` to **SERVER**. This causes the Diagnostics Server to use its own time as the **Time Source**.

Note: It is recommended that you do not change the value of the **timemanager.time_source** property in `<diagnostics_server_install_dir>/etc/server.properties` once data has been captured and persisted using the designated **Time Source**. Changing the **Time Source** after data has been captured could result in a significant corruption to the data that has been captured and persisted. This is because the data that was persisted may have been captured while the Diagnostics Server was not synchronized with GMT. If the data that is captured later is captured while the Diagnostics Server is synchronized with GMT, it is possible that data could get re-aggregated multiple times or could get recorded into time buckets where it does not belong.

Configuring the Diagnostics Server for a Large Installation

If you are using a Diagnostics Server in Mediator mode with more than 20 probes, it is recommended that you make modifications to the default configuration of the Diagnostics Server.

Note: These changes to the configuration are not needed for the Diagnostics Server in Commander mode unless it also has probes assigned to it so that it also serves as a Diagnostics Server in Mediator mode.

Adjusting the Heap Size

The size of the heap can impact the performance of the Diagnostics Server in Mediator mode. If the heap is too small you may experience problems where the Diagnostics Server in Mediator mode “hangs” for periods of time. If the heap is too large it is possible that the Diagnostics Server in Mediator mode will experience long garbage collection delays.

The default value for the heap size is 512 MB. The heap size is set in the **server.nanny** file located at:

<diagnostics_server_install_dir>\nanny\windows\dat\nanny\
for Windows, or

<diagnostics_server_install_dir>/nanny/solaris/launch_service/dat/nanny/
for Solaris.

Use the following VM argument to set the size:

-Xmx???m

If you encounter problems with the Diagnostics Server in Mediator mode “hanging”, you can increase the heap size specified by updating the value specified in the **-Xmx???m** option.

To adjust the heap size of the Diagnostics Server in Mediator mode:

- 1** Determine the amount of heap that the Diagnostics Server in Mediator mode will need based upon the following table:

Number of Probes	Recommended Heap Size
0 - 10	512 MB
11 - 20	700 MB
20 - 30	1,400 MB

Note: The recommended heap size should not exceed more than 75% of the physical memory of the machine. If a machine has 1 G, then the heap size should not exceed 768 MB.

- 2** Open the **server.nanny** file that is to be edited. This file is located at:

<diagnostics_server_install_dir>\nanny\windows\dat\nanny\
for Windows, or

<diagnostics_server_install_dir>/nanny/solaris/launch_service/dat/nanny/
for Solaris.

- 3 On the `start_<platform>` line that is appropriate, replace the heap size specified in the `-Xmx???m` option with the optimal heap size that you calculated:

-Xmx???m

Continuing the previous example, the current heap size, represented by "???" is replaced with 768 MB.

-Xmx768m

Before you modify this line in the `server.nanny` file it will look like this:

```
start_nt="C:\MercuryDiagnostics\Server\jre\bin\javaw.exe" -server -Xmx512m
-Dsun.net.client.defaultReadTimeout=70000
-Dsun.net.client.defaultConnectTimeout=30000
"-javaagent:C:\MercuryDiagnostics\Server\probe\lib\probeagent.jar"
-classpath "C:\MercuryDiagnostics\Server\lib\mediator.jar;
C:\MercuryDiagnostics\Server\lib\loading.jar;
C:\MercuryDiagnostics\Server\lib\common.jar;
C:\MercuryDiagnostics\Server\lib\mercury_picocontainer-1.1.jar"
com.mercury.opal.mediator.util.DiagnosticsServer
```

After you modify this line in the `server.nanny` file it will look like this:

```
start_nt="C:\MercuryDiagnostics\Server\jre\bin\javaw.exe" -server -Xmx768m -
Dsun.net.client.defaultReadTimeout=70000
-Dsun.net.client.defaultConnectTimeout=30000
"-javaagent:C:\MercuryDiagnostics\Server\probe\lib\probeagent.jar"
-classpath "C:\MercuryDiagnostics\Server\lib\mediator.jar;
C:\MercuryDiagnostics\Server\lib\loading.jar;
C:\MercuryDiagnostics\Server\lib\common.jar;
C:\MercuryDiagnostics\Server\lib\mercury_picocontainer-1.1.jar"
com.mercury.opal.mediator.util.DiagnosticsServer
```

Overriding the Default Diagnostics Server Host Name

In situations where a firewall or NAT is in place, or where the host for the Diagnostics Server in Mediator mode has been configured as a multi-homed device, it may not be possible for the Diagnostics Server in Commander mode to communicate with the Diagnostics Server in Mediator mode using the host name that was assigned when the Diagnostics Server in Mediator mode was installed. The **registered_hostname** property allows you to override the default host name that the Diagnostics Server in Mediator mode uses to register itself with the Diagnostics Server in Commander mode.

To override the default host name for a Diagnostics Server in Mediator mode set the **registered_hostname** property located in `<diagnostics_server_install_dir>/etc/server.properties` to an alternate machine name or IP address that will allow the Diagnostics Server in Commander mode to communicate with the Diagnostics Server in Mediator mode.

Changing the Default Diagnostics Server Port

If the configuration of the Diagnostics Server host will not allow the default Diagnostics port to be used, you may choose a different port for the Diagnostics Server communications with the probes and other Diagnostics Servers.

If you decide to use an alternative port number after you have deployed Diagnostics, you must update the properties in the following table for each of the indicated components in your deployment with the new port number in order to ensure that the proper communications can take place.

Component Type	Properties
Diagnostics Server in Commander mode	<diagnostics_server_install_dir>/etc <ul style="list-style-type: none"> • webservice.properties - jetty.port • mediator.properties - commander.url • probe/etc/dispatcher.properties - registrar.url
Diagnostics Server in Mediator mode	<diagnostics_server_install_dir>/etc <ul style="list-style-type: none"> • mediator.properties - commander.url <diagnostics_server_install_dir>/probe/etc <ul style="list-style-type: none"> • dispatcher.properties - registrar.url
Probes	<probe_install_dir>/etc <ul style="list-style-type: none"> • dispatcher.properties - registrar.url

Configuring the Diagnostics Server for Multi-Homed Environments

The machines that host the Diagnostics Server can be configured with more than one Network Interface Card (NIC), and the Diagnostics Server process listens on all interfaces on its host. Some customer environments do not allow applications to listen on all network interfaces on a machine. If your environment has this restriction, use the following instructions to configure the Diagnostics Server to listen on specific network interfaces.

Setting the Event Host Name

If the Diagnostics Server host has multiple network interfaces, and you want to specify the hostname that the Diagnostics Server will listen on, you need to set the **event.hostname** property.

This property can be found in:

`<diagnostics_server_install_dir>/etc/server.properties`

Uncomment the property, **event.hostname**, and specify the hostname value.

By default, the **event.hostname** property is not set. This means that the Diagnostics Server will listen on all hostnames.

Modifying the jetty.xml File

The **jetty.xml** file has a section which defines the interfaces on which the Diagnostics Server is permitted to listen. By default, the **jetty.xml** file included with the Diagnostics Server has no listeners defined and the Diagnostics Server listens on all of the interfaces.

To configure the Diagnostics Server to listen on specific network interfaces on a machine:

- 1 Open `<diagnostics_server_install_dir>/etc/jetty.xml` and locate the following line:

```
<Configure class="org.mortbay.jetty.Server">
```

- 2 Add the following block of code after this line, changing the `<Set name="Host">.....</Set>` to contain the NIC's IP address.

```
<Call name="addListener">
  <Arg>
    <New class="org.mortbay.http.SocketListener">
      <Set name="Host">127.0.0.1</Set>
      <Set name="Port"><SystemProperty name="jetty.port" default="2006"/></Set>
      <Set name="MinThreads">1</Set>
      <Set name="MaxThreads">5</Set>
      <Set name="MaxIdleTimeMs">30000</Set>
      <Set name="LowResourcePersistTimeMs">5000</Set>
      <Set name="ConfidentialPort">8443</Set>
      <Set name="IntegralPort">8443</Set>
    </New>
  </Arg>
</Call>
```

- 3 Repeat the previous step adding a new copy of the block of code and setting the IP address for the NIC for each interface on which the Diagnostics Server is to listen.

Make sure that the `</Configure>` tag follows the listener code for the last NIC.

Note: Make sure that components that access the Diagnostics Server can resolve the hostnames of the Diagnostics Server to the IP address that you specify in the **jetty.xml** file for the host values. It is possible that some systems may resolve the host name to a different IP address on the Diagnostics Server host. For more information, see “Overriding the Default Diagnostics Server Host Name” on page 250.

Sample jetty.xml File

The following example shows the **jetty.xml** file for the Diagnostics Server, where the Diagnostics Server will listen on loopback and one IP address on the system.

```

<!-- Configure the Jetty Server -->
<!-- ===== -->
<Configure class="org.mortbay.jetty.Server">
<!-- ===== -->
<!-- Configure the Request Listeners -->
<!-- ===== -->
<Call name="addListener">
  <Arg>
    <New class="org.mortbay.http.SocketListener">
      <Set name="Host">127.0.0.1</Set>
      <Set name="Port"><SystemProperty name="jetty.port" default="2006"/></Set>
      <Set name="MinThreads">1</Set>
      <Set name="MaxThreads">5</Set>
      <Set name="MaxIdleTimeMs">30000</Set>
      <Set name="LowResourcePersistTimeMs">5000</Set>
      <Set name="ConfidentialPort">8443</Set>
      <Set name="IntegralPort">8443</Set>
    </New>
  </Arg>
</Call>

<-Listen on specific IP Address on this machine for incoming Commander calls->
<Call name="addListener">
  <Arg>
    <New class="org.mortbay.http.SocketListener">
      <Set name="Host">10.241.3.109</Set>
      <Set name="Port"><SystemProperty name="jetty.port" default="2006"/></Set>
      <Set name="MinThreads">1</Set>
      <Set name="MaxThreads">5</Set>
      <Set name="MaxIdleTimeMs">30000</Set>
      <Set name="LowResourcePersistTimeMs">5000</Set>
      <Set name="ConfidentialPort">8443</Set>
      <Set name="IntegralPort">8443</Set>
    </New>
  </Arg>
</Call>
</Configure>

```

Reducing Diagnostics Server Memory Usage

The Transaction Timeout Period is a safety mechanism that is used to prevent the Diagnostics Server from using excessive amounts of memory because it is holding on to old data for too long. The Diagnostics Server holds on to all of the information that it receives for a transaction until it receives the End of Transaction Notification (ELT), which tells the Diagnostics Server the transaction is complete. The timeout period for a transaction is reset each time the Diagnostics Server receives data for the transaction.

If the machine that the Diagnostics Server in Commander mode is running on is overloaded (CPU is heavily loaded or there are too many transactions per second for it to handle), or if there are network connectivity issues between the Load Generators or Mercury Business Availability Center and the Diagnostics Server in Commander mode, or between Business Process Monitor and Mercury Business Availability Center, the Diagnostics Server may not receive the ELT that lets it know when a transaction has ended. If the ELT is not received by the time the transaction timeout period expires, the Diagnostics Server assumes that the ELT is not coming and proceeds to process the data for the transaction and free up the memory that the transaction data is using.

The **correlation.txn.timeout** property sets the duration of the transaction timeout period. If you are experiencing out of memory conditions in the Diagnostics Server, you may want to reduce the transaction timeout period so that the Diagnostics Server waits less time for the end of a transaction. Use caution when adjusting the value of this property because multiple probes may be sending data to the Diagnostics Server, and a active transaction may be idle in one Diagnostics Server, so setting the value of this property too low could cause transactions to be reported incorrectly. It is recommended that if you need to reduce the value of this property, set it to 90 seconds more than the longest transaction in your test.

LoadRunner / Performance Center Diagnostics Server Assignments

Default Diagnostics Server Assignment

By default, a probe that is selected for a LoadRunner or Performance Center run will use the Diagnostics Server that is specified in its `<probe install directory>/etc/dynamic.properties`.

Diagnostics Server Assignment Override

It is possible to override the probe configuration when the probe is started for a run. To do so, you modify a mapping file on the Diagnostics Server in Commanding Mode, allowing you to override the Diagnostics Server assignment for a probe.

This can be useful when you are running Diagnostics in a combined LoadRunner / Performance Center and Mercury Business Availability Center environment. In such a case, you may want the probes to use different Diagnostic Servers when they are in a LoadRunner / Performance Center run than they do for Mercury Business Availability Center monitoring. You can do so by following these Diagnostics Server override instructions.

It may also be more convenient to use this mechanism than to edit the probe configuration file.

Note: When the probe is not in a run, it will always use the Diagnostics Server specified in its `<probe install directory>/etc/dynamic.properties file`.

To override the Diagnostics Server Assignment for a probe, modify the `server_assignment.properties` file in the `<diagnostics_server_install_dir>\etc` directory on the Diagnostics Server in Commander mode host machine.

The format of the **server_assignment.properties** file is:

<ProbeID> = <Server.id>

- ▶ Replace <ProbeID> with the ID of the probe.
- ▶ Replace <Server.id> with the ID of the Diagnostics Server.

The **server_assignment.properties** file is dynamically read at the start of each LoadRunner / Performance Center run. This means that changes made to this file become effective without having to restart the Diagnostics Server in Commander mode.

Configuring Diagnostics Using the Diagnostics Server Configuration Pages

The Diagnostics Server Configuration pages provide a way for you to set the property values that control how the Diagnostics Server communicates with the other Diagnostics components, and how it processes the data that it receives from the probes.

Note: To ensure that you are entering valid property values you should use these pages to update the Diagnostics Server properties rather than editing the property files directly.

For information about viewing and modifying Diagnostics using the Diagnostics Server Configuration pages, see Appendix A, “Diagnostics Server Administration Page.”

Configuring the Diagnostics Server for the LoadRunner Offline File

For each LoadRunner scenario or Performance Center test that is run, the Diagnostics Server in Mediator mode produces a file that is needed for LoadRunner Offline analysis which contains the J2EE data captured during the scenario. The size of this file can grow to be quite large. You must ensure that you have enough disk space to hold the LoadRunner Offline file on both the Diagnostics Server in Mediator mode host machine where the file is stored temporarily while the scenario is running and the Load Runner controller host machine where the file is stored when the scenario has ended.

Estimating the Size of the LoadRunner Offline File

Estimating the size of the offline file is highly dependent upon the data that is being captured and the rate at which the data is captured.

To estimate the size of the LoadRunner offline file:

- 1** Run a load test for 5 minutes and monitor the size of the offline file created by the Diagnostics Server in Mediator mode when Load Runner scenario is started.

Locate the offline file on the Diagnostics Server in Mediator mode host machine in `<diagnostics_server_install_dir>/data/<newest directory>`. The offline file has an extension of “.inuse”.

- 2** After 5 minutes note the size of the offline file.
- 3** Extrapolate the size of the offline file after one hour by multiplying the size of the offline file from the previous step by 12.
- 4** Determine the anticipated size of the offline file at the end of your load test by multiplying the one hour file size calculated in the previous step by the number of hours that you expect your actual load test to run.
- 5** Determine if the Diagnostics Server in Mediator mode host machine and the Controller host machine have enough disk space to accommodate the anticipated offline file size.

Reducing the Size of the LoadRunner Offline File

If you are concerned about the size of the offline file, you can reduce the file size by increasing the offline aggregation periods for the Diagnostics Server in Mediator mode. This will reduce the level of granularity in the offline data and thus reduce the size of the offline files.

The default settings for these properties are **5s** (5 seconds) which causes the Diagnostics Server in Mediator mode to aggregate all data into 5 second time slices. Increasing the value of these properties will make the offline file smaller because fewer data points are required to be stored when the aggregation period is longer. For example, increasing the offline aggregation period properties to **45s** should reduce the file size by roughly 50-75%.

Note: The impact on the size of the offline file size that will be achieved by adjusting the offline aggregation period is highly dependent upon the behavior of your application and the specifics of your load test.

Use the following steps modify the Diagnostics Server in Mediator mode offline aggregation period properties `bucket.lr.offline.duration` and `bucket.lr.offline.sr.duration` in `<diagnostics_server_install_dir>/etc/mediator.properties`.

To reduce the size of the offline files by increasing the Diagnostics Server in Mediator mode offline aggregation periods:

- 1** Ensure that the Diagnostics Server in Mediator mode is not participating in any active LoadRunner / Performance Center runs. This is necessary because the Diagnostics Server in Mediator mode must be restarted before the property changes described in the following steps will take effect.
- 2** Access the Mediator Configuration Page from your browser by navigating to the following URL:

`http://<diagnostics_server_hostname>:8081/configuration/Aggregation?level=60`

- 3** Increase the Offline VU Aggregation Period by increasing the setting for the **Load Runner / Performance Center Offline VU Aggregation Period** property. The value of this property must be a multiple of 5. For example set it to "45s".

- 4 Increase the Offline Server Request Aggregation Period by increasing the value of the **Load Runner / Performance Center Offline Server Request Aggregation Period** property. The value of this property must be a multiple of 5. For example set it to “45s”.
- 5 Update the Diagnostics Server in Mediator mode with the revised property values by clicking **Submit** at the bottom of the page.

A message will appear at the top of the page to indicate that the changes were saved along with a reminder to restart the Diagnostics Server in Mediator mode. As a further reminder the **Restart Mediator** button is displayed.

For more information on updating property values from the Configuration Page and a screen image showing the command buttons, see “Configuring Diagnostics Using the Diagnostics Server Configuration Pages” on page 437.

To cause the configuration changes to take effect restart the Diagnostics Server in Mediator mode by clicking **Restart Mediator**.

15

Advanced .NET Probe Configuration

This chapter provides instructions for configuring the .NET Probe.

This chapter describes:	On page:
Automatically Discovering ASP.NET Applications	264
Discovering the Classes and Methods in an Application	265
Customizing the Instrumentation for ASP.NET Applications	267
Restarting IIS	270
Enabling the Mercury Diagnostics Probe for .NET	271
Enabling the Mercury Diagnostics Profiler for .NET	272
Configuring Latency Trimming and Throttling	272
Configuring Depth Trimming	276
Configuring the .NET Probe for Light-Weight Memory Diagnostics	278
Disabling Logging	280
Overriding the Default Probe Host Machine Name	280
Modifying the .NET Probe Configuration File	281

Automatically Discovering ASP.NET Applications

The .NET Probe installer attempts to detect ASP.NET applications on the machine where the probe is installed. The installer configures the probe to capture basic ASP/ADO/MSMQ workload for each of the ASP.NET applications detected.

The .NET Probe installer discovers applications by inspecting the IIS configuration and looking for virtual directory entries that might refer to ASP.NET applications. In some instances, the ASP.NET applications may have been installed in a manner that prevents them from being detected. An example of a situation that could cause the installer to fail to detect an ASP.NET application is when an ASP.NET application has been installed as a Web directory instead of a virtual directory.

You can force the probe to rescan the IIS configuration if you have modified an existing ASP.NET application deployment or installed new ASP.NET applications. To force the probe to rescan the IIS configuration and update the **probe_config.xml** file, select **Start > Mercury Diagnostics .NET Probe > Rescan ASP.NET Applications**.

For each ASP.NET application that is automatically detected, the installer creates an appdomain reference in the **probe_config.xml** file located in the **<probe_install_dir>/etc** directory. Each of the appdomains referenced in the **probe_config.xml** file contains a reference to a points file template. The referenced points file is the application-specific points file template that the installer creates for each detected application. The points file defines the instrumentation that controls the workload that the probe will capture for each application. You must modify the instrumentation in the points file to enable the probe to capture application-specific custom business logic.

Discovering the Classes and Methods in an Application

If you want to monitor the performance of an application that you are not familiar with, you can use the Reflector automatic discovery tool that is installed with the .NET Probe to find the classes and methods in the application that you want to add to the probes instrumentation. The Reflector is located at `<probe_install_dir>\bin\reflector.exe`.

To discover classes and methods using Reflector:

- 1 Locate the installation directory for the application that you want to monitor.
- 2 Locate the folder in the application installation directory where the .dll files are stored.
- 3 Open a command prompt and change the directory to the folder where the .dll files for your application are stored.
- 4 Run the Reflector against all of the .dll files and .exe files in the current directory by executing the following the command at the command prompt:

```
<probe_install_dir>\bin\Reflector.exe
```

You can limit the Reflector to certain .dll and .exe files by adding additional parameters to the command. The above example is the same as entering the command:

```
<probe_install_dir>\bin\Reflector.exe *.dll *.exe
```

This command explicitly tells the Reflector to check all of the .dll and .exe files in the target directory.

To limit the Reflector to specific files only, you could enter the following:

```
<probe_install_dir>\bin\Reflector.exe WorkHorse.dll Utility.dll
```

This command explicitly tells the Reflector to check only the two .dll files specified.

The following example shows the commands that you might execute if you have an application called PetShop that has .dll files located in a bin folder:

```
C:\>cd "c:\Program Files\Microsoft\PetShop\Web\bin"

C:\Program Files\Microsoft\PetShop\Web\bin>
C:\MercuryDiagnostics\".NET Probe"\bin\Reflector.exe
```

- 5** The Reflector displays a report of the assemblies, namespaces, classes, and methods found in the .dll files that you specified.

```
-----
Assemblies:
-----
C:\Program Files\Microsoft\PetShop\web\bin\PetShop.BLL.dll
C:\Program Files\Microsoft\PetShop\web\bin\PetShop.DAL.dll
C:\Program Files\Microsoft\PetShop\web\bin\PetShop.web.dll
-----
Namespaces:
-----
(8 classes) PetShop.BLL
(6 classes) PetShop.DALFactory
(17 classes) PetShop.web
(11 classes) PetShop.web.Controls
(2 classes) PetShop.web.ProcessFlow
(1 classes) PetShop.web.webComponents
-----
(8 classes) Namespace: PetShop.BLL
-----
PetShop.BLL.Account (10 Methods)
  Equals          System.Boolean(System.Object)
  Finalize        System.Void()
  GetAddress      PetShop.Model.AddressInfo(System.String)
  GetHashCode     System.Int32()
  GetType        System.Type()
  Insert          System.Void(PetShop.Model.AccountInfo)
  MemberwiseClone System.Object()
  SignIn         PetShop.Model.AccountInfo(System.String,system.String)
  ToString       System.String()
  Update         System.Void(PetShop.Model.AccountInfo)

PetShop.BLL.Cart (17 Methods)
  Add             System.Void(System.String)
  Equals         System.Boolean(System.Object)
  Finalize       System.Void()
  get_Count     System.Int32()
  get_Item      PetShop.Model.CartItemInfo(System.Int32)
  get_Total     System.Decimal()
  GetCartItems System.Collections.ArrayList()
  GetEnumerator System.Collections.IEnumerator()
  GetHashCode   System.Int32()
  GetInStock   System.Int32(System.String)
  GetOrderLineItems System.Collections.ArrayList()
  GetType      System.Type()
  MemberwiseClone System.Object()
```

Note: You can redirect the output from the Reflector to a file, as shown in the following example:

```
<probe_install_dir>\bin\Reflector.exe sys*.dll > <report_name>.txt
```

The output from Reflector is redirected to the file that you specify.

- 6 Use the information in the report to customize the instrumentation for your application, as described in “Customizing the Instrumentation for ASP.NET Applications” on page 267.

Customizing the Instrumentation for ASP.NET Applications

When the .NET Probe is installed, the **ASP.NET.points** file is created with the standard instrumentation that is applied to all ASP.NET processing on the monitored server. You may create custom points files to capture performance metrics for the business logic that has been implemented through application-specific classes. If you do not know which classes and methods in the application you want to monitor, you may use the Reflector tool installed with the probe to analyze the .dll files in the application to discover the classes and methods. See “Discovering the Classes and Methods in an Application” on page 265.

To let the .NET Probe know that you want to apply the custom instrumentation to an application, you must update the **probe_config.xml** file so that the customized points file is associated with the application.

To associate a customized points file with an application:

- 1** Create a points file with the instrumentation for the application specific classes. To create a points file, copy an existing points file in the `<probe_install_dir>/etc` folder.

Note: If the application was auto-detected during the installation or during a rescan of IIS, a points file template will already exist for the application.

- 2** Customize the points file so that the probe captures custom business logic for your applications.

The following example illustrates how to modify the points file so that the probe captures IBuySpy custom code:

```
[IBuySpy Callee]
class = !IBuySpy.*
method = !.*
signature =
scope =
ignoreScope =
layer = Custom.IBuySpy
```

- 3** Add an `<appdomain>` tag for the application, and its `<points>` tag to the ASP.NET `<process>` tag in the `probe_config.xml` file which is located in the `<probe_install_dir>/etc` directory.

```
<appdomain name="your app name">
  <points file="your app.points"/>
</appdomain>
```

The example below illustrates this step. A custom points file has been created for the MSPetsShop application. The file has been named **MSPetShop.points**. The **<appdomain>** tag for the application, and the points file have been added to the ASP.NET **<process>** tag in the **probe_config.xml** file.

```
<?xml version="1.0" encoding="utf-8"?>

<probeconfig>
  <id probeid="" probegroup="Umatilla"/>

  <credentials username="" password=""/>
  <profiler authenticate=""><authentication username="" password=""/></profiler>

  <diagnosticserver url="http://issaquah:2006"/>
  <mediator host="issaquah" port="2612"/>
  <webserver start="35000" end="35100"/>
  <modes am="true"/>

  <instrumentation><logging level="" threadids="no"/></instrumentation>

  <lwmd enabled="true" sample="1m" autobaseline="1h" growth="10" size="10"/>

  <process name="ASP.NET">
    <logging level=""/>
    <points file="ASP.NET.points"/>
    <appdomain name="MSPetShop">
      <points file="MSPetShop.points"/>
    </appdomain>
  </process>
</probeconfig>
```

4 Restart IIS as instructed in “Restarting IIS” below.

Restarting IIS

After you have modified the instrumentation and configuration for the probe, you must restart IIS. To restart IIS from the command line or from the **Start > Run** menu, type **iisreset** and press ENTER. This command restarts the Web publishing service but does not immediately start the probe.

Your applications will be instrumented, and the probe will register with the Diagnostics Server, the next time that a Web page is requested for the application.

If you are running the probe in a Production environment (with Mercury Business Availability Center or Diagnostics Standalone), or in a Test environment (with LoadRunner or Performance Center), you can verify that the probe is connected by opening a browser and browsing to <http://<Diagnostics Server Host>:<Commander Port>/registrar>. Select **View > Edit Registered Components** and look for your application listed in the **Name** column.

Note: ASP.NET is designed to restart applications under various circumstances, including when it has detected that applications have been redeployed, or when applications are exceeding the configured resource thresholds.

When ASP.NET restarts an application that is being monitored by a .NET Probe, the probe is deactivated and a new probe is started. While this is occurring, there may be a period of overlap where there are multiple probes simultaneously registered with the Diagnostics Server in Commander mode and connected to the Diagnostics Server in Mediator mode. This condition could cause LoadRunner / Performance Center and Mercury Business Availability Center to report errors during the application restart sequence.

Enabling the Mercury Diagnostics Probe for .NET

The .NET Probe is enabled when it is installed. After you restart your Web server, the .NET Probe will begin to gather performance information from the moment a URL in your application has been accessed.

You may disable the .NET Probe so that it will not start and not gather performance metrics.

To disable a .NET Probe:

Select **Start > Mercury Diagnostics .NET Probe > Disable Mercury .NET Probe**.

To enable a .NET Probe that was disabled:

Select **Start > Mercury Diagnostics .NET Probe > Enable Mercury .NET Probe**.

Note: Disabling the probe only disables probe metrics. It does not disable system metrics. The process of enabling or disabling system metrics is controlled through the standard Windows services manager.

Enabling the Mercury Diagnostics Profiler for .NET

When the .NET Probe is installed to work with a Mercury Diagnostics Server, the Mercury Diagnostics Profiler for .NET is disabled by default. You must configure the probe to enable the .NET Diagnostics Profiler.

To enable the .NET Diagnostics Profiler for a probe:

Add the **pro** attribute and set it to **true**. The **pro** attribute is located in the `<modes>` tag in the `probe_config.xml` file, which is located in the `<probe_install_dir>/etc` directory.

```
<probeconfig>
...
  <modes enterprise="true" pro="true"/>
...
</probeconfig>
```

Configuring Latency Trimming and Throttling

When the .NET Probe determines that it is running out of resources because the Diagnostics Server is not keeping up with the amount of data that the probe is capturing, it can automatically reduce the number of methods that it captures using a process called *latency trimming*. By default, latency trimming is disabled so that all methods are captured.

When latency trimming is enabled, the .NET Probe trims the number of methods captured by ignoring methods with a latency below a certain minimum latency threshold. The idea behind trimming is that it is better to miss capturing methods with lower latency that are less likely to be of interest so that the probe will be able to continue to run, and be able to capture the more interesting methods with higher latencies.

Enabling Latency Trimming

By default, trimming is disabled for the .NET Probe. You must change the probe configuration in order to enable trimming.

To enable Latency Trimming:

Add the **latency** element to the `<probe_install_dir>/etc/probe_config.xml` configuration file, as shown in the following example:

```
<trim>  
    <latency enabled="true" />  
</trim>
```

The attribute of the latency element that turns on latency trimming is **enabled**. Latency trimming is enabled when **enabled** is set to true. When **enabled** attribute is set to false, latency trimming is disabled. The default value for this attribute is false.

For a description of attributes and elements of the **latency** element, see “<latency> element”, on page 299.

Note: Because of threading and buffering, partial information about a method that was trimmed may be transmitted to the Diagnostics Server. When the Diagnostics Server detects that it received only partial information for a method, it issues a warning message. You should ignore these warning messages unless you expected that the information for all methods was to be captured.

Setting Latency Trimming Thresholds

By default, the latency trimming thresholds are set so that those methods with a latency less than 2 ms will be trimmed, and those methods with a latency greater than 100 ms will never be trimmed.

You can set the minimum trimming threshold by adjusting the value of the **min** attribute. You can set the maximum trimming threshold by adjusting the value of the **max** attribute. These attributes are specified in the **latency** element in the `<probe_install_dir>/etc/probe_config.xml` configuration file.

```
<trim>
  <latency enabled="true" min="50" max="100" />
</trim>
```

The attributes of the latency element that control the trimming thresholds are:

► **min**

Sets the minimum latency threshold. When latency trimming is enabled, methods with a latency less than or equal to the value of this attribute are trimmed. If you do not specify a value for this attribute, the default value of 2 ms is used.

The lower the value of the **min** attribute the greater the chance that the performance of your application will be adversely impacted. A lower value means that fewer methods will be trimmed because more low-latency methods will be captured.

If the information for all methods must be captured, disable latency trimming by setting **latency enabled** equal to false.

► **max**

Sets the maximum latency threshold. When latency trimming is enabled, methods with a latency greater than or equal to the value of this attribute will never be trimmed. The default value for this attribute, if you do not specify a value, is 100 ms.

For a description of the attributes and elements of the **latency** element, see “<latency> element” on page 299.

Configuring Latency Trimming Throttling

Latency trimming is not throttled by default. Without throttling, the methods that fall below the minimum method latency threshold are always trimmed.

When throttling is turned on, the trimming is automatically adjusted based on the percentage of the probe resources that are being used up by the Diagnostics Server processing backlog. Initially, trimming starts without any throttling, and methods that fall below the minimum method latency threshold are trimmed.

If the percentage of probe resources used increases above a set throttling increment threshold, the effective trimming threshold is incremented so that methods with higher latency will be trimmed. If the percentage of probe resources used increases above the threshold again, the effective trimming threshold is incremented once more so that methods with even higher latency will be trimmed. If the percentage of probe resources used drops below the throttling decrement threshold, the effective trimming threshold is decremented so that the methods with lower latencies will once again be captured.

The effective trimming threshold cannot be incremented above the maximum method latency threshold, and it cannot be decremented below the minimum method latency threshold.

Below is an example of the **latency** element in the `<probe_install_dir>/etc/probe_config.xml` configuration file that includes the throttling attributes:

```
<trim>
  <latency enabled="true" min="50" max="100" >
  <throttle="true" incrementthreshold="75">
  <decrementthreshold="50" increment="2"/>
</trim>
```

The attributes of the **latency** element that control throttling are:

► **throttle**

Throttling is enabled when this attribute is set to **true**. When this attribute is set to **false** throttling is disabled. The default value for this attribute is **false**.

► **increment**

Sets the amount that the effective trimming threshold will be incremented when the percentage of probe resources used exceeds the **incrementthreshold**. Sets the amount that the effective trimming threshold will be decremented when the **decrementthreshold** is crossed. The default value for this attribute is 2 ms.

► **incrementthreshold**

When the percentage of probe resource usage rises to the value of this attribute or higher, throttling is triggered so that the effective trimming threshold is incremented. The default value for this attribute is 75 percent.

► **decrementthreshold**

When the percentage of probe resource usage falls to the value of this attribute or lower, throttling is triggered so that the effective trimming threshold is decremented. The default value for this attribute is 50 percent.

For a description of the attributes and elements of the **latency** element, see “<latency> element” on page 299.

Configuring Depth Trimming

The .NET Probe can automatically reduce the number of methods that it captures using a process called *depth trimming*. When the Diagnostics Server is not keeping up with the amount of data that the probe is capturing, the probe can use depth trimming to help prevent it from running out of resources. By default, depth trimming is disabled.

When depth trimming is enabled, the .NET Probe trims the number of methods captured by ignoring methods that are called at a stack depth that is greater than the maximum stack depth threshold. Those that are called at a stack depth less than or equal to the stack depth threshold will be captured. The idea behind trimming is that it is better to miss capturing methods further down in the call stack, that are less likely to be of interest, so that the probe will be able to continue to run and be able to capture the more interesting methods that occur higher in the call stack.

For example, if the stack depth threshold is 3, and the following method calls are made:

```
/login.do calls a() calls b() calls c()
```

only the /login.do, a, and b methods will be captured, and method c will be trimmed.

Below is an example of the **depth** element in the `<probe_install_dir>/etc/probe_config.xml` configuration file that includes the trimming attributes:

```
<trim>
  <depth enabled="true" depth="10" />
</trim>
```

The attributes of the **depth** element that control trimming are:

► **enabled**

Depth trimming is enabled when this attribute is set to true. When this attribute is set to false depth trimming is disabled. The default value for this attribute is false.

► **depth**

Sets the threshold that will be used for depth trimming. Methods that are called at or below the value of this attribute are trimmed when depth trimming has been enabled. The default value for this attribute is 25.

Setting **depth** to a lower value can significantly reduce the overhead of capture. For a description of the attributes and elements of the **depth** element, see “<depth> element” on page 298.

Configuring the .NET Probe for Light-Weight Memory Diagnostics

The .NET Probe has three different levels of Light-Weight Memory Diagnostics (LWMD):

- ▶ LWMD Off (Do not capture heap or collections metrics)
- ▶ Capture heap metrics only
- ▶ Capture heap metrics and collections metrics

When the .NET Probe is installed, the default configuration for the .NET Probe is to have Light-Weight Memory Diagnostics (LWMD) turned on only for the capture of heap metrics.

If you want to enable LWMD to capture collection metrics, you must change the probe's configuration using the instructions below.

Note: Enabling the probe to capture collections metrics may incur additional overhead on the host for your application.

To enable the capture of Collections metrics:

- 1 Uncomment the following lines in the **ASP.NET.points** file for ASP.NET applications:

Note: For non-ASP.NET applications you will need to add the lines specified below to the appropriate points file.

Before:

```
;[LWMD]  
;keyWord = lwmd
```

After:

```
[LWMD]
keyWord = lwmd
```

- 2 Verify the setting of **lwmd enabled** element in `<probe_install_dir>/etc/probe_config.xml`. This is the element that controls the collection of heap metrics.

```
<lwmd enabled="true"/>
```

To disable LWMD:

- 1 Comment out the following lines in the **ASP.NET.points** file for ASP.NET applications:

Note: For non ASP.NET applications you will need to comment out the lines specified below in the appropriate points file.

Before:

```
[LWMD]
keyWord = lwmd
```

After:

```
:[LWMD]
:keyWord = lwmd
```

- 2 Set the **lwmd enabled** element in `<probe_install_dir>/etc/probe_config.xml` to false:

```
<lwmd enabled="false"/>
```

Disabling Logging

You can disable probe application logging by changing the **logging level** tag of the ASP.NET process section of the **probe_config.xml** file, as shown in the following example:

```
<process name="ASP.NET">  
    <logging level="off"/>  
</process>
```

You can disable probe instrumentation logging by changing the **logging level** tag of the instrumentation section, as shown in the following example:

```
<instrumentation>  
    <logging level="off" />  
</instrumentation>
```

Overriding the Default Probe Host Machine Name

The **registered_hostname** property allows you to override the default host machine name that the probe uses to register itself with the Diagnostics Server in Commander mode. In situations where a firewall or NAT is in place or where your probe host machine has been configured as a multi-homed device, it may not be possible for the Diagnostics Server in Commander mode to communicate with the probe unless you override the default host machine name.

To override the default host machine name for a probe, set the **registered_hostname** attribute, located in the .NET Probe **<registrar>** tag of the **probe_config.xml** file, to an alternate machine name or IP address that will allow the Diagnostics Server in Commander mode to communicate with the probe. Below is an example of the override:

```
<registrar url="http://foo:2006/registrar" registered_hostname="bar"/>
```

Note: Setting the `registered_hostname` attribute because of a NAT or firewall is only an issue for a test environment where you are using LoadRunner, Performance Center, or Diagnostics Standalone.

However, if you should set the `registered_hostname` in a production environment where you are using Mercury Business Availability Center or Diagnostics Standalone, the name that you specify will be shown as the host name in System Health.

Modifying the .NET Probe Configuration File

You control the configuration of the .NET Probe by modifying the elements and attributes in the configuration file:

`<probe_install_dir>/etc/probe_config.xml`.

The sections that follow describe the elements and attributes that make up the configuration file, and how you can use them to configure your probe.

<probeconfig> element

Purpose

Provides single containing root element for the probe configuration.

Attributes

None.

Elements

Number of Occurrences	1
Parent Elements	None
Child Elements	appdomain, bufferpool, credentials, diagnosticserver, eventserverhost, id, instrumentation, ipaddress, logging, lwmd, modes, points, process, profiler, sample, trim, webserver

<id> element**Purpose**

Provides probe id and probe group id.

Attributes

Attribute	Valid Values	Default	Description
probeid	Letters, digits, underscore, dash, period	\$(AppDomain).NET	Names the probe used by LoadRunner / Performance Center and System Health
probegroup		Default	Defines the grouping for probes used for reporting metrics.

Elements

Number of Occurrences	1 per parent
Parent Elements	probeconfig
Child Elements	none

<credentials> element**Purpose**

Supplies credentials that are used to validate for communication with the Diagnostics Server.

Attributes

Attributes	Valid Values	Default	Description
username		none	Used to validate with the Diagnostics Server
password		none	Used to validate with the Diagnostics Server
authenticate	true, false	true	Enables and disables authentication. This should be disabled for backward compatibility with releases prior to D4

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<filter> element**Purpose**

Filters out certain metrics that would skew the results or not be representative of the processing being monitored.

Attributes

Attributes	Valid Values	Default	Description
firstserverrequest	true, false	false	Enables/disables skipping the collection of metrics for the first time a particular server requests (URL) gets run after application startup.

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<profiler> element**Purpose**

Contains settings for the Profiler feature of the probe.

Attributes

Attributes	Valid Values	Default	Description
authenticate	true, false	none	Enables/Disables authentication of incoming Profiler connection requests
register	true, false	false	Tells the probe to register even if it is in Profiler only mode
samples	number	60	Tells the Profiler how many samples to keep for lwmd/heap trending
best	number	1	The number of fastest instance trees to keeps
worst	number	3	The number of slowest instance trees to keeps

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	authentication

<authentication> element

Purpose

List of authenticated user names and passwords.

Attributes

Attributes	Valid Values	Default	Description
username		admin	
password		admin	Passwords must be generated using the passgen utility in the <probe_install_dir>\bin directory

Elements

Number of Occurrences	zero to many
Parent Elements	profiler
Child Elements	none

<modes> element**Purpose**

Tells the probe which mode(s) to run in.

Attributes

Attributes	Valid Values	Default	Description
enterprise	true false	Depends on mode chosen in installation. <ul style="list-style-type: none"> • true if pro is false • false if pro is true 	Probe in Enterprise mode (probe is working with Diagnostics Server)
pro	true false	Depends on mode chosen in installation. <ul style="list-style-type: none"> • true if enterprise is false • false if enterprise is true 	Probe in Profiler mode

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<diagnosticserver> element**Purpose**

Contains connection and settings information related to the Diagnostics Server which are used for enterprise mode.

Attributes

Attributes	Valid Values	Default	Description
url	Registrar URL. http:// <host>: <port>	none	URL to connect to registrar
delay	number	2	Number of seconds to wait before registering
keepalive	number	15	Number of seconds between keepalives
proxy	URL of proxy	none	Registrar connection proxy.
registeredhost name	string	none	Name of host to register as (external name for firewall traversing)
register_byip	true, false	false	Register using ipaddress instead of hostname

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<mediator> element**Purpose**

Specifies the diagnostics server that is in the Mediator mode to which events are to be sent when in the enterprise mode.

Attributes

Attributes	Valid Values	Default	Description
host	host name	none	Name of mediator
port	number	2612	Mediator port
block	true/false	false	Block until mediator connection established
ipaddress			local ipaddress to use when connecting to the eventserver.
localportstart	number	4000	Beginning of port range to use for tcp event channel connection to the Diagnostics Server in Mediator mode. Used only when ipaddress is specified.
localportend	number	5000	End of port range to use for tcp event channel connection to the Diagnostics Server in Mediator mode. Used only when ipaddress is specified.

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<webserver> element

Purpose

Specifies the local Web server properties for communication with the probe.

Attributes

Attributes	Valid Values	Default	Description
start	number	35000	Starting port for webserver
end	number	35100	Ending port for webserver
ipaddress			local ipaddress to run webserver on.

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<bufferpool> element**Purpose**

Configures the bufferpool behavior.

Attributes

Attributes	Valid Values	Default	Description
size	number	65536	Size of each buffer
buffers	number	512	Number of buffers in pool
sleep	number	1000	Number of milliseconds between flush checks
expires	number	1000	Number of milliseconds before buffer expires

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<lwmd> element

Purpose

Configures the light-weight memory diagnostics feature.

Attributes

Attributes	Valid Values	Default	Description
enabled	true false	false	Enables lwmd capturing
sample	string	1m	Sample interval (h-hour/m-minute/s-second)
autobaseline	string	1h	Auto baseline interval
growth	number	15	Number of collections to growth track
size	number	15	Number of collections to size track
include	string	none	Comma separated list of collection classes to include to the exclusion of all others.
exclude	string	none	Comma separated list of collection classes to exclude
manualbaseline	string	none	Manual baseline time

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	exclude, include

<instrumentation> element**Purpose**

Contains logging configuration for instrumenter.

Attributes

None.

Elements

Number of Occurrences	1 per parent
Parent Elements	probeconfig, process
Child Elements	logging

<logging> element**Purpose**

Sets the logging level for the probe.

Attributes

Attributes	Valid Values	Default	Description
level	off info severe warning debug events	off	Sets the logging level for the probe

Elements

Number of Occurrences	
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<logging> element**Purpose**

Sets the logging level for instrumenter.

Attributes

Attributes	Valid Values	Default	Description
level	assert break info severe warning debug points checksum ch sig cil classmap ilasm symbol deepmode load	off	Sets the logging level for instrumentation
threadids	true false	true	Include thread ids in log

Elements

Number of Occurrences	
Parent Elements	instrumentation
Child Elements	none

<points> element**Purpose**

List a points file for inclusion in instrumentation.

Attributes

Attributes	Valid Values	Default	Description
file	string	none	Name of instrumentation points file

Elements

Number of Occurrences	zero or more
Parent Elements	appdomain, process
Child Elements	none

<sample> element

Purpose

Sets the sampling type and rate.

Attributes

Attributes	Valid Values	Default	Description
method	percent, count, period	percent	Sets the sampling method <ul style="list-style-type: none"> • for percent rate must be 0-100 • for count rate must be >1 • for period rate must be one of standard Diagnostics time strings (3m for 3 minutes, 4s for 4 seconds, and so forth)
rate	number	0	Sets the sampling rate for percent type

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<symbols> element**Purpose**

Limits the number of unique URIs and SQL strings that can be captured to control the amount of memory consumed.

Attributes

Attributes	Valid Values	Default	Description
maxuri	number	1000	Sets the top limit for number of unique URIs that can be captured.
maxuriname	string	Maximum number of unique URIs exceeded	
maxsql	number	1000	Sets the top limit for number of unique URIs that can be captured.
maxsqlname	string	Maximum number of unique SQLs exceeded	

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	none

<trim> element

Purpose

Configures the trimming feature to reduce data volume.

Attributes

None.

Elements

Number of Occurrences	1 per parent
Parent Elements	appdomain, probeconfig, process
Child Elements	depth, latency

<depth> element

Purpose

Configures depth trimming.

Attributes

Attributes	Valid Values	Default	Description
enabled	true false	false	Enables depth trimming
depth	number	25	Sets the depth for depth trimming

Elements

Number of Occurrences	1
Parent Elements	trim
Child Elements	none

<latency> element**Purpose**

Configures latency trimming.

Attributes

Attributes	Valid Values	Default	Description
enabled	true false	false	Enables latency trimming
throttle	true false	false	Enables latency trimming throttling
min	number	2	Minimum latency threshold.
max	number	100	Maximum latency threshold
increment	number	2	Threshold increment
increment threshold	number	75	The percentage of the buffer usage before the throttling should be incremented
decrement threshold	number	50	The percentage of the buffer usage before the throttling should be decremented

Elements

Number of Occurrences	1
Parent Elements	trim
Child Elements	none

<cputime> element**Purpose**

Controls the **cputime** setting property.

Attributes

Attributes	Valid Values	Default	Description
mode	none, serverrequest, method	none	

Elements

Number of Occurrences	1
Parent Elements	<probeconfig>, <process>, or <appdomain>
Child Elements	none

<rum> element**Purpose**

Controls the settings (enabled/header) for Real User Monitoring.

Attributes

Attributes	Valid Values	Default	Description
enable	true false	false	Enables Real User Monitoring
header	""	"X-Mercury-Diag-Response-Color"	The value for header can be any string but the string must be configured the same for all components

Elements

Number of Occurrences	1
Parent Elements	<probeconfig>, <process>, or <appdomain>
Child Elements	none

<xvm> element

Purpose

Controls the cross vm settings.

Attributes

None.

Elements

Number of Occurrences	1
Parent Elements	<probeconfig>, <process>, or <appdomain>
Child Elements	<ws>

<ws> element

Purpose

Controls Web services correlation sampling.

Attributes

None.

Elements

Number of Occurrences	1
Parent Elements	<xvm> element
Child Elements	<sample>

<appdomain> element**Purpose**

Builds an appdomain inclusion list for processes that host multiple appdomains.

The value of none means that it applies to all.

Attributes

Attributes	Valid Values	Default	Description
name	string	none	Name that the .NET AppDomain settings apply to

Elements

Number of Occurrences	zero or more
Parent Elements	process
Child Elements	bufferpool, credentials, diagnosticserver, mediator, id, ipaddress, logging, lwmd, modes, points, profiler, sample, trim, webserver, symbols, filter

<process> element

Purpose

Provide an inclusion filter list of which applies to the probe.

None means none.

Attributes

Attributes	Valid Values	Default	Description
name	string	none	Name of the .NET process that these setting apply to
monitorheap	string	none	Enables the Heap tab and Heap Breakdown processing when the value is set to true.

Elements

Number of Occurrences	zero or more
Parent Elements	probeconfig
Child Elements	appdomain, bufferpool, credentials, diagnosticsserver, mediator, id, instrumentation, ipaddress, logging, lwmd, modes, points, profiler, sample, trim, webserver, filter, symbols

16

Advanced J2EE Probe and Application Server Configuration

This chapter provides instructions for configuring the J2EE Probe and the application server using the properties and settings provided for situations that require more advanced configurations.

This chapter describes:	On page:
Advanced Configuration Directory	306
Configuring the Probe to Work with Mercury Products	307
Configuring the Probes for Multiple Application Server JVM Instances	312
Disabling the J2EE Diagnostics Profiler	322
Specifying Probe Properties as Java System Properties	322
Controlling Probe Logging	323
Setting the Probe Host Machine Name	325
Controlling Automatic Method Trimming on the Probe	326
Controlling Probe Throttling	329
Configuring a Probe for a Proxy Server	332
Configuring Reverse HTTP for a Probe in MMS	332
Configuring the Application Server for Deep Diagnostics Allocation Capture	334
Diagnostics Probe Administration Page	334

This chapter describes:	On page:
J2EE Probe Security and User Management	337
Using BEA JRockit Mission Control	341

Advanced Configuration Directory

The following bullet points will help guide you to the probe configuration topic that will enable you to configure the probe to address particular issues in your environment.

- ▶ If you have a probe that was installed to work with one Mercury product and you now would like to work with other products, see “Configuring the Probe to Work with Mercury Products” on page 307.
- ▶ If your application server is using multiple JVMs or if you want to capture data for multiple JVMs you must perform additional probe configuration steps described in “Configuring the Probes for Multiple Application Server JVM Instances” on page 312.
- ▶ If you have a probe that you want to prevent others from using in Profiler mode, see “Disabling the J2EE Diagnostics Profiler” on page 322.
- ▶ If you have more than one JVM using a single probe installation, you may need to set some of the probe properties in the Java system properties in the application start up commands. For more information, see “Specifying Probe Properties as Java System Properties” on page 322.
- ▶ If you want to have log messages posted to the probe logs for lower level messages what can adjust the log level as described in “Controlling Probe Logging” on page 323.
- ▶ If you have more than one probe installed on the same host you can make sure that the log messages for each probe are stored in a different file as explained in “Changing the Log Directory for a Probe” on page 324.
- ▶ If you want to examine the performance of processing that would normally be trimmed from the metrics reported in Diagnostics you can reduce the level of trimming, or turn off trimming completely as described in “Controlling Automatic Method Trimming on the Probe” on page 326.

- In some situations, you may want to turn off probe throttling so that you can be sure that you are getting all of the available performance information without regard to the performance impact on the application itself. For information on throttling and how to control it see, “Controlling Probe Throttling” on page 329.
- If there is a proxy present between the probe and the Diagnostics Server in Commander mode you must set the correct property to tell the probe the URL of the Diagnostics Server in Commander mode. For more information, see “Configuring a Probe for a Proxy Server” on page 332.
- If you have installed a J2EE Probe that is to be used in a Mercury Managed Services (MMS) environment you are encouraged to disable the reverse http (rhttp) communication between the probe and the Diagnostics Server in Mediator mode. For more information, see “Configuring Reverse HTTP for a Probe in MMS” on page 332.
- If you are capturing Deep Diagnostics Allocation workload, you may need to modify your JVM settings to ensure the capture works properly. For more information, see “Configuring the Application Server for Deep Diagnostics Allocation Capture” on page 334.

Configuring the Probe to Work with Mercury Products

Adding a Mercury Product to the Probe Configuration

The J2EE Probe is a lifecycle probe that can be configured to monitor your applications from development through implementation and production. The J2EE Probe can be configured to work with several Mercury products or to work as a Diagnostics Profiler without other Diagnostics components or Mercury products.

The products with which the J2EE Probe can work is determined by the value of the **active.products** property located in the property file `<probe_install_dir>/etc/probe.properties`.

The value of the **active.products** property is set at the time that you install the J2EE Probe. For information on probe installation see Chapter 6, “Installing the Mercury Diagnostics Probe for J2EE.”. To enable the probe to capture data for different Mercury products you must set the value of the **active.products** property by editing the property file and restarting the application server.

Note: If you want to use the J2EE Probe that was downloaded with the Mercury Diagnostics Profiler for J2EE with other Mercury products, contact Mercury Customer Support. The instructions that follow will not work for a probe installed for the J2EE Diagnostics Profiler without intervention from Mercury Customer Support.

To see Diagnostics data in the user interface of the interfacing Mercury Products there are additional configuration steps that must be completed. See the chapters in this guide that correspond to the Mercury product that you are using to complete the configuration.

The sections that follow provide instructions for configuring each product mode.

PRO Product Mode – Mercury Diagnostics Profiler for J2EE

For Mercury Diagnostics Profiler for J2EE, the product mode is PRO. In this mode, the probe gathers additional metrics and presents those metrics on a user interface that it makes available through a URL on the probe host.

If you are running the J2EE Probe as part of the J2EE Diagnostics Profiler in Development Mode, there are restrictions placed upon the probe to limit the load that the probe can handle.

If you are running the J2EE Probe as part of Diagnostics Standalone, or along with another Mercury product, the Profiler is enabled without the load restrictions.

Enterprise Product Mode

When configured in Enterprise mode, the probe works with the Mercury products such as Mercury Business Availability Center, LoadRunner, Performance Center, and as Mercury Diagnostics Standalone. In Enterprise mode, the probe may also be configured to work with Mercury Deep Diagnostics.

To configure the probe for Enterprise mode:

- 1** Set the **active.products** property to Enterprise and restart the Application Server.
- 2** Configure the probe to register with the Diagnostics Server in Commander mode.

One of the functions of the Diagnostics Server in Commander mode is to keep track of the Diagnostics components so that it can facilitate communication between them and keep you informed about the status and health of the components.

To configure the probe to register with the Diagnostics Server, set the host name and port using the **registrar.url** property which can be found in the property file: `<probe_install_dir>\etc\dispatcher.properties`.

Below is an excerpt from **dispatcher.properties** showing the **registrar.url** property:

```
## the URL of the registrar
registrar.url=http://host01.company.com:2006/registrar/
```

- 3** Configure the probe to connect to the Diagnostics Server in Mediator mode.

The probe must be able to transmit the processing metrics that it gathers to the Diagnostics Server in Mediator mode in order for Mercury Business Availability Center, LoadRunner, and Performance Center to be able to receive, process, and display the reports. In LoadRunner and Performance Center, Diagnostics assigns a Diagnostics Server to the probe. In Mercury Business Availability Center, you must tell the probe which Diagnostics Server to work with.

To configure the probe to communicate with the Diagnostics Server in Mediator mode, set the host name and the port using the **mediator.host.name** and **mediator.port.number** properties which can be found in `<probe_install_dir>\etc\dynamic.properties`.

The default Diagnostics Server Mediator port is 2612. If, when installing the Diagnostics Server in Mediator mode, you set the Diagnostics Server Mediator port to a value other than the default, make sure to use that same port number when you set the **mediator.port.number** property for the probe.

Below is an excerpt from the **dynamic.properties** file showing these properties:

```
#the host the Topaz mediator is running on
mediator.host.name=host01.company.com

#the port the Topaz mediator is listening on (default is 2612)
mediator.port.number=2612
```

DD Mode–Deep Diagnostics

Mercury Deep Diagnostics (DD) for J2EE does not include a probe in its installation; instead, it uses the probe installed with Mercury's LoadRunner 8.1, Performance Center 8.1, or Mercury Business Availability Center 6.x. The probe installed with these other products may be configured to work with Mercury Deep Diagnostics for J2EE by adding **DeepDiagnostics** to the **active.products** property. A probe can run in standalone DD mode, and can also be configured so that it will be able to gather additional metrics for the Enterprise mode.

To configure the probe and Application Server for Deep Diagnostics mode:

- 1** Ensure that you have installed the Deep Diagnostics System Under Test agent on the application server machine. If the DD agent is not there, run the Deep Diagnostics setup, and install the DD SUT.

When you install DD, you are asked to provide the location of the probe installation directory, `<probe_install_dir>`. This information is used to update properties that DD uses to contact the probe and to associate the probe with the application.

- 2 Verify that the Deep Diagnostics Server can see the installed DD agent. (In the Deep Diagnostics Console, under the **host tree** node, see the **global** tab.)
- 3 Ensure that you have configured the probe and application server as described in the *Mercury Diagnostics Installation and Configuration Guide*.
- 4 Locate the directory where the probe was installed, `<probe_install_dir>`.
- 5 Set the following properties in `<probe_install_dir>/etc/capture.properties` as follows:
 - ▶ Set the **timeserver.host** property to the host name for the machine where the Deep Diagnostics server was installed.
 - ▶ Set **timer.synchronize.with.server.clock** = false.
 - ▶ Set **use.native.timestamps** = true.
- 6 Set the following properties in `<probe_install_dir>/etc/probe.properties`:
 - ▶ Add DeepDiagnostics to the value of **active.products**. See the instructions provided in the property files for the correct syntax.
To set the product mode so that the probe is configured for both LoadRunner / Performance Center and Deep Diagnostics, the property should be:

`active.products=Enterprise, DeepDiagnostics`
 - ▶ Set the **id** property in **probe.properties** to identify the probe to Deep Diagnostics. This value must exactly match the Application Definition used to specify instrumentation in the Deep Diagnostics.
 - ▶ Restart the application server.

Note: If the probe is used in Deep Diagnostics mode simultaneously with AD or AM, and a different “id” for the probe is necessary for those modes, you can specify the **deepdiagnostics.id** property in **probe.properties** to indicate which Deep Diagnostics application definition this probe is associated with.

Removing a Product from the Probe Configuration

The product mode for the probe is configured using the **active.products** property located in the file `<probe_install_dir>\etc\probe.properties`.

To configure the probe so that it will no longer support a product, update the **active.products** property so that the product is no longer included in the property value.

For example if a probe has been configured for Deep Diagnostics and Mercury Business Availability Center, the **active.products** property would be set as follows:

```
active.products=Enterrise, DeepDiagnostics
```

To reconfigure the probe so that Deep Diagnostics is no longer included, for example, remove the product from the property value:

```
active.products=Enterprise
```

To remove all products except the Mercury Diagnostics Profiler, set the property as follows:

```
active.products = PRO
```

Configuring the Probes for Multiple Application Server JVM Instances

When your application server is using multiple JVMs, or when you want to capture data for multiple JVMs, you must perform additional probe configuration steps. There are two options; you may install one probe for each JVM on a host, or you may install one probe that will be shared by all of the JVMs.

This section describes:

- ▶ Configuring Multiple JVMs to Use a Single Probe Installation
- ▶ Configuring Separate Probe Installations For Each JVM
- ▶ Configuring Multiple JVMs to Use a Single Probe for Deep Diagnostics

Configuring Multiple JVMs to Use a Single Probe Installation

To allow multiple JVMs to share a single probe installation, you must configure a separate instance of the probe for each JVM. This configuration enables the following:

- ▶ Establishment of communication between the JVM and the probe
- ▶ Identification of the probe by the JVM
- ▶ Instrumentation of the JVM to instruct the probe about the performance metrics it will monitor

To configure a J2EE Probe to work with multiple JVMs:

- 1 When a probe is being used to monitor multiple JVM versions, the JRE Instrumenter must be run once for each JVM version to enable the probe to monitor the events of the applications that are running on the JVMs.

If you have not already run the JRE Instrumenter for each of the JVM versions that the probe will be working with, do so now. See “Running the JRE Instrumenter” on page 165.

Note: You must run the JRE Instrumenter even if you let the installer instrument the JRE when the probe was installed. The JRE Instrumenter prepares the applications for multi-JVM support.

Note: If you are using multiple JVM versions and have run the JRE Instrumenter multiple times, be sure to include two paths in the **-Xbootclasspath** parameter exactly as indicated in the output from the Instrumenter. The first path is for the specific JVM and the second path for the default “boot” directory. See the following example:

```
-Xbootclasspath/p:/path/to/javaprobe/classes/IBM/1.3.1:/path/to/javaprobe/
classes/boot
```

- 2 Specify the range of ports from which the probe can automatically select. The J2EE Probe communicates using the mini Web server. A separate port is assigned for probe communications for each JVM that a probe is monitoring. By default, the port number range is set to 35000 - 35100. You must increase the port number range when the probe is working with more than 100 JVMs.

Note: If a firewall separates the probe from the other Diagnostics components, you must configure the firewall to allow communications using the ports in the range that you specify. For more information, see Appendix , “Configuring Diagnostics Components to Work with a Firewall.”

If you have configured the firewall to allow probe communications on a range of ports that is different than the default, make sure to update the port range values discussed in the following bullets accordingly.

- a Locate the **webservice.properties** file in the folder `<probe_install_dir>/javaprobe/etc`.
- b Set the following properties to adjust the range of ports that are available for probe communications.
 - The minimum port in the port number range is set using the following property:
jetty.port=35000
 - The maximum port in the port number range is set using the following property:
jetty.max.port=35100
- 3 Assign a custom probe Identifier to the probe for each JVM, using the Java command line.

`-Dprobe.id=<Unique_Probe_Name>`

The probe ids defined on the Java command line override the probe names that are defined in the **probe.properties** file using the probe’s **id** property.

The following is an example of a WebLogic startup script before adding the **probe.id** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -classpath "%CLASSPATH%"
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer -Dbea.home="C:\bea"
-Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server
```

The following is an example of a WebLogic startup script after adding the **probe.id** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m"
-Xbootclasspath/p:C:\Program Files\Mercury Interactive\Diagnostics for
J2EE\Probe\classes\Sun\1.4.1_03;C:\Program Files\
Mercury Interactive\Diagnostics for J2EE\Probe\classes\boot"
-classpath "%CLASSPATH%"
-Dprobe.id=<Unique_Probe_Name> -Dweblogic.Domain=petstore
-Dweblogic.Name=petstoreServer
-Dbea.home="C:\bea" -Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server
```

Note: The command-line properties must be entered on one line, without any breaks.

- 4 Specify the points file that the probe will use. By default LoadRunner/Performance Center will set the `points.file.name` to `auto_detect.points`. You may want to specify that a custom points file be used when you need to use more than one custom instrumentation plan, or where you have several JVMs on the same machine using a single probe instance, and one or more of the JVMs needs specific methods and classes included in a layer support custom instrumentation.

```
-Dprobe.points.file.name="<Custom_AutoDetect_Points_File>"
```

Configuring Separate Probe Installations For Each JVM

When there are multiple JVMs on a single host, you may install a separate probe for each JVM instance. To use a separate probe for each JVM, you must install the probe multiple times, and define an instance of each probe by setting the probe's **id** property in the **probe.properties** file in each probe's installation directory.

To define an instance of each installed probe:

- 1 If you have not already manually run the JRE Instrumenter for the JVMs that the probe will be working with, do so now. To run the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.
- 2 Locate the **probe.properties** file in the `<probe_install_dir>/javaprobe/etc` directory.

For example:

```
C:\Program Files\Mercury Interactive\common\JavaProbe\etc\probe.properties
```

- 3 Assign a name to the **id** property that will be unique on the server and on the Diagnostics Server as follows:

```
id=<uniqueProbeName>
```

When the probe instance is started, a log file is created in the `<probe_install_dir>/javaprobe/log` directory where the log messages for the probe are stored.

Note: This configuration enables you to perform the diagnostics; but you will have to spend some time configuring each probe installation. Using a single probe for multiple JVMs can provide you with the same diagnostics detail, while allowing you to spend less time configuring the components.

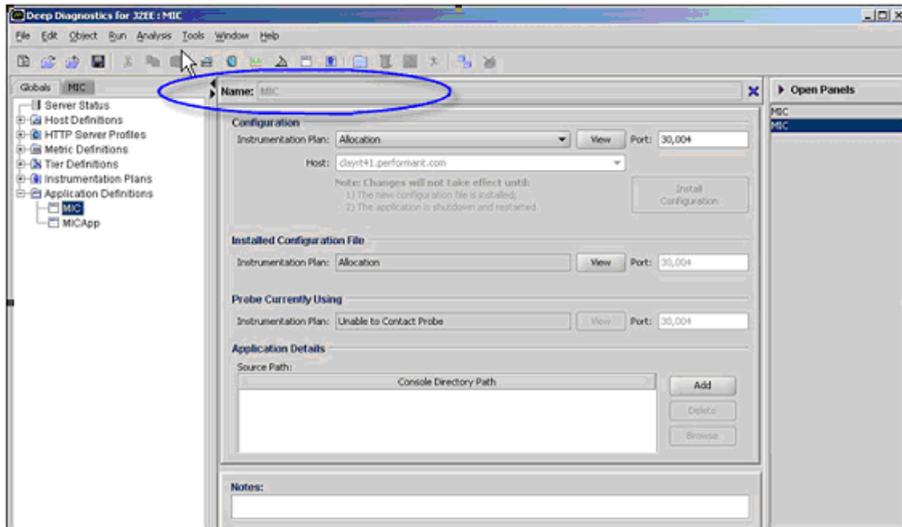
Configuring Multiple JVMs to Use a Single Probe for Deep Diagnostics

When a probe is used by multiple JVMs, each instance of the application server must identify with an application definition in Deep Diagnostics. This is accomplished by defining the application definition using the Deep Diagnostics Console. The application definition associates the port number that the probe will use to communicate with Deep Diagnostics and the host name of the computer on which the probe is installed.

Note: This setup is the same if the application server instances are sharing a common JRE (for example, when multiple WebLogic JVM instances are using a shared file system).

To configure Deep Diagnostics and the JVM:

- 1 Create a new application definition, or access an existing application definition from the Deep Diagnostics Console.
 - a Access the Application Definition window by selecting an application definition from the tree view on the Globals tab.
 - b Type the name for the Application Definition in the **Name** field.



- 2 Explicitly specify the port that the probe is to use for communicating with the Deep Diagnostics server.

In the Configuration section of the Application Definition window, in the **Port** field, type the number of the port that the probe will use to communicate with Deep Diagnostics. Deep Diagnostics updates the **port** property in the points file.

3 Instrument each JVM that is being monitored by the probe.

- a** For each JVM version that is being monitored by the probe:
- Run the JRE Instrumenter. For details on running the JRE Instrumenter, see “Running the JRE Instrumenter” on page 165.
 - Copy the `<probe_install_dir> classes/boot` directory and rename the directory to a name that indicates the JVM version that you ran the JRE Instrumenter for.
 - For example, if you just ran the JRE Instrumenter for the WebLogic JVM you may want to rename the boot directory to:

```
<probe_install_dir>\classes\bootwls
```

- b** Add the `Xbootclasspath` parameter to the Java command line that starts the application server. Make sure that the `Xbootclasspath` parameter points to the `<probe_install_dir> classes/boot` version that you created for the application server type that is appropriate for the startup script that you are modifying.

The parameter must be placed at the beginning of the Java parameters following any JIT options such as `-hotspot` or `-classic`.

Note: When modifying the `-xbootclass` parameter. Be sure to use quotes if there are spaces in the path that you specify.

The following is an example of the `Xbootclasspath` parameter where `<probe_install_dir>` is to be replaced with the path to the directory where the probe was installed. The boot directory is the copy of the directory that was created for the JVM.

```
-Xbootclasspath/p:"C:\Program Files\Mercury Interactive\Diagnostics for J2EE\Probe\classes\Sun\1.4.1_03;C:\Program Files\Mercury Interactive\Diagnostics for J2EE\Probe\classes\boot"
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

The following is an example of a WebLogic startup script before adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m -classpath "%CLASSPATH%"  
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer  
-Dbea.home="C:\bea"  
-Dweblogic.management.password=%WLS_PW%  
-Dweblogic.ProductionModeEnabled=%STARTMODE%  
-Dcloudscape.system.home=./samples/eval/cloudscape/data  
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server
```

The following is an example of a WebLogic startup script after adding the **Xbootclasspath** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m  
-Xbootclasspath/p:"C:\Program Files\Mercury Interactive\Diagnostics for  
J2EE\Probe\classes\Sun\1.4.1_03;C:\Program Files\  
Mercury Interactive\Diagnostics for J2EE\Probe\classes\boot"  
-classpath "%CLASSPATH%"  
-Dweblogic.Domain=petstore -Dweblogic.Name=petstoreServer  
-Dbea.home="C:\bea"  
-Dweblogic.management.password=%WLS_PW%  
-Dweblogic.ProductionModeEnabled=%STARTMODE%  
-Dcloudscape.system.home=./samples/eval/cloudscape/data  
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server"
```

- 4** Point each of the application server instances to the probe directory, and assign a unique probe Identifier that will be used by the application server to reference the probe. To define the unique probe Identifier for each application server, add the following system property to the startup script for each instance, and to the **JavaProbe/etc/probe.properties** file:

```
-Dprobe.id=<Unique_DD_Probe_Name>
```

Following is an example of a WebLogic startup script before adding the **probe.id** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m
-Xbootclasspath/p:"C:\Program Files\
Mercury Interactive\Diagnostics for J2EE\Probe\classes\Sun\1.4.1_03;C:\Program
Files\Mercury Interactive\Diagnostics for J2EE\Probe\classes\boot"
-classpath "%CLASSPATH%" -Dweblogic.Domain=petstore -Dweblogic.Name=pet-
storeServer
-Dbea.home="C:\bea" -Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1/lib/weblogic.policy" weblogic.Server
```

Note: The startup script examples are shown with line breaks. The actual scripts do not have line breaks and the text of the commands will wrap on your screen as necessary.

Following is an example of a WebLogic startup script after adding the **probe.id** parameter:

```
"%JAVA_HOME%\bin\java" -hotspot -ms64m -mx64m "-Xbootclasspath/p:C:\Program Files\Mercury Interactive\Diagnostics for J2EE\Probe\classes\Sun\1.4.1_03;C:\Program Files\Mercury Interactive\Diagnostics for J2EE\Probe\classes\boot"
-classpath "%CLASSPATH%"
-Dprobe.id=<Unique_DD_Probe_Name> -Dweblogic.Domain=petstore
-Dweblogic.Name=petstoreServer -Dbea.home="C:\bea"
-Dweblogic.management.password=%WLS_PW%
-Dweblogic.ProductionModeEnabled=%STARTMODE%
-Dcloudscape.system.home=./samples/eval/cloudscape/data
-Djava.security.policy=="C:\bea\wlserver6.1\lib\weblogic.policy" weblogic.Server
```

The system creates a separate log file for each probe name, so the log messages generated by the probe for a particular application server appear in its own log file.

Disabling the J2EE Diagnostics Profiler

You can disable the Mercury Diagnostics Profiler for J2EE on a J2EE Probe so that it cannot be accessed accidentally. When the J2EE Diagnostics Profiler is disabled, the user interface cannot be accessed from the J2EE Diagnostics Profiler URL: http://<probe_host>:<probeport>/profiler.

To disable the J2EE Diagnostics Profiler, set the **disable.profiler** property in `<probe_install_dir>/etc/probe.properties` to true.

The default value for **disable.profiler** is false. To enable the J2EE Diagnostics Profiler once it has been disabled, change the value of the **disable.profiler** property from true to false.

Specifying Probe Properties as Java System Properties

All of the J2EE Probe properties, except for those defined in the **dynamic.properties** property file, can be specified as Java System properties on the startup command-line for the application server. This is very useful when there is more than one JVM using a single probe installation.

To specify a probe property as a Java System property, prepend the letter **D** and the first part of the properties file name to the property name. This is best explained using some examples:

- ▶ To set the **id** property in **probe.properties** from the startup command, you concatenate the **D** and **probe** from the property file name, and then tack on the name of the property that you are specifying, that is, **id**, as follows:

```
-Dprobe.id=SomeId
```

- ▶ To set the **active.products** property in **probe.properties** from the startup command, you concatenate the **D** and **probe** from the property file name, and then tack on the name of the property that you are specifying, that is, **active.products**, as follows:

```
-Dprobe.active.products=AD,AM,DeepDiagnostics
```

- ▶ To set the **registrar.url** property in **dispatcher.properties** from the startup command, you concatenate the **D** and **dispatcher** from the property file name, and then tack on the name of the property that you are specifying, that is, **registrar.url**, as follows:

```
-Ddispatcher.registrar.url=http://host01.company.com:2006/registrar
```

Controlling Probe Logging

You can control the level of the messages that the probe logs and change the location where the log messages are posted using the probe properties.

Controlling the Log Message Level

The level of messages from the probe that are logged to the standard output is controlled by the **lowest_printing_level** property in the property file `<probe_install_dir>/etc/logging.properties`. The default setting for this property is **OFF** which prevents almost all probe messages from being logged to the console.

You can adjust the logging level dynamically by changing the value assigned to the **lowest_printing_level** property. The level of messages that are logged changes as soon as you save the property file.

The valid values for the **lowest_printing_level** property are:

Property Value	Description
OFF	No messages are logged.
DEBUG	All messages are logged.
INFO	Info, Severe, and Warning messages are logged.
WARN	Warning and Severe messages are logged.
SEVERE	Severe messages are logged.

Changing the Log Directory for a Probe

The default location for the log directory for a probe is `<probe_install_dir>/log`. When you have more than one probe on the same host, you can change the location of the log directory for each probe using the **log.dir** property. This property can be set in two ways:

- The value of the **log.dir** property can be set in the property file `<probe_install_dir>/etc/probe.property`.
- The value of the **log.dir** property can be specified on the startup command-line for the application server as a JAVA system property as shown in the following example:

```
-Dprobe.log.dir=/path/to/log
```

For more information on specifying the `log.dir` property on the startup command line, see “Configuring a Probe for a Proxy Server” on page 332.

Setting the Probe Host Machine Name

The Probe's host name is used to register the Probe with the Diagnostics Server in Commander mode. The Diagnostics Server in Commander mode uses the Probe's host name to communicate with the Probe and displays it along with the system metrics for the server that the Probe is monitoring in the Diagnostics views.

The Probe normally can detect the host name of the machine that is its host. In some situations, the server configuration is faulty and the Probe cannot detect the correct host name. In situations where a firewall or NAT is in place or where your Probe host machine has been configured as a multi-homed device, it may not be possible for the Probe to properly detect its host.

If the Probe cannot detect its host name, you can instruct the Probe to get the host name via a reverse DNS lookup based on the socket connection or you can specify the host name using a Probe property.

Instructing the Probe to use Reverse DNS Lookup

If the configuration of the Probe's host prevents the Probe from detecting the host name, you can instruct the Probe to detect the host name using a reverse-DNS lookup by setting the `server.host.name.lookup` property. This property is located in the `<probe_install_dir>/etc/dispatcher.properties` file.

The default value for the `server.host.name.lookup` property is 'false'. This tells the Probe to do the lookup without using reverse-DNS. Set this property to 'true' to instruct the Probe to use reverse-DNS lookup.

Manually Specifying the Probe Host Name

The `registered_hostname` property allows you to manually set host machine name for the Probe and stop the Probe from doing the automatic lookup.

To set default host machine name for a Probe, set the `registered_hostname` property, located in the J2EE Probe property file, `<probe_install_dir>/etc/dispatcher.properties`, to a machine name or IP address.

When you set the **registered_hostname** property the automatic lookup of the host name is disabled.

Note: Setting the **registered_hostname** property because of a NAT or firewall is only an issue for a test environment where you are using LoadRunner, Performance Center, or Diagnostics Standalone.

However, if you should set the **registered_hostname** in a production environment where you are using Mercury Business Availability Center or Diagnostics Standalone, the name that you specify will be shown as the host name in System Health.

Controlling Automatic Method Trimming on the Probe

The default configuration for the probe includes settings that control the trimming of methods. Trimming can be controlled based upon how long the method takes to execute, which is known as *latency*, and by the *stack depth* of the method call. The default configuration instructs the probe to trim both by latency and by depth.

You may want to reduce the level of trimming, or turn off trimming completely, for certain Diagnostics situations. You can control the trimming using the **minimum.method.latency** and **maximum.stack.depth** properties in `<probe_install_dir>/etc/capture.properties`.

Controlling Latency Trimming

Methods that complete with latency greater than or equal to the value of the **minimum.method.latency** property are captured, and those that complete with latency less than this limit are trimmed to avoid incurring the overhead for methods that are less likely to be of interest.

Methods that are the root for a call tree are not trimmed. Methods that threw an exception are not trimmed, unless the probe is facing a severe throttling situation due to an inability to send events to the Diagnostics Server in Mediator mode.

Because of threading and buffering behavior, partial information about a method that was trimmed may be saved in the Profiler buffers (Development mode) or transmitted to the Diagnostics Server (Production and Test mode). When the Diagnostics Server detects that it only received partial information for a method, it issues a warning message. You should ignore this message unless your run requires that the information for all methods be captured.

If the information for all methods must be captured, lower the value of the **minimum.method.latency** property or set it to zero.

Consider the following when setting the **minimum.method.latency** property:

- ▶ The lower the value of the **minimum.method.latency** property, the greater the chance that the performance of your application will be adversely impacted.
- ▶ Depending on your platform, and whether native timestamps are being used (**use.native.timestamps** = false), it may not be useful to specify this value in increments of less than 10 ms.
- ▶ For Mercury Business Availability Center, LoadRunner, and Performance Center capture, the throttle mechanism will automatically increase the effective minimum method latency value if the Diagnostics Server is unable to keep up with the number of events being sent.
- ▶ For Deep Diagnostics, the throttle mechanism will automatically increase the effective minimum method latency value if the probe is unable to write events out to the workload log file fast enough. For more information on throttling, see “Controlling Probe Throttling” on page 329.

Controlling Depth Trimming

Methods that are called at a stack depth less than or equal to the value of the **maximum.stack.depth** property will be captured. Those that are called at a stack depth greater than this limit will be trimmed to avoid incurring the overhead for methods that are less likely to be of interest.

For example:

- If **maximum.stack.depth** is 3
- /login.do calls a() calls b() calls c()
- Only /login.do, a, and b will be captured.

The following should be considered when setting the **maximum.stack.depth** property.

- Setting a low **maximum.stack.depth** can significantly reduce the overhead of capture.
- In AM mode, it is not useful to have a **maximum.stack.depth** configured higher than the Diagnostics Server's depth trim. The Diagnostics Server's depth trimming is set using **trimming.type=depth** in **<Diagnostics Server_install_dir>/etc/server.properties**. This property is disabled by default.

Controlling Probe Throttling

About Throttling

The default configuration for the J2EE Probe is set to minimize the performance impact of collecting Diagnostics information. In addition, the probe is able to scale back automatically on the amount of information that it captures when it detects that its processing has started to have a negative impact on the performance of your application. The process for scaling back the amount of information that the probe captures is called *throttling*.

The probe throttles the amount of Diagnostics information that it collects by skipping over method calls that occur relatively quickly. The probe determines exactly what “relatively quickly” means based upon the amount of data that it is collecting. The skipped method calls can have a duration that starts at the configured minimum latency trim value, which defaults to 51 milliseconds. As load increases, the minimum latency trim value is increased. When the load decreases, the trim value is reduced to the configured minimum value.

The probe properties that are used to control throttling are listed below. These properties are found in `<probe_install_dir>/etc/capture.properties`.

► **gentle.reserve.buffer.count**

The gentle reserve buffers are temporarily allocated for workload spikes, while the load throttle measures are deploying. By default, the **gentle.reserver.buffer.count** property is set to the same value as the **maximum.private.buffer.count** property.

► **hard.reserve.buffer.count**

The hard reserve buffers are allocated for workload spikes when it is determined that the gentle reserve buffering cannot keep up. By default, the **hard.reserve.buffer** count is set to the same value as the **maximum.private.buffer.count** property.

► **buffer.wait.time**

The **buffer.wait.time** property is used to control the length of time that the probe will wait for an event to be buffered. This property tells the probe what to do when all throttle attempts are exhausted and an event cannot be buffered:

- -1 = indicates that the processing should wait for as long as it takes
- 0 = indicates that the event should be dropped immediately
- # = indicates that the processing should wait for # milliseconds before dropping the event

Turning Off Throttling

When you are using the probe to monitor a production system, the preferred behavior is to allow the probe to throttle automatically. However, when diagnosing some performance issues, you may want to sacrifice the performance of your application so that you can be sure that every diagnostic event is captured.

To configure the probe so that it will not throttle, you must edit the following properties in the `<probe_install_dir>/etc/capture.properties` file:

- Set **gentle.reserve.buffer.count** to 0
- Set **hard.reserve.buffer.count** to 0
- Set **buffer.wait.time** to -1

Note: When you are turning off throttling, be sure to set all three of the properties as instructed.

Tuning Throttling

By reviewing the probe logs, you can determine when throttling is being started and stopped. If throttling is being engaged more than you would like you should check to make sure that the probe thread configuration is synchronized with the number of threads that the application server has been configured to allow. If the application server has more threads than the probe can handle, throttling is engaged sooner and may not be disengaged.

You can check the maximum number of threads that can generate events into the probe by checking the log messages after the probe is started. A message similar to the following message should be logged:

```
2006-02-01 12:37:30,203 INFO class com.mercury.opal.capture.util.BufferPool [main] BufferPool ready: buffer size=6600000 (max 66 threads), # events in throttle overflow=4000000
```

In this example the maximum number of threads is 66.

The maximum number of threads that can generate events can be influenced using the **maximum.private.buffer.count** property, in property file `<probe_install_dir>/etc/capture.properties`.

If the maximum number of threads indicated by the log message is less than the number of threads that your application server can allocate, you can increase the value of the property. Be careful when increasing this value because this causes the memory overhead of the probe to increase. If the memory overhead is of concern, the value of the **maximum.buffer.size** and **minimum.buffer.size** properties may need to be decreased by an amount proportionate to the increase in the maximum buffer count.

Configuring a Probe for a Proxy Server

Important! This section only applies if you are using the probe with a Diagnostics Server.

There are two properties that are used to tell the probe the URL of the Diagnostics Server in Commanding mode. The property that you set depends upon whether there is a proxy present or not.

► **registrar.url in dispatcher.properties**

The **registrar.url** property in `<probe_install_dir>\etc\dispatcher.properties` is set when you install the probe. When there is a direct connection between the probe and the URL of the Diagnostics Server in Commander mode, the value of this property is used by the probe.

► **registrar.url in webserver.properties**

In the presence of a proxy, you must set the **registrar.url** property in the `<probe_install_dir>\etc\webserver.properties` file to indicate the URL of the Diagnostics Server in Commander mode.

Configuring Reverse HTTP for a Probe in MMS

When you install a J2EE Probe that is to be used in a Mercury Managed Services (MMS) environment you are encouraged to disable the reverse http (rhttp) communication between the probe and the Diagnostics Server in Mediator mode. This configuration prevents other customers in the MMS environment from inadvertently getting the metrics from your applications.

The following instructions have been provided so that you can disable or enable rhttp after the probe has been installed.

Disabling Reverse HTTP

If you neglected to disable rhttp when the probe was installed you can disable it manually using the following instructions.

To disable rhttp, comment out the **dispatcher.objects** property value in `<probe_install_dir>/etc/modules.properties` as shown in the following example:

```
#####
## Dispatcher Module properties
#####
dispatcher.objects=\
  com.mercury.diagnostics.common.net.InternalCommSecurityManager \
  com.mercury.diagnostics.common.modules.HostNameResolver, \
# com.mercury.opal.capture.dispatcher.RhttpClient
```

Enabling Reverse HTTP

If you inadvertently disabled rhttp for a probe, you can enable rhttp manually using the following instructions.

To enable rhttp, uncomment the **dispatcher.objects** property value in `<probe_install_dir>/etc/modules.properties` as shown in the following example:

```
#####
## Dispatcher Module properties
#####
dispatcher.objects=\
  com.mercury.diagnostics.common.net.InternalCommSecurityManager \
  com.mercury.diagnostics.common.modules.HostNameResolver, \
com.mercury.opal.capture.dispatcher.RhttpClient
```

Configuring the Application Server for Deep Diagnostics Allocation Capture

If you are going to use allocation capture, you must scope the capture to reduce the number of events. You may also have to increase the size of the permanent generation (perm gen) in the heap so that the capture will work properly. Capturing everything from allocation tier can cause the VM to crash if the perm gen is not sized properly.

The following example shows how to set the size of the permanent generation to 64m:

```
-XX:MaxPermSize=64m
```

Diagnostics Probe Administration Page

You can use the Diagnostics Probe Administration page to configure J2EE Probe settings, and set J2EE Probe and J2EE Diagnostics Profiler user privileges. You access the Diagnostics Probe administration page directly from your browser.

Accessing the Diagnostics Probe Administration Page

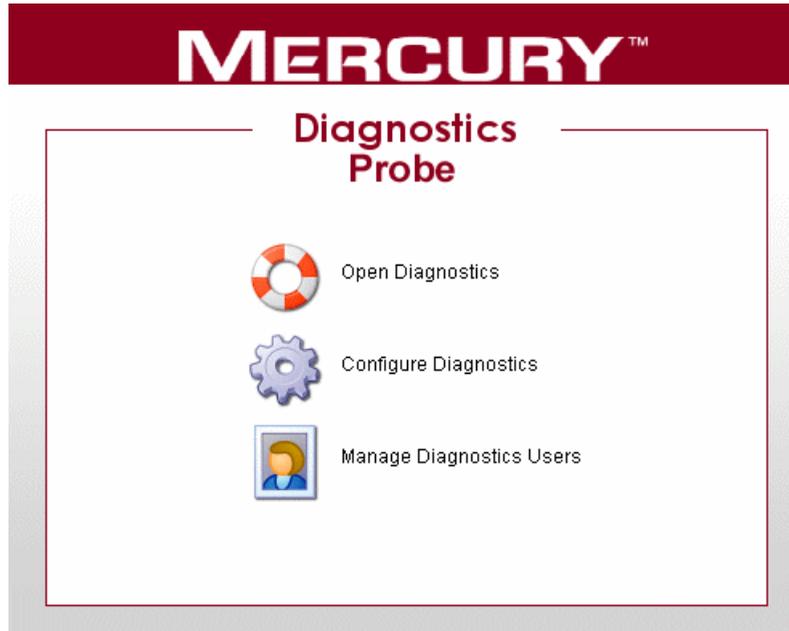
You open the Diagnostics Probe administration page inside your browser.

To access the Diagnostics Probe administration page:

- 1 In your browser, navigate to http://<probe_host>:<probeport>.

A probe is assigned to the first available port, beginning at 35000.

The Administration page opens.



2 Select the menu option for the activity that you want to perform.

- ▶ **Open Diagnostics.** Opens the J2EE Diagnostics Profiler.
- ▶ **Configure Diagnostics.** Opens the Components pages, which has a link to the J2EE Diagnostics Profiler, and to the User Administration page.
- ▶ **Manage Diagnostics Users.** Opens the User Administration page where you can add and maintain the security information, and user privileges. For more information, see “Understanding User Privileges” on page 337.

Diagnostics Probe Components Page

From the Components page you can open the J2EE Diagnostics Profiler, and access the User Administration page.

To access the Components page:

- 1 Open the Diagnostics Probe Administration page as described in “Accessing the Diagnostics Probe Administration Page” on page 334.
- 2 Click **Configure Diagnostics**.
- 3 If prompted, enter your user name and password.

This must be a valid user name, and must have both **View** and **Change** privileges. For information about valid user names and passwords, see “Understanding User Privileges” on page 337.

The Components page opens.



MERCURY [™]	
Components	
Component Name	Component Description
query	Query API - allows you to download diagnostics data in HTML, XML or as Java objects
profiler	Mercury Diagnostics Profiler - low-overhead profiling for J2EE.
security	User Management

- 4 Click one of the listed options, as follows:
 - **query**. For internal use by developers.
 - **profiler**. Opens the J2EE Diagnostics Profiler application.
 - **security**. Opens the User Administration page where you manage user privileges. For more information, see “J2EE Probe Security and User Management,” below.

J2EE Probe Security and User Management

Product security is controlled by creating user names and granting the user privileges so that they can perform various functions within the application.

Before you can configure the probe settings or access the J2EE Diagnostics Profiler, you must log on with a valid user name and password. The user name that is used to log on must have permission to view the requested data.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

Understanding User Privileges

There are two privilege levels defined for the Mercury Diagnostics Probe for J2EE:

- ▶ **View.** The user can view Diagnostics data from the UI. The user can also modify his own password. Access to the **Configure Diagnostics** option is not allowed.
- ▶ **Change.** The user can access the **Configure Diagnostics** option to update J2EE Probe configuration properties, and manage user information

Each privilege level stands alone. There is no inheritance of privileges from one level to the next. You must grant a user both of the privileges that are necessary to perform the functions that he needs to perform.

There are two default user names defined for the J2EE Probe, as described in the following table:

Default User Names	Privileges	Description
user	View	Can only view the data from the UI. Can manage own user password.
admin	View, Change	Can view data, configure the J2EE Probe, and manage user information. This is the default user.

You can use these default user names to access the J2EE Probe and J2EE Diagnostics Profiler functionality.

The passwords for the default user names are the same as the user names. For example, for the user name **admin**, the password is **admin**.

You can modify the password or privileges for the default user names to suit your needs. You can define new user names to suit your needs.

Accessing the User Administration Page

You manage user privileges on the Diagnostics Probe User Administration page. You can add new user names, delete user names, and update the security information for existing user names.

To access the User Administration page:

- 1 Open the Diagnostics Probe Administration page as described in “Accessing the Diagnostics Probe Administration Page” on page 334.
- 2 Click **Manage Diagnostics Users**.
- 3 If prompted, enter your user name and password.

This must be a valid user name, and must have at least **View** privileges, and **Change** privileges to manage more complex security settings. For information about valid user names and passwords, see “Understanding User Privileges” on page 337.

The User Administration page opens.

Managing User Details

You add new user names, update the security information for an existing user name, and delete users from the User Administration page.

Users with **View** privileges only, can maintain their own password. Users with both **View** and **Change** privileges can add new users and maintain the security information for any user.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

Adding and Deleting Users

Users with both **View** and **Change** privileges can add new users, grant user privileges, and delete users.

To add a user:

- 1 On the User Administration page, click **Create User**.
- 2 In the **New User Name** box, type a user name for the new user and click **OK**. The new user appears in the list of user names.
- 3 Under **Privileges**, select the appropriate privileges for the new user, as described in “Understanding User Privileges” on page 337.
- 4 Under **Change Password**, in the **Password** box, type a password for the new user, and confirm it by retyping it in the **Confirm Password** box.
- 5 In the **Password for <current user>** box, type the password of the user currently logged on.
- 6 Click **Save Changes** to save changes.

To delete a user:

- 1 On the User Administration page, in the **Password for <current user>** box, type the password for the user currently logged on.
- 2 Click the **Delete user** button.
- 3 A message box opens asking you if you want to delete the selected user. Click **OK** to delete.



Updating Privileges of Existing Users

Users with both **View** and **Change** privileges can update the password and privileges of existing users, including the default users. A user that has only **View** privileges can modify his own password, but not his privileges.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

To update an existing user's details:

- 1 On the User Administration page, under **Privileges**, select the relevant privileges for a particular user.
- 2 To modify the user's password, type the new password in the appropriate **Password** and **Confirm Password** boxes.
- 3 In the **Password for <current user>** box, type the password of the user currently logged on.
- 4 Click **Save Changes** to save the changes.

Using BEA JRockit Mission Control

If the application that you are monitoring is running on the BEA JRockit JVM and if the J2EE Probe has been configured to work with a Diagnostics Server, then you have access to a BEA JRockit Mission Control license that allows you to use Mission Control to monitor your applications performance.

When the probe that is monitoring an application running on the BEA JRockit JVM starts up and connects with the Diagnostics Server, it will ask the Diagnostics Server for the BEA license file if it does not already have one. The probe stores the BEA license file to enable you to license and begin using BEA JRockit Mission Control.

17

Probe Instrumentation

This chapter explains the process of assigning classes and methods to layers, using an advanced process called probe instrumentation.

This chapter describes:	On page:
About Instrumentation and Capture Points Files	343
Understanding Points in the Capture Points File	345
Understanding Diagnostics Layers	352
Defining Points Using Code Snippets	356

About Instrumentation and Capture Points Files

The instrumentation for a probe is specified in the Capture Points file. Instrumentation refers to the process of assigning classes and methods to *layers* using the *points* in the Capture points file. The probe gathers performance metrics for classes and methods and assigns the metrics to layers based upon the instructions that were provided in the points.

When you install the probe, a predefined Capture Points file is provided with the default points. You can find the Capture Points file at:
`<probe_install_dir>\etc\auto_detect.points`

You can modify the Capture Points file to change existing points or add new custom points. For example, to capture performance metrics for a class that contains J2EE/.NET functionality but does not directly implement a J2EE/.NET component, you add a new point to the Capture Points file.

Another example would be where you want to monitor a class that is of special interest to you in a separate layer. In this case you would modify the existing point to exclude the class of interest and then add a new point to instrument the target class for a new custom layer.

Below is an example of a Capture Points file that shows the points for JSP-`_jspService` and Servlet-all.

```
[JSP-_jspService]
;----- implements HttpJspPage Interface -----
;-- If an application does not have an implementation of this interface,
;-- an identical point MUST be added for whatever class
;-- translates the service() method call into the _jspService() method call.
class      = javax.servlet.jsp.HttpJspPage
; (See HttpCorrelation point for ignore documentation)
ignore_cl  = !jsp_servlet\.*
method     = !(service)|(_jspService)
signature  = !.*
deep_mode  = hard
layer     = Web Tier/JSP

[Servlet-all]
;----- extends HttpServlet -----
class      = javax.servlet.http.HttpServlet
; (See HttpCorrelation point for ignore documentation)
; In addition, ignore class we know we are not interested in.
ignore_cl  = javax.servlet.http.HttpServlet, weblogic.servlet.JSPServlet,
com.bea.netuix.servlets.manager.UIServlet, com.bea.netuix.servlets.manager.PortalServlet,
com.sap.engine.services.servlets_jsp.server.servlet.InvokerServlet,
com.sapportals.portal.prt.dispatcher.Dispatcher,
com.ibm.wps.engine.Servlet
ignore_tree = org.apache.jasper.runtime.HttpJspBase
method     = !(service)|(do.*)
signature  = !.*
deep_mode  = hard
layer     = Web Tier/Servlet
```

Note: The information in this section is a high-level description of the Capture Points file. For detailed information about instrumenting an application, see the white paper that was included with the product documentation, *Advanced Instrumentation for Mercury Diagnostics for J2EE*.

Understanding Points in the Capture Points File

In order to capture metrics for the classes and methods in your application, the classes and methods must be specified in a point in Capture Points file. A point provides instructions for identifying the classes and methods and assigns the classes and methods to layers. The layers are used to group the performance metrics in a meaningful way by the area of the system where the classes and metrics were being used.

You may customize the points by modifying an existing point or creating a new point to gather metrics in a way that is more meaningful for your applications.

The following example shows a template for a single point definition in the Capture Points file. The arguments that make up the point definition are defined in the tables that follow.

```
[Point-Name]
;-----
class          = <class/package name/s to capture>
method         = <method name/s to capture>
signature      = <signature/s of method/s>
ignore_cl      = <classes to ignore>
ignore_method  = <method prototypes to ignore>
deep_mode      = <soft or hard mode>
scope          = <comma separated list of methods>
detail         = <args,caller,return>
code           = <code snippet for which you want to capture metrics>
code-key       = <key that you generate to secure the code snippet>
tierDefinition = <hex_identifier>
layer          = <layer name>
layerType      = <layer type>
merge_classes  = <class names>
merge_url_to   = <comma separated list of methods>
ignore_tree    = args,caller,return
```

Mandatory Point Entries

Every point definition must have entries for the following arguments:

Argument	Description
class	The class argument specifies the name of the class to be captured. The name should include the full package name using periods between the package levels. A regular expression may be used. Regular expressions are identified by starting the expression with an exclamation mark.
method	The method argument specifies the name of the method to be captured. To be successful, the method name must match a method on the class specified by the class argument. As with the class specification, a regular expression may be used.
signature	The signature argument specifies the signature (parameter and result types) of the method using standard Java encoding for method signatures. In most cases, you will not care about the signature; so you will use the expression: <code>!*</code> to specify any signature. A regular expression may be used.

Optional Point Entries

Point definition may contain entries for one or more of the following arguments:

Argument	Description
ignore_cl	The ignore_cl argument specifies a comma-separated list of classes to ignore. It makes sense to use this only if the class argument is a regular expression. Any class matching one of the classes specified with ignore_cl will not be instrumented. The ignore_cl argument provides the “except these” portion of the specification.
ignore_method	Similar to the ignore_cl argument ignore_method provides a mechanism for specifying methods to ignore.

Argument	Description
deep_mode	<p>The deep_mode argument specifies how subclasses are handled. This is a real time saver but not without danger. Three valid values are accepted for this argument:</p> <ul style="list-style-type: none"> ▶ none - A value of “none” is like not specifying a deep_mode argument. It has no effect on how subclasses are handled. ▶ soft - A value of “soft” requests that for every class or interface matching the class, method, and signature entries, any subclasses or sub-interfaces that also implement the matching method and signature should also be instrumented. ▶ hard - A value of “hard” requests that for every class or interface matching the class, method, and signature entries, any subclasses or sub-interfaces at any depth should have all their methods instrumented. Caution: “hard” mode can lead to extensive instrumentation and very high overhead.
scope	<p>The scope argument constrains the context in which instrumentation is performed. While each of the arguments previously defined arguments specify what to instrument, the scope argument specifies where it should be instrumented. Scope only constrains the caller of a method specified in a point. Without a scope specification, all calls to matching methods are captured. With a scope specification, only calls made from within a specific set of methods are captured. Scope values are a comma-separated list of package, class, and method names in standard Java notation. Note: Using scope requires caller side instrumentation. Regular expression may be used as well.</p>
layer	<p>The layer argument specifies the name of the layer under which the metrics captured by this point are grouped. AD/AM products use this to present layer aggregated information to the user.</p>

Argument	Description
detail	<p>The detail argument specifies more specific capture instructions. It is a comma-separated list of the following keywords that indicate the additional details to be captured:</p> <ul style="list-style-type: none"> ▶ caller - causes caller side instrumentation to be performed. If this keyword is not specified, the default instrumentation, callee side instrumentation, is performed. ▶ args:n - causes the n-th argument value for the captured method to be converted to a string (via the <code>toString()</code> method), and used as the aggregation parameter in the layer argument.(like a SQL query argument). The value for n can be 1 through 256. See the description of the layer argument for more details. ▶ args:0 - causes either <code>toString()</code> on the callee object or the class name of the callee object to be captured, depending on whether or not it is a static method. ▶ args:code - causes the code snippet specified in the code argument to be used as the capture argument. ▶ disabled - causes the point to be placed in a disabled state when the probe is started. This means that the instrumentation is done, but the point will not capture data. All points for a specific layer can be dynamically enabled/disabled by using the Profiler URL <a href="http://<probe_install_dir>:<probe_port>/inst/layer">http://<probe_install_dir>:<probe_port>/inst/layer. ▶ outbound - flags the method so it will appear on the "Outbound Calls" screen ▶ ws-operation - is an internal keyword that you should not use.

Argument	Description
code	The code argument specifies the snippet of code for which you want to capture performance metrics. For information about code snippets see “Defining Points Using Code Snippets” on page 356.
code-key	The code-key argument specifies the secure code key that you generated for the code snippet that you entered for the code argument. Information about code snippets and code keys can be found at “Defining Points Using Code Snippets” on page 356.

AD & AM Only Point Entries

Every AM/AD point definition may have entries for the following arguments:

Argument	Description
<p>layer</p>	<p>The layer argument specifies the name of the layer and sub layers under which the metrics collected by the instrumentation for this point are to be grouped. The layer name appears first. If the name is followed by a / (slash) the name following the slash is a sub-layer. Each sub-layer can also have sub-layers as well.</p> <p>In the UI, this means, for example, that all Web/* layers will show up under the Web layer, and a drilldown will exist from Web to its sub-layers, such as JSP and Struts.</p> <p>Instead of specifying the actual sub-layer name you can specify <code>\${ARG}</code> to request layer name substitution. <code>\${ARG}</code> is replaced with the value of the method argument that is specified in the detail args:n argument.</p> <p>In the following example <code>\${ARG}</code> will be replaced by the second argument of the execute method.</p> <pre>method = execute detail = args:2 deep_mode = soft layer = Business Tier/SAP R3/\${ARG}</pre>
<p>layerType</p>	<p>Points with this argument result in a special layer breakdown for portal components. All necessary instances of these are pre-configured in the auto-detect.points file, and it is important that they are not removed.</p> <p>Other values for layerType are not supported at this time</p>

Deep Diagnostics Only Point Entries

Every point definition that is to be used for DD must have entries for the following arguments:

Argument	Description
detail	<p>This argument specifies additional detail that Deep Diagnostics should capture. It is a comma-separated list of keywords that indicate the additional details to be captured, as follows:</p> <ul style="list-style-type: none"> • args. causes the method argument values for the methods that are captured to be captured as well. • return. If “return” is specified, then the return values of instrumented methods are captured.
tierDefinition	<p>This argument is automatically generated by Deep Diagnostics. It associates the point specification with a Deep Diagnostics tier definition.</p> <p>Important! Do not modify this argument. Do not attempt to provide this value manually. (It is a unique identifier in hexadecimal notation.)</p>

Understanding Diagnostics Layers

Mercury Diagnostics groups the performance metrics for classes and methods into *layers* and *sub-layers* based upon the instructions provided in the points file. The layers have been defined to group together the performance metrics for processing in the application that were performed using similar system resources. The layers make it easier for you to isolate and identify the areas of the system that may be contributing to performance issues.

The following tables list the layers and sub-layers that have been defined for typical J2EE and .NET classes and methods, and indicate the points in the Capture Points file where the classes and methods are assigned to the layers. This information should help you interpret and customize the points in the Capture Points file.

Additional, platform specific layers have also been defined in the Capture Points file. These layers are, for the most part, sub-layers of the top-level parent layers defined in the tables below. For information about the platform specific layers, see the *Mercury Diagnostics User's Guide*.

J2EE Layers

Layer	Sub-Layers	Parent Layer	Defined in Points
Web Tier	JSP Servlets Struts		JSP-_jspService Servlet-all Struts-Action
JSP		Web Tier	JSP-_jspService
Servlets		Web Tier	Servlet-all
Struts		Web Tier	Struts-Action
Business Tier	EJB		EJB-Entity-all EJB-Session-all
EJB	Entity Bean Session Bean	Business Tier	EJB-Entity-all EJB-Session-all
Entity Bean		EJB	EJB-Entity-all
Session Bean		EJB	EJB-Session-all
Directory Service	JNDI		JNDI-lookup
JNDI		Directory Service	JNDI-lookup
Database	JDBC		JDBC-Connection-create JDBC-Connection-prepare JDBC-Statement-execute JDBC-PreparedStatement-execute JDBC-ResultSet JDBC-DataSource-getConnection JDBC-XADataSource-getConnection JDBC-Driver-connect JDBC-OCI-Statement-execute

Layer	Sub-Layers	Parent Layer	Defined in Points
JDBC	Execute Connections	Database	JDBC-Connection-create JDBC-Connection-prepare JDBC-Statement-execute JDBC-PreparedStatement-execute JDBC-ResultSet JDBC-DataSource-getConnection JDBC-XADataSource-getConnection JDBC-Driver-connect JDBC-OCI-Statement-execute
Execute		JDBC	JDBC-Connection-create JDBC-Connection-prepare JDBC-Statement-execute JDBC-PreparedStatement-execute JDBC-ResultSet JDBC-OCI-Statement-execute
Connections		JDBC	JDBC-DataSource-getConnection JDBC-XADataSource-getConnection JDBC-Driver-connect
Messaging	JMS		JMS-MessageProducer JMS-MessageListener JMS-MessageConsumer
JMS	Producer Listener Consumer	Messaging	JMS-MessageProducer JMS-MessageListener JMS-MessageConsumer
Producer		JMS	JMS-MessageProducer
Listener		JMS	JMS-MessageListener
Consumer		JMS	JMS-MessageConsumer

.NET Layers

Layer	Sub-Layers	Parent Layer	Defined in Points
Web Tier	ASP.NET		
Database	ADO		ADO 1 ADO 2 ADO 3 ADO 4
ADO	Execute Connection	Database	ADO 1 ADO 2 ADO 3 ADO 4
Execute		ADO	ADO 1 ADO 2 ADO 3
Connection		ADO	ADO 4
Messaging	MSMQ		MSMQ Synchronous MSMQ Asynchronous MSMQ ReceiveCompleted EventHandler MSMQ PeekCompleted EventHandler
MSMQ		Messaging	MSMQ Synchronous MSMQ Asynchronous MSMQ ReceiveCompleted EventHandler MSMQ PeekCompleted EventHandler

Portal Layers

Diagnostics groups the performance metrics for the classes and method calls associated with processing for portals into Portal Component layers. Each Portal Component layer is broken down to layers for the portal lifecycle methods. For more information about portal layers, see the *Mercury Diagnostics User's Guide*.

Defining Points Using Code Snippets

Securing Code Snippets

By default, you must specify a valid code-key along with the code snippet before the probe will use the code snippet during instrumentation. Requiring the code-key prevents someone from accidentally introducing instrumentation that could significantly increase the overhead of the probe.

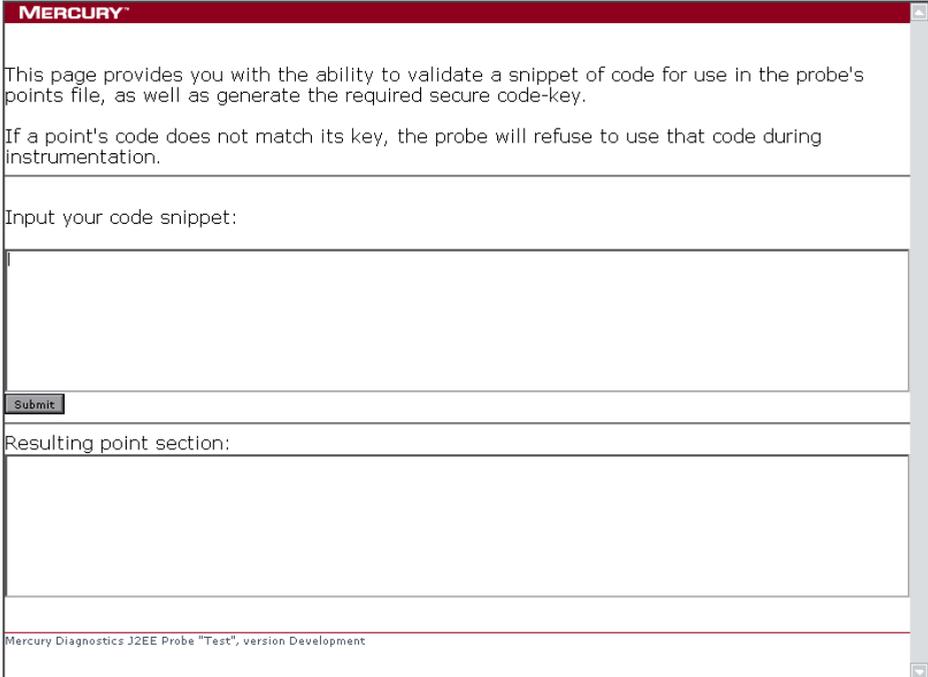
When you generate the code-key, Diagnostics checks the syntax of the code snippet to make sure that it is valid before it generates the key. When Diagnostics instruments your application, it checks the value entered for the code-key argument to make sure that it matches the code-key that it calculates for the code snippet that it finds in the code argument. If the code-keys do not match, Diagnostics ignores the code snippet and does not create the instrumentation point.

Generating the Code-Key

The J2EE Probe is installed with tool that generates the code-key from the code snippet that you input.

To generate a code-key:

- 1 Open the page at the following URL in your browser: <http://<probe-host>:<probe-port>/inst/code-key>. Diagnostics displays the page where you can validate the code snippet syntax and generate the code-key as shown in the following example:



The screenshot shows a web browser window with a red title bar labeled "MERCURY". The page content includes:

- A paragraph: "This page provides you with the ability to validate a snippet of code for use in the probe's points file, as well as generate the required secure code-key."
- A paragraph: "If a point's code does not match its key, the probe will refuse to use that code during instrumentation."
- A label: "Input your code snippet:"
- A large empty text input field.
- A "Submit" button.
- A label: "Resulting point section:"
- A large empty text area for the result.
- A footer: "Mercury Diagnostics J2EE Probe 'Test', version Development"

- 2 Enter the code snippet that you specified in the code argument in the auto_detect_points file into the **Input your code snippet** text box and click **Submit**

Note: The code snippet must include all of the text following the code = argument name.

- 3 Diagnostics presents the results of the code snippet validation and the code-key generation in the **Resulting point section** text box.

If the code snippet was valid, Diagnostics displays the value of both the code-key and code arguments. Enter these value in the points file.

If the code snippet was not valid, Diagnostics displays an error message that indicates the problem that was detected. Correct the problem and click **Submit** again to validate the corrected code.

Disabling the Code-Key Security Check

Diagnostics verifies that the value of the code-key argument matches the value that it generates when it is instrumenting the application by default. It is possible to disable this security check by inserting the **require.code.security.key** property into the `<probe_install_dir>/etc/inst.properties` file with a value of `false`.

Note: You should be very careful when using this property. If you disable this check, you could allow unexpected processing overhead and unpredictable performance monitoring results.

18

Using Solution Templates

This chapter describes how to work with solution templates in Mercury Diagnostics.

This chapter describes:	On page:
About Solution Templates	359
Understanding Solution Template Data	360
Viewing Solution Template Information	361

About Solution Templates

Solution templates are predefined, out-of-the-box instrumentation dashboard and report templates for leading application middleware servers and frameworks. Mercury Diagnostics currently provides solution templates for:

- ▶ **WebLogic:** The WebLogic Portal Server solution template displays information for BEA WebLogic Portal Server WL 8.1 SP3, SP4.
- ▶ **WebSphere:** The WebSphere Portal Server solution template displays information for WebSphere 5.1.
- ▶ **Oracle 11i:** The Oracle 11i solution template displays information for the Web-based applications framework.
- ▶ **SAP:** The SAP solution template displays information for iViews in the SAP Web Application Server 6.40 SP 10, or later.

Understanding Solution Template Data

Mercury Diagnostics groups the performance metrics for classes and methods into *layers* and *sub-layers*, based upon the resources that the application invokes to perform the processing. These layers make it easier to isolate and identify the areas of the system that may be contributing to performance issues.

The solution templates let you display business services associated with various middleware servers and frameworks in new layers across the Diagnostics application.

Note: For more information about Diagnostics layers and about how classes and methods are assigned to layers in the Capture Points file, see Chapter 17, “Probe Instrumentation.”

Viewing Solution Template Information

Mercury Diagnostics displays unique information for each of the middleware servers and frameworks supported by solution templates.

Note: You can also view graphs of WebLogic and SAP portal data in Mercury Diagnostics on the **Portal** screen. For more information about using the Portal screen, refer to the *Mercury Diagnostics User's Guide*.

WebLogic

The WebLogic Portal Server Solution Template lets you display the following business services in new layers across the Diagnostics application:

Service	Layer	Description
Entitlements	Portal/BEA/Entitlement	entitlement API's
Content management	Portal/BEA/Content	content management API's
User profiles	Portal/BEA/UserProfile	user profile API's
Personalization	Portal/BEA/Personalization	personalization API's
Portal components	Portal/BEA/Portlet	portal component rendering API's

Note: The Portal/BEA/Web Services layer, representing BEA Web Service, is commented out of the points file by default.

The following example of a screen in Mercury Diagnostics displays layers that represent WebLogic business services.



WebSphere

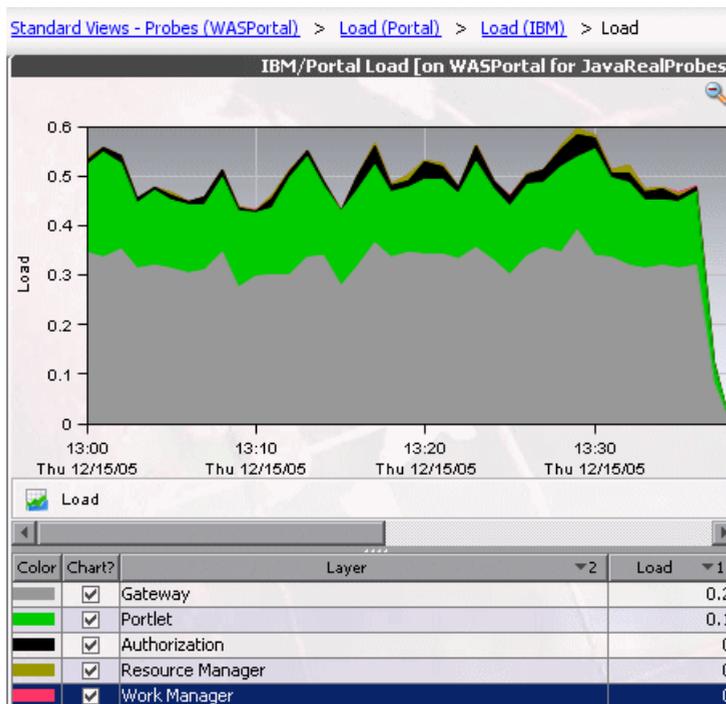
The WebSphere Portal Server Solution Template lets you display the following business services in new layers across the Diagnostics application:

Service	Layer
WAS Portal Gateway Servlet	Portal/IBM/Gateway
WAS Portal Authentication	Portal/IBM/Authentication
WAS Portal PortalAuthorization	Portal/IBM/Authorization
WAS Portal Resource Manager	Portal/IBM/ResourceManager
WAS Portal Portlet Adapter Interface Actions	Portal/Jetspeed/Portlet/Action
WAS Portlet Lifecycle	Portal/Jetspeed/Portlet/Lifecycle
WAS Portal Rendering	Portal/IBM/Portlet/Rendering
WAS Portal Work Manager	Portal/IBM/WorkManager
WAS Portal Portlet Services	Portal/IBM/Portlet/Services
WAS Portal Phases	Portal/IBM/Portlet/Phases
WAS Portal Services	Portal/IBM/Portlet/Services
WAS Portal Portlet Filter	Portal/IBM/Portlet/Filter

Note: The Portal/IBM/Portlet/Phases layer, representing WAS Portal Phases, is commented out of the points file by default. Portal Phases do not need to be captured unless you want to understand how the container processes each phase of the portlet's Init-Action-Render cycle.

The Portal/IBM/Portlet/Services layer, representing WAS Portal Services, is also commented out of the points file by default. Many portal services such as the AccessControlService service and PortletFilterService are exposed in other layers, but you might want to uncomment this section if you are interested in some of the other services provided by the portal container.

The following example of a screen in Mercury Diagnostics displays layers that represent WebSphere business services.



Oracle

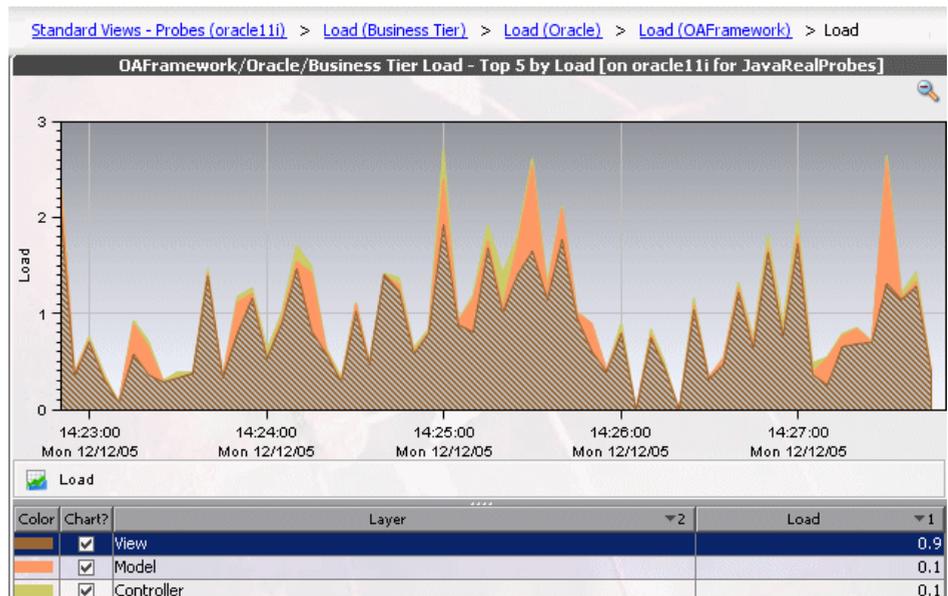
The Oracle 11i Solution Template lets you display the following two layers found underneath the root **Business Tier** layer:

- ▶ Caching
- ▶ OAFramework (Oracle Application Framework)

You can drill down further into the OAFramework layer to display the following three layers:

- ▶ Model
- ▶ View
- ▶ Controller

The MVC layers (Model, View, and Controller) relate to the different parts of the Oracle 11i application framework. The following example of a screen in Mercury Diagnostics displays the MVC layers.



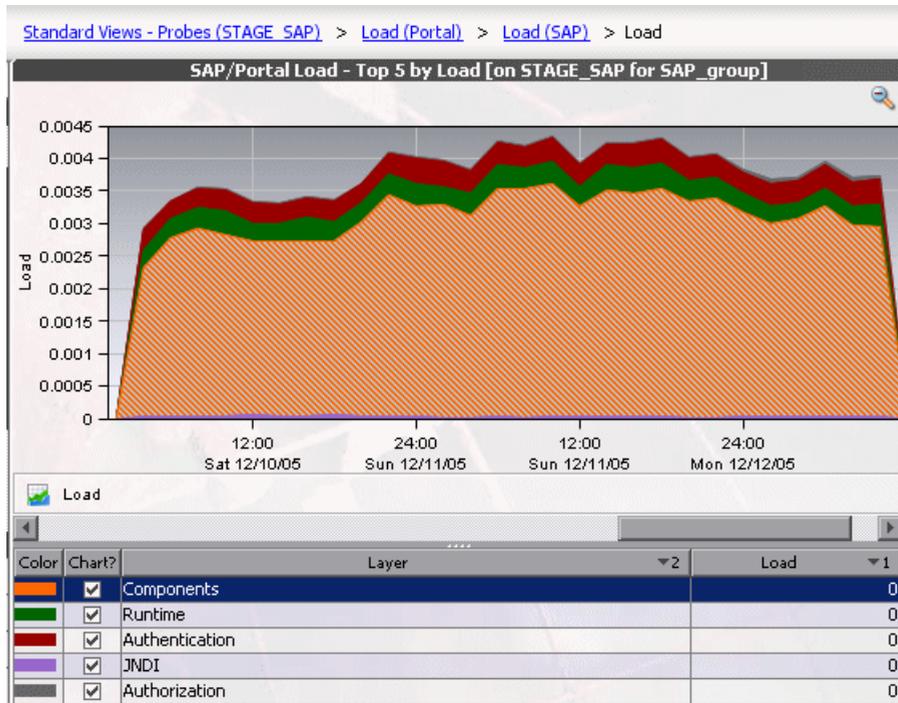
SAP

The SAP Portal Server Solution Template lets you display the following business services in new layers across the Diagnostics application:

Service	Layer	Comment
Runtime	Portal/SAP/Runtime	General runtime part of infrastructure
Authorization	Portal/SAP/Authorization	Core application component dedicated to authorization services
Authentication	Portal/SAP/Authentication	Core application component dedicated to authentication services
JNDI Services:	Portal/SAP/JNDI	Core application component dedicated to JNDI services
Components Runtime	Portal/SAP/Components/Runtime	General runtime part of Portal components
Content Generation (Portal Components)	Portal/SAP/Components/ContentGeneration	Content generation part of Portal components
Content Generation (Service Components)	Portal/SAP/Components/Service	Content generation Service components (Note: Commented out. Optimize before use)
Content Generation (Request Object Specific)	Portal/SAP/Components/Request	Content generation Request object specific (Note: Commented out. Optimize before use)
Content Generation (Profile Components)	Portal/SAP/Components/Profile	Content generation Profile components
Content Generation (Config Components)	Portal/SAP/Components/Config	Content generation configuration components

Service	Layer	Comment
Content Generation (Response Object Specific)	Portal/SAP/Components/Response	Content generation Response object specific
R3 Content Generation	Business Tier/SAP R3	Content generation specific to R3 backend
Dynpage Content Generation	Web Tier/SAP/DynPage	Content generation specific to DynPage
JSP Dynpage Content Generation	Web Tier/SAP/JSPDynPage	Content generation specific to JSP DynPage
Page Processor Component	Web Tier/SAP/Page Processor	Content generation specific to the Page Processor

The following example of a screen in Mercury Diagnostics, displays layers that represent SAP business services. The names of the layers are equivalent to the names of the remote function calls (RFCs).



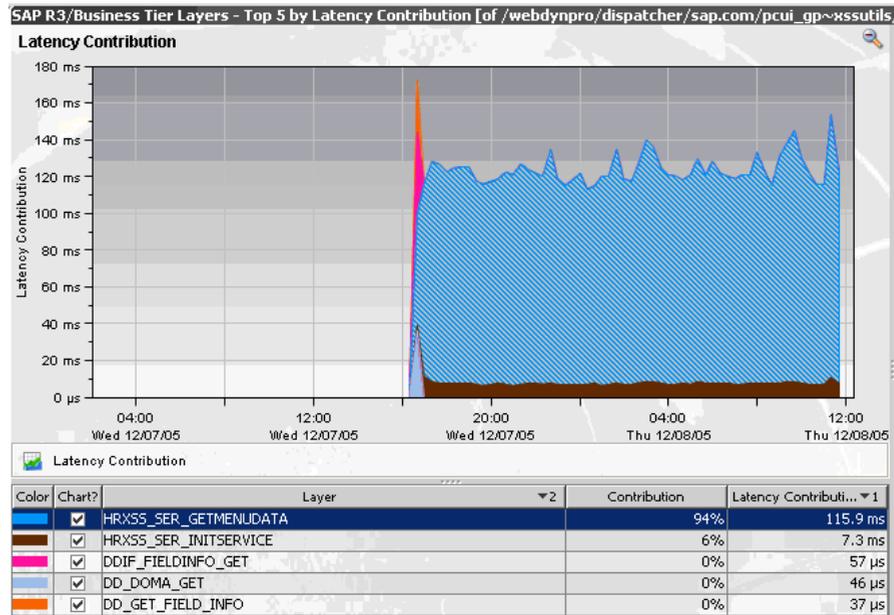
SAP R3 Layer

In a typical SAP Web Application Server deployment (WAS), the SAP WAS is connected to a backend SAP R3 Server. Communication with the SAP R3 Server takes place through RFCs, using SAP JCO (Java Connector). Mercury Diagnostics captures each distinct call and presents it as sub-layers.

You access these sub-layers by drilling down from the Business Tier layer. For example, if you start from the Server Summary screen in Mercury Diagnostics, you can drill down to the SAP R3 sub-layers as follows:

Standard views - Server Summary > Server Requests > Layers (Business Tier) > Layers (SAP R3) > Layers.

The following example of a screen in Mercury Diagnostics, displays SAP R3 sub-layers.



Part VI

Advanced Component Communication

19

Enabling HTTPS Between Diagnostics Components

This chapter describes the configuration steps that you must perform to enable HTTPS communications between the Mercury Diagnostics components.

This chapter describes:	On page:
About Configuring HTTPS Communications	374
Enabling Incoming HTTPS Communication for Diagnostics Components	375
Enabling Outgoing HTTPS Communication from Diagnostics Components	379
Enabling HTTPS Communications for the Mercury Business Availability Center Server	382

Note: The configuration instructions described in this chapter are intended for experienced users with in-depth knowledge of Mercury Diagnostics. Please use caution when modifying any configuration settings for the Diagnostics components.

About Configuring HTTPS Communications

To configure HTTPS Communications between the Diagnostics components you must enable SSL on the Diagnostics Server in Commander mode and on all Diagnostics Servers and probes connected to the Diagnostics Server in Commander mode. You must also configure each component so that it will be able to communicate with components that have the proper certificates.

The instructions for configuring each type of component contain details of the following main steps:

- ▶ Generate a keystore on the component.
- ▶ Export the certificate from the keystore
- ▶ Obfuscate passwords that provide access to the keystore.
- ▶ Copy the component's certificate to the Diagnostics components that will initiate communication.
- ▶ Configure the component's security properties to enable SSL and provide the passwords necessary for HTTPS communication to take place.

Enabling Incoming HTTPS Communication for Diagnostics Components

The following instructions provide you with the steps necessary to configure the Diagnostics Server or J2EE Probe to receive incoming HTTPS communications from other Diagnostics components.

Notes:

- ▶ Enabling SSL and HTTPS Communications for the J2EE Probe is supported on SUTs with the Sun and JRockit JVMs. However, if you are using a JVM version prior to 1.4, you must download and install the Sun JSSE Optional Package onto the SUT server before you can enable SSL. Other JSSE implementation, such as IBM's are not supported.
 - ▶ The Diagnostics Server in Commander mode and the Diagnostics Server in Mediator mode must also be configured to communicate via HTTPS when you configure the probe.
-

To configure the Diagnostics Server or J2EE Probe for incoming HTTPS connections:

- 1** Generate a keystore in the `<diagnostics_server_install_dir>/etc` directory or `<J2EE_probe_install_dir>/etc` directory using the following command:
 - a** Replace `<diagnostics_server_install_dir>` with the path to the installation directory for the Diagnostics Server or replace `<J2EE_probe_install_dir>` with the path to the installation directory for the J2EE Probe.
 - b** Replace `<diagnostics_server_hostname>` with the machine name for the host of the Mercury Diagnostics Server or replace `<probe_hostname>` with the machine name for the host of the J2EE Probe. This value cannot be the server's IP address.
 - c** Replace each occurrence of `<password>` with the same password string. If you want, you may assign different passwords to the **storepass** and the **keypass**.

Command for Diagnostics Server:

```
<diagnostics_server_install_dir>/_jvm/bin/keytool -keystore  
<diagnostics_server_install_dir>/etc/keystore -storepass <password> -alias SERVER  
-genkey -keyalg RSA -keypass <password> -dname  
"CN=<diagnostics_server_hostname>, OU=Diagnostics, O=Mercury, L=Mountain  
View, S=CA, C=USA" -validity 3650
```

After you execute this command, a keystore is created in
<diagnostics_server_install_dir>/etc/keystore with an entry called
SERVER for the host of the Diagnostics Server.

Command for J2EE Probe:

```
<probe_install_dir>/_jvm/bin/keytool -keystore <probe_install_dir>/etc/keystore -  
storepass <password> -alias PROBE -genkey -keyalg RSA -keypass <password> -  
dname "CN=<probe_hostname>, OU=Diagnostics, O=Mercury, L=Mountain View,  
S=CA, C=USA" -validity 3650
```

After you execute this command, a keystore is created in
<probe_install_dir>/etc/keystore with an entry called **PROBE** for the host
of the J2EE Probe.

- 2 Export the certificate for the SERVER or PROBE entry in the keystore using the following command.
 - a Replace **<diagnostics_server_install_dir>** with the path to the installation directory for the Diagnostics Server, or replace **<probe_install_dir>** with the path to the installation directory for the J2EE Probe.
 - b Replace **<password>** with the string that you assigned as the **storepass** password when you created the keystore.
 - c Replace **<server_certificate_name>** or **<probe_certificate_name>** with the name that you would like to assign to the certificate file. It is recommended that you assign a certificate name that will make it easy to recognize the component for which the certificate was created.

For Diagnostics Servers, use `diag_server_commander.cer` or `diag_server_mediator.cer`.

For J2EE Probes, include the type of the probe and the host name for the probe so that it will be easy to recognize the component for which the certificate was created. For example: `j2EE_probe_<probe_hostname>`.

Command for Diagnostics Server:

```
<diagnostics_server_install_dir>/_jvm/bin/keytool -keystore
<diagnostics_server_install_dir>/etc/keystore -storepass <password> -alias
SERVER -export -rfc -file
<diagnostics_server_install_dir>/etc/<server_certificate_name>.cer
```

After this command runs, a certificate file with the name assigned in `<server_certificate_name>` is created in the `<diagnostics_server_install_dir>/etc` directory for the Diagnostics Server, for example, `diag_server_commander.cer`.

Command for J2EE Probe:

```
<probe_install_dir>/_jvm/bin/keytool -keystore <probe_install_dir>/etc/keystore -
storepass <password> -alias PROBE -export -rfc -file
<probe_install_dir>/etc/<probe_certificate_name>.cer
```

After this command runs, a certificate file called `j2ee_probe_<probe_hostname>.cer` is created in the `<probe_install_dir>/etc` directory for the J2EE Probe.

Note: The certificate file must be imported to the host machines for each of the Diagnostics components that are expected to initiate communications with the Diagnostics Server or J2EE Probe. The instructions for importing the certificate file to each Diagnostics component are provided below.

- 3 Using the command in the following example, generate an obfuscated version of the **storepass** and the **keypass** passwords that you assigned when you created the keystore.
 - a Replace **<diagnostics_server_install_dir>** with the path to the installation directory for the Diagnostics Server, or replace **<probe_install_dir>** with the path to the installation directory for the J2EE Probe .
 - b Replace **<password>** with the string that you assigned as the password when you created the keystore.

```
<diagnositcs_server_install_dir>/_jvm/bin/java -cp  
<diagnositcs_server_install_dir>/lib/ThirdPartyLibs.jar org.mortbay.util.Password  
<password>
```

The output from the obfuscation is shown in the following example. In this example, the password string was “testpass”. The output consists of three lines. The original string that was to be obfuscated and two lines depicting the obfuscated password. Only the line that begins with “OBF” is used to set the properties in the following step of this process.

```
testpass  
OBF:1ytc1vu91v2p1y831y7v1v1p1v11yta  
MD5:179ad45c6ce2cb97cf1029e212046e81
```

Note: If you did not use the same password for **keypass** and **storepass**, you must run this command twice to create an obfuscated version for each password.

- 4 Change the following properties in the file `<diagnostics_server_install_dir>/etc/security.properties` for the Diagnostics Server in Commander mode, or `<probe_install_dir>/etc/security.properties` for the Mercury Diagnostics Probe for J2EE.
 - a Set `enableSSL=true`.
 - b Set `keyStorePassword=<obfuscated_password>`
 - c Set `keyPassword=<obfuscated_password>`

Note: The value entered for `<obfuscated_password>` must include the entire “OBF” line that was output from the command in the previous step. For example:

```
keyStorePassword=OBF:1ytc1vu91v2p1y831y7v1v1p1vv11yta
```

Enabling Outgoing HTTPS Communication from Diagnostics Components

The following instructions provide you with the steps necessary to configure the Diagnostics Server and J2EE Probe to send outgoing HTTPS communications to the other Diagnostics components.

You perform these steps for each component that will connect to the Diagnostics Server or J2EE Probe, where the Diagnostics Server or J2EE Probe will initiate communication.

To enable the Diagnostics Server in Commander mode for outgoing communication to the Diagnostics Server in Mediator mode.

- 1 Copy the certificate file from `<diagnostics_server_install_dir>/etc/diag_server_mediator.cer` on the Diagnostics Server to `<diag_server_install_dir>/etc/diag_server_mediator.cer` on the Diagnostics Server in Commander mode.

- 2 Change the value of the **trusted.certificate** property in the file `<diagnostics_server_install_dir>/etc/security.properties` for the Diagnostics Server in Commander mode.

Set **trusted.certificate=diag_server_mediator.cer**. If there are already other certificate files included in the value of this property, add the certificate file to the end of the list separated from the preceding value by a comma.

- 3 For incoming Diagnostics Server communication, indicate the URL for the Diagnostics Server by updating the following property in the file `<diagnostics_server_inst_dir>/etc/server.properties`.

Set **commander.url** to **https://<diagserver_hostname>:8443**

To enable the Diagnostics Server in Mediator mode for outgoing communication to the Diagnostics Server in Commander mode:

- 1 Copy the certificate file from `<diagnostics_server_install_dir>/etc/diag_server_commander.cer` on the Diagnostics Server in Commander mode to `<diag_server_install_dir>/etc/diag_server_commander.cer` on the Diagnostics Server in Mediator mode.

- 2 Change the value of the **trusted.certificate** property in the file `<diagnostics_server_install_dir>/etc/security.properties` for the Diagnostics Server in Mediator mode.

Set **trusted.certificate=diag_server_commander.cer**. If there are already other certificate files included in the value of this property, add the certificate file to the end of the list separated from the preceding value by a comma.

- 3 For incoming Diagnostics Server communication, indicate the URL for the Diagnostics Server by updating the following property in the file `<diagnostics_server_inst_dir>/etc/server.properties`.

Set **commander.url** to **https://<diagserver_hostname>:8443**

To enable the Diagnostics Server in Mediator mode for outgoing communications to the probes:

- 1** Copy the certificate file from `<probe_install_dir>/etc/j2ee_probe_<probe_host>.cer` for each probe to `<diag_server_install_dir>/etc/j2ee_probe_<probe_host>.cer` on the Diagnostics Server in Mediator mode.
- 2** Change the value of the `trusted.certificate` property in the file `<diagnostics_server_install_dir>/etc/security.properties` for the Diagnostics Server in Mediator mode.

Set `trusted.certificate=j2ee_probe_<probe_host>.cer` If there are already other certificate files included in the value of this property, add the certificate file to the end of the list separated from the preceding value by a comma.

To enable the J2EE Probe for outgoing communications to the Diagnostics Server in Mediator mode:

- 1** Copy the certificate file from `<diagnostics_server_install_dir>/etc/diag_server_mediator.cer` on the Diagnostics Server in Mediator mode to `<probe_install_dir>/etc/diag_server_mediator.cer` on the J2EE Probe.
- 2** Change the value of the `trusted.certificate` property in the file `<probe_install_dir>/etc/security.properties` for the J2EE Probe.

Set `trusted.certificate=diag_server_mediator.cer`. If there are already other certificate files included in the value of this property, add the certificate file to the end of the list separated from the preceding value by a comma.

- 3** For incoming J2EE Probe communication, Indicate the URL for the Diagnostics Server in Mediator mode by updating the following property in the file `<probe_inst_dir>/etc/dispatcher.properties`.

Set `registrar.url` to

`https://<diagserv_mediatormode_hostname>:8443/commander/registrar/`

Enabling HTTPS Communications for the Mercury Business Availability Center Server

The following instructions will guide you through the process of configuring Mercury Business Availability Center for HTTPS communication with Diagnostics.

To enable HTTPS communications between the Diagnostics Server in Commander mode and Mercury Business Availability Center:

- 1** Copy the Diagnostics certificate file, **diag_server_commander.cer**, from the Diagnostics Server in Commander mode installation directory, **<diagnostics_server_install_dir>/etc/**, to the Mercury Business Availability Center host.
- 2** Import the copied certificate, **diag_server_commander.cer**, into the Mercury Business Availability Center server cacert keystore by running the following command on the Mercury Business Availability Center host:

```
<BAC_server_install_dir>/_jvm/bin/keytool -import -file  
<copied_diag_certificate_directory>/diag_server_commander.cer -keystore  
<BAC_server_install_dir>/jre/lib/security/cacerts -alias SERVER
```

- Replace **<BAC_server_install_dir>** with the path to the installation directory for Mercury Business Availability Center.
- Replace **<copied_diag_certificate_directory>** with the path to the copied Diagnostics certificate file.

Type **changeit** when you are prompted to enter the keystore password.

Type **yes**, instead of the default **no** when you are asked if the certificate should be trusted.

- 3** Copy the Mercury Business Availability Center certificate file, **<BAC_certificate_file.cer>**, to the Diagnostics Server host.

- 4** Import the copied certificate into the Diagnostics Server cacert keystore by running the following command on the Diagnostics Server host.

```
<diagnostics_server_install_dir>/_jvm/bin/keytool -import -file
<copied_BAC_certificate_directory>/<BAC_certificate_file.cer> -keystore
<diagnostics_server_install_dir>/JRE/lib/security/cacerts
```

- Replace **<diagnostics_server_install_dir>** with the path to the installation directory of the Diagnostics Server in Commander mode.
- Replace **<copied_BAC_certificate_directory>** with the path to the copied Mercury Business Availability Center certificate file.

When you are prompted to enter the keystore password, type the string that you assigned as the **storepass** password when you created the keystore.

Type **yes**, instead of the default **no**, when you are asked if the certificate should be trusted:

- 5** Point the Mercury Business Availability Center server to the HTTPS port on the Diagnostics Server in Commander mode.
- a** Open the Diagnostics Administration in Mercury Business Availability Center by selecting **ADMIN > Diagnostics**.
 - b** Click the **Registration** tab.
 - c** Locate the Diagnostics Server details section.
 - d** Provide the following information in the appropriate fields:
 - Enter the host name for the Diagnostics Server in Commander mode exactly as it was specified in the **CN** parameter when you created the keystore for the Diagnostics Server in Commander mode. See “Enabling Incoming HTTPS Communication for Diagnostics Components” on page 375
 - Enter HTTPS for the protocol
 - Enter 8443 for the Web port of the Diagnostics Server.
 - e** Click **Submit Configuration**.

20

Configuring Diagnostics Components for HTTP Proxy

This chapter describes the configuration steps that you must perform to enable HTTP proxy communications between the Mercury Diagnostics components.

This chapter describes:	On page:
Enabling HTTP Proxy Communications for the Diagnostics Servers	386
Enabling HTTP Proxy Communications for the J2EE Probe	387
Enabling HTTP Proxy Communications for a .NET Probe	387

Note: The configuration instructions described in this chapter are intended for experienced users with in-depth knowledge of Mercury Diagnostics. Please use caution when modifying any configuration settings for the Diagnostics components.

Enabling HTTP Proxy Communications for the Diagnostics Servers

The following section describes how to configure the Diagnostics Servers in Commander mode and Diagnostics Servers in Mediator mode to communicate with each other through an HTTP proxy.

To configure the Diagnostics Servers for HTTP proxy communications:

- 1** Set the following properties in `<diagnostics_server_install_dir>/etc/server.properties`:
 - Set **proxy.host** to the host name of the proxy server.
 - Set **proxy.port** to the port of the proxy server.
 - Set **proxy.protocol** to the protocol to use for the proxy server (http).
 - Set **proxy.user** to the user used to authenticate the proxy server.
 - Set **proxy.password** to the password used to authenticate the proxy server.
- 2** Restart the Diagnostics Server. For instructions, see “Starting and Stopping the Diagnostics Server” on page 52.

Enabling HTTP Proxy Communications for the J2EE Probe

The following section describes how to configure the J2EE Probe to communicate with Diagnostics Server in Commander mode through an HTTP proxy:

To configure the J2EE Probe for HTTP proxy communications:

- 1** Set the following properties in `<J2ee_install_dir>/etc/dispatcher.properties`:
 - ▶ Set **proxy.host** to the host name of the proxy server.
 - ▶ Set **proxy.port** to the port of the proxy server.
 - ▶ Set **proxy.protocol** to the protocol to use for the proxy server (http).
 - ▶ Set **proxy.user** to the user used to authenticate the proxy server.
 - ▶ Set **proxy.password** to the password used to authenticate the proxy server.
- 2** Restart the instrumented application VM.

Enabling HTTP Proxy Communications for a .NET Probe

The following section describes how to configure the .NET Probe to communicate with the Diagnostics Server in Commander mode through an HTTP proxy:

To configure the .NET Probe for HTTP proxy communications:

- 1** Set the **proxy** property in `<probe_install_dir>/etc/probe_config.xml` to point to the Diagnostics Server host:

```
<registrar url="http://<diagserver_host_name>:2006/registrar/"
proxy="http://proxy:8080" />
```

- 2** Restart the instrumented application process.

21

Configuring Diagnostics Components to Work with a Firewall

This chapter describes the configuration steps that you must perform to enable Mercury Diagnostics to work correctly in an environment where a firewall is present. This additional configuration is required when the firewall separates the probes from the other Diagnostics components or the components of LoadRunner, Performance Center, or Mercury Business Availability Center.

This chapter describes:	On page:
Overview of Configuring Diagnostics for a Firewall	390
Collating Offline Analysis Files over a Firewall	393
Installing and Configuring the Mercury MI Listener	393
Configuring the Diagnostics Server in Mediator mode to Work with a Firewall	394
Configuring LoadRunner and Performance Center to Work with Diagnostics Firewalls	400

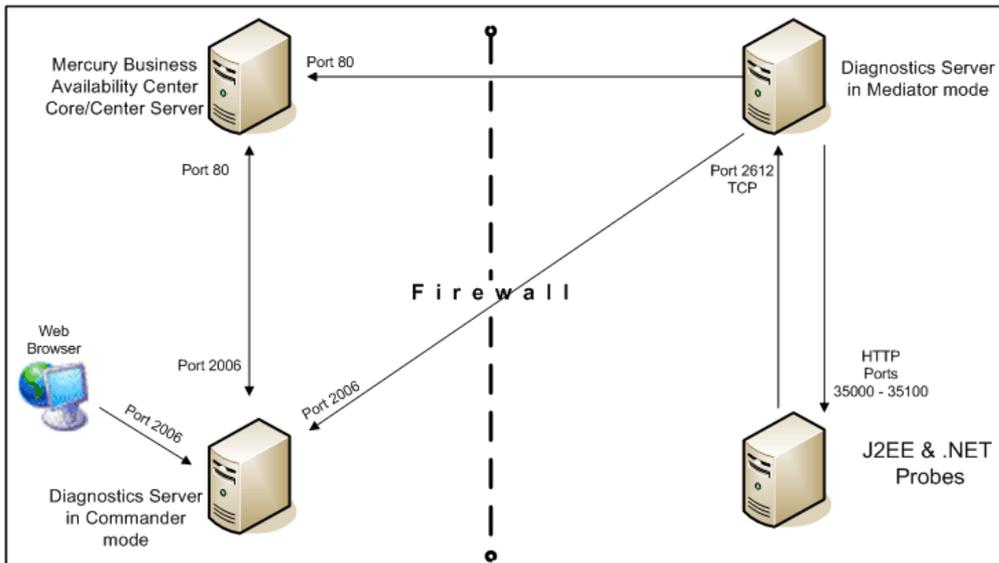
Note: The configuration instructions described in this chapter should be used only by experienced users with in-depth knowledge of Mercury Diagnostics. Please use caution when modifying any configuration settings for the Diagnostics components.

Overview of Configuring Diagnostics for a Firewall

Configuring Diagnostics for a firewall differs depending on which Mercury Product is part of the Diagnostics integration.

Mercury Business Availability Center

The diagram below shows a typical Diagnostics topology where a firewall separates the probe from the other Diagnostics and Mercury Business Availability Center components.



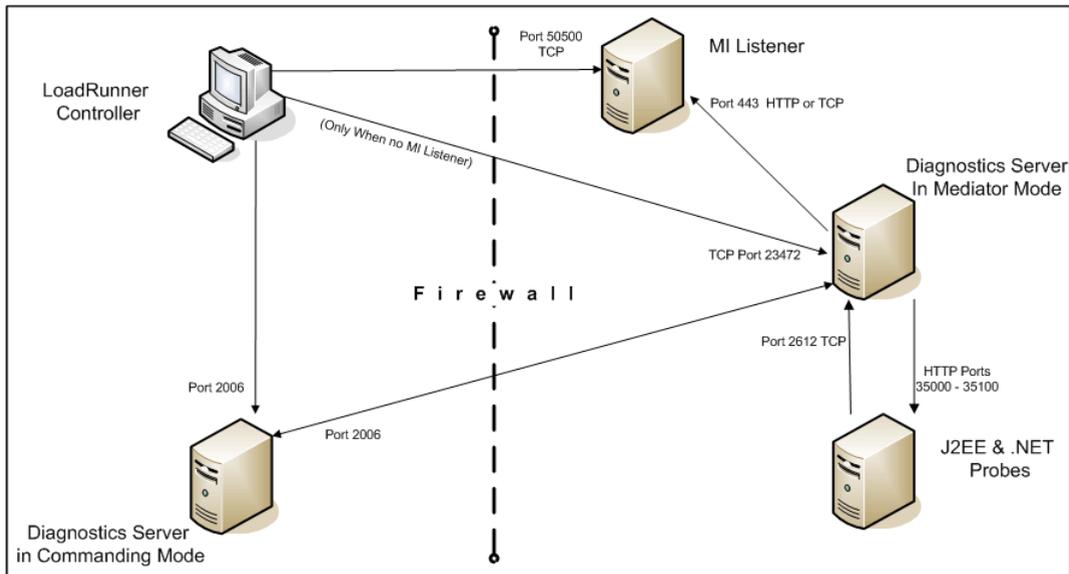
To configure the firewall to enable the communications between the Diagnostics components, open the ports that will allow:

- ▶ HTTP requests from the Mercury Business Availability Center Core/Center server(s) to the Diagnostics Server in Commander mode, on port 2006.
- ▶ HTTP requests from the Diagnostics Server in Commander mode to Mercury Business Availability Center core server on port 80.
- ▶ HTTP requests from the Diagnostics Server in Mediator mode to ports 35000-35100 of probe.

- ▶ HTTP requests from the Web Browser Client machine to the Diagnostics Server in Commander mode on port 2006.
- ▶ HTTP requests from the Diagnostics Server in Mediator mode to the Diagnostics Server in Commander mode on port 2006.
- ▶ TCP requests from the probe to the Diagnostics Server in Mediator mode on port 2612.
- ▶ HTTP requests from the probe to the Diagnostics Server in Mediator mode on port 2006.
- ▶ HTTP requests from the Diagnostics Server in Commander mode to the probes on port 35000-35100. The actual ports on which you must allow communications will depend upon the port numbers that you enabled when you configured the probe and the number of instrumented VMs. For information on setting the probe port range, see “Configuring the Probes for Multiple Application Server JVM Instances” on page 312.

LoadRunner and Performance Center

The diagram below shows a typical Diagnostics topology where a firewall separates the probe from the other Diagnostics and LoadRunner components.



Note: LoadRunner is used in this diagram for illustrative purposes. The same information would apply to Performance Center.

You must configure the firewall to allow the Diagnostics components to communicate with each other.

To configure the firewall to enable the communications between the Diagnostics components, open the ports that will allow:

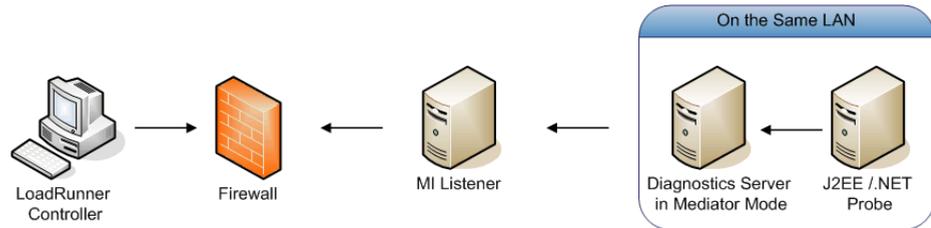
- ▶ HTTP requests from the Diagnostics Server in Mediator mode to the Diagnostics Server in Commander mode on ports 2612 and 2006.
- ▶ TCP requests from the probe to the Diagnostics Server in Mediator mode on port 2612.
- ▶ HTTP requests from the probe to the Diagnostics Server in Mediator mode on port 2006.
- ▶ HTTP requests from the Diagnostics Server in Commander mode to the probes on port 35000-35100. The actual ports on which you must allow communications will depend upon the port numbers that you enabled when you configured the probe and the number of instrumented VMs. For information on setting the probe port range, see “Configuring the Probes for Multiple Application Server JVM Instances” on page 312.

Note: In addition to the above topology, if you are using the LoadRunner Analysis Tool to view offline J2EE results, see “Collating Offline Analysis Files over a Firewall” to properly configure the Controller and the Diagnostics Servers in Mediator mode for offline file retrieval.

Collating Offline Analysis Files over a Firewall

During a LoadRunner / Performance Center load test, the Diagnostics Servers that have probes reporting to them generate an offline analysis file on their host machine. The offline analysis file is retrieved by LoadRunner / Performance Center when it collates the results of a load test.

If there is a firewall between the LoadRunner / Performance Center Controller and the Diagnostics Server involved in a load test, you must configure the Controller and the Diagnostics Server to use the MI Listener utility to enable the transfer of the offline analysis file. The MI Listener utility comes with LoadRunner / Performance Center and should be installed on a machine inside your firewall as shown in the following diagram.



To configure the Controller to access Diagnostics Servers that are behind a firewall:

- Install and configure the Mercury MI Listener.
- Configure the Diagnostics Servers to work with a firewall.
- Configure LoadRunner / Performance Center to work with a firewall.

Installing and Configuring the Mercury MI Listener

The Mercury MI Listener component is the same component that is used to serve Load Generators that are outside of a firewall. For more information about how to configure the MI Listener for LoadRunner, refer to the Mercury LoadRunner Controller User's Guide. For more information about how to configure the MI Listener for Performance Center, refer to the Mercury Performance Center System Configuration and Installation Guide.

Configuring the Diagnostics Server in Mediator mode to Work with a Firewall

To configure the Diagnostics Server in Mediator mode so that it can work across a firewall, you must complete the following additional configuration steps. If you have not yet installed and configured the Diagnostics Server in Mediator mode you must do so prior to attempting these steps. For instructions on installing the Diagnostics Server in Mediator mode, see Chapter 2, “Installing the Mercury Diagnostics Server.”

To configure the Diagnostics Server in Mediator mode for a firewall on a Windows machine:

- 1 Modify the `<diagnostics_server_install_dir>\nanny\windows\.`
- 2 Modify the `<diagnostics_server_install_dir>\windows\nanny\dat\nanny\server.nanny` file.

- ▶ Locate the `start_nt` line:

```
start_nt="C:\MercuryDiagnostics\Server\jre\bin\javaw.exe" -server
-Xmx512m ....
```

- ▶ Remove the quotes that surround the `java.exe` path and add the caret (^) after `javaw.exe` as shown below:

```
start_nt=C:\MercuryDiagnostics\Server\jre\bin\javaw.exe^ -server
-Xmx512m...
```

- 3 Modify the `<diagnostics_server_install_dir>/etc/server.properties` file.
 - ▶ Uncomment the `dispatcher.offline.locationprefix` property.
 - ▶ Set the value of the property to `C:/Documents and Settings/Default User/Local Settings/Temp/MOfflineFiles` as shown in the following example. Make sure that you use the correct drive letter and path for your installation.

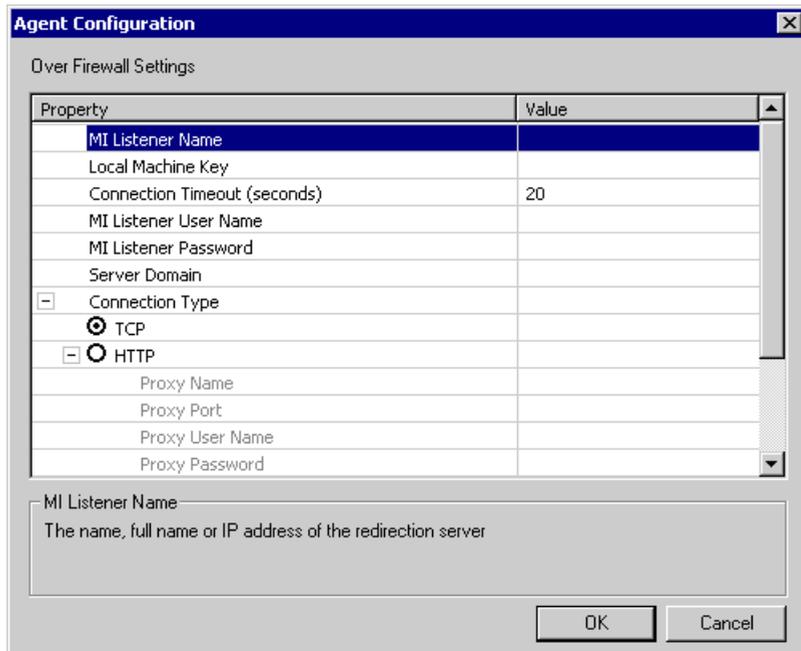
```
dispatcher.offline.locationprefix= C:/Documents and Settings/Default
User/Local Settings/Temp/MOfflineFiles
```

- 4 Launch the Agent Configuration by running `<diagnostics_server_install_dir>/nanny/windows/bin/AgentsConfig.exe`.

The Agent Configuration process opens the following dialog box opens:



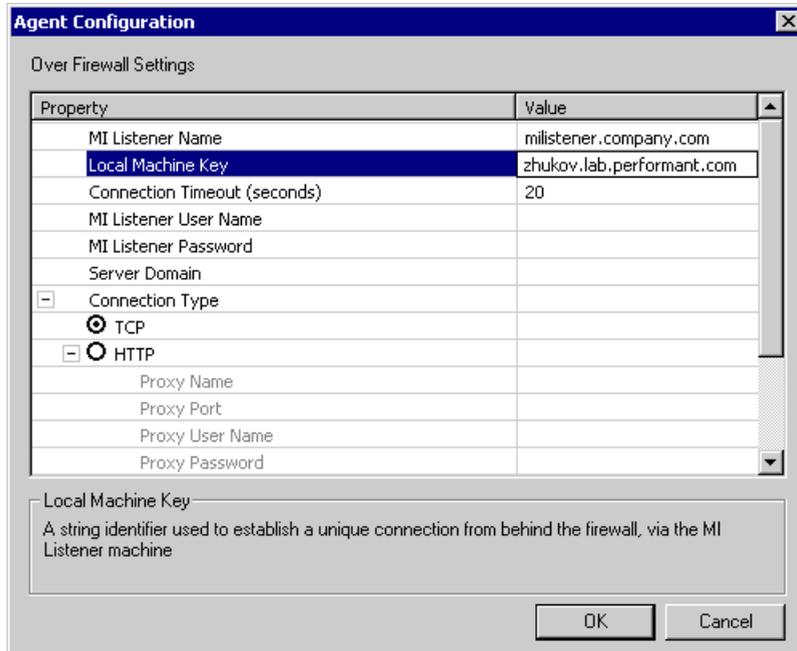
- 5 Select **Enable Firewall Agent**. The **Settings** button becomes enabled.
- 6 Click **Settings**. The Agent Configuration process opens the Agent Configuration Settings dialog box.
- 7 In Value column of the MI Listener Name property, enter the host name or IP address of the machine where the MI Listener was installed.



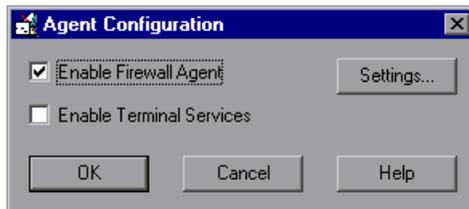
- 8 For the **Local Machine Key** property, enter the machine name of the host of the Diagnostics Server in Mediator mode.

Use the System Health Monitor to determine the machine name for the host of the Diagnostics Server in Mediator mode. For more information on the System Health Monitor, see Appendix C, “Using the System Health Monitor.”

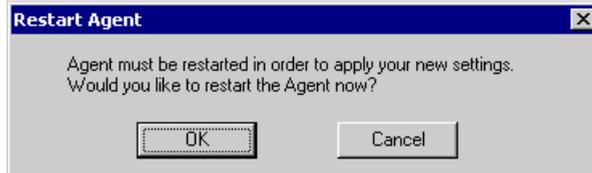
- 9 Click OK to close the setting dialog box.



- 10 Click OK again to close the Agent Configuration dialog box.



- 11 The Restart Agent dialog box opens. Click **OK** to restart the Agent.



- 12 Restart the Mercury Diagnostics Server service.

To configure the Diagnostics Server in Mediator mode for a firewall on a Solaris machine:

- 1 Modify the `<diagnostics_server_install_dir>/nanny/solaris/merc_asl/process.asl` file by adding the following lines:

```
nanny.exe=
nanny=
java=
```

- 2 Modify the `<diagnostics_server_install_dir>/nanny/solaris/dat/nanny/server.nanny` file:

- ▶ Locate the `start_solv4` line:

```
start_solv4= "/opt/Diagnostics/Server/jre/bin/java" -server -Xmx512m...
```

- ▶ Remove the quotes that surround the java path and add the caret (^) after java as shown below:

```
start_solv4= /opt/Diagnostics/Server/jre/bin/java^ -server -Xmx512m...
```

- 3 Modify the `<diagnostics_server_install_dir>/nanny/solaris/dat/br_Inch_server.cfg` file.

Change the value of the `FireWallServiceActive` property to 1.

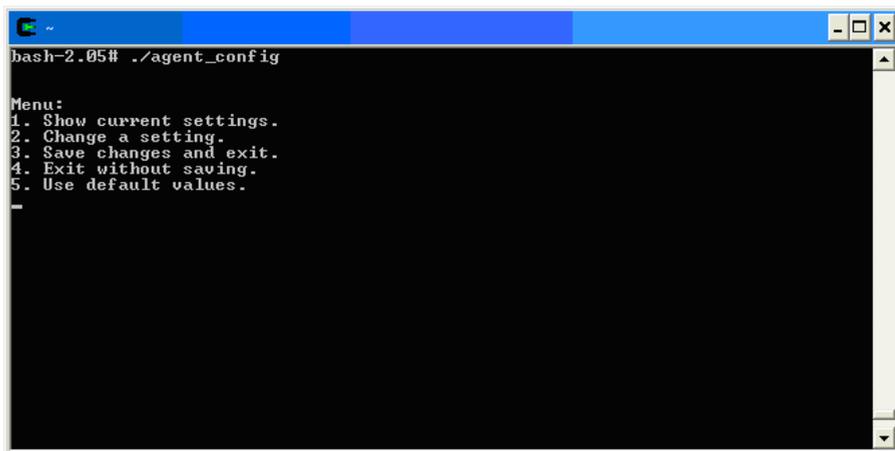
- 4 Modify the `<diagnostics_server_install_dir>/etc/server.properties` file.
 - ▶ Uncomment the `dispatcher.offline.locationprefix` property.
 - ▶ Set the value of the property file to `/tmp/MOfflineFiles` as shown in the following example. The letter for your drive may be different but the path to the directory must be the same.

```
dispatcher.offline.locationprefix= /tmp/MOfflineFiles
```

- 5 Run the following commands to launch the Agent Configuration utility:

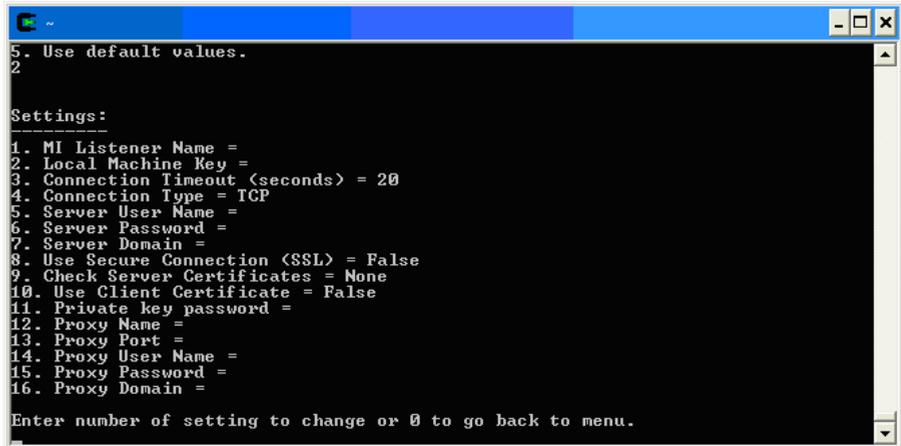
```
export LD_LIBRARY_PATH=.  
export M_LROOT=<diagnostics_server_install_dir>  
cd $M_LROOT/nanny/solaris/bin  
./agent_config
```

- 6 In the Agent Configuration Utility window, press **2**, **Change a Setting**.



7 A list of settings appears.

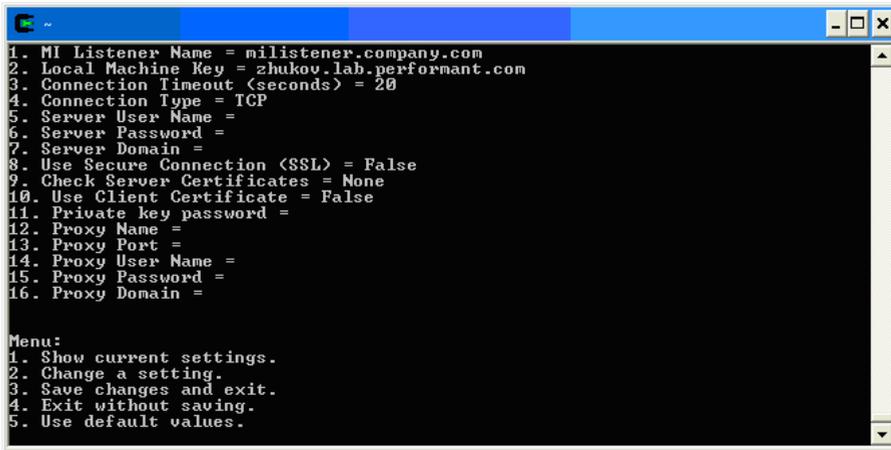
Press **1** to select **MI Listener Name**, and enter the machine name or IP address of the MI Listener host.

A screenshot of a Windows command prompt window. The window title bar is blue and contains the text "~" and standard window control buttons (minimize, maximize, close). The command prompt shows a menu with the following text:

```
5. Use default values.  
2  
  
Settings:  
-----  
1. MI Listener Name =  
2. Local Machine Key =  
3. Connection Timeout (seconds) = 20  
4. Connection Type = TCP  
5. Server User Name =  
6. Server Password =  
7. Server Domain =  
8. Use Secure Connection (SSL) = False  
9. Check Server Certificates = None  
10. Use Client Certificate = False  
11. Private Key password =  
12. Proxy Name =  
13. Proxy Port =  
14. Proxy User Name =  
15. Proxy Password =  
16. Proxy Domain =  
  
Enter number of setting to change or 0 to go back to menu.
```

- 8 Press **2** to select **Change a Setting** and enter the machine name of the host of the Diagnostics Server in Mediator mode as displayed in the Host: field on the System Health Monitor.

Use the System Health Monitor to determine the machine name for the host of the Diagnostics Server in Mediator mode. For more information on the System Health Monitor, see Appendix C, “Using the System Health Monitor.”



```
1. MI Listener Name = milistener.company.com
2. Local Machine Key = zhukov.lab.performant.com
3. Connection Timeout (seconds) = 20
4. Connection Type = ICP
5. Server User Name =
6. Server Password =
7. Server Domain =
8. Use Secure Connection (SSL) = False
9. Check Server Certificates = None
10. Use Client Certificate = False
11. Private key password =
12. Proxy Name =
13. Proxy Port =
14. Proxy User Name =
15. Proxy Password =
16. Proxy Domain =

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

- 9 Press **3**, **Save changes and exit**, to complete the updates.

- 10 Restart the Diagnostics Server in Mediator mode.

For more information about restarting the Diagnostics Server in Mediator mode, see “Starting and Stopping the Diagnostics Server” on page 52.

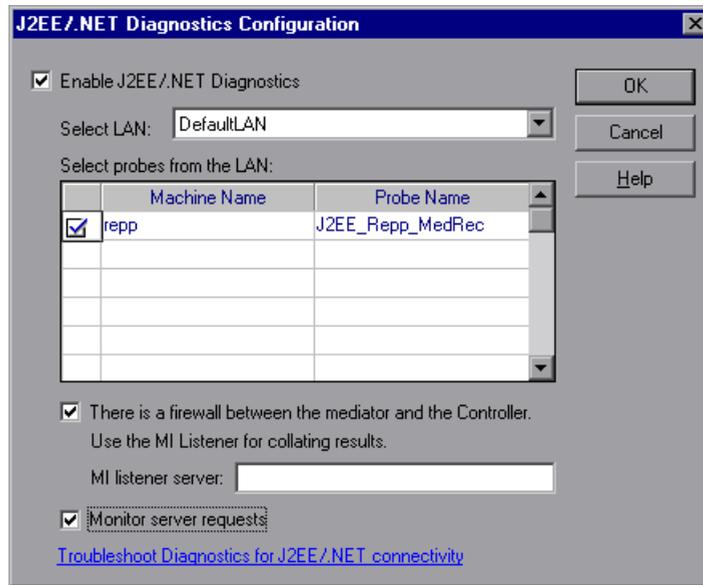
Configuring LoadRunner and Performance Center to Work with Diagnostics Firewalls

After the MI Listener has been installed and your Mediator machines have been configured, you must update the Diagnostics Configuration in LoadRunner / Performance Center so that the application knows to use the MI Listener when it is transferring the offline data from a Mediator that is outside of a firewall. This section describes:

- ▶ Configuring LoadRunner to work with a Diagnostics firewall
- ▶ Configuring Performance Center to work with a Diagnostics firewall

To configure LoadRunner to work with a Diagnostics firewall:

- 1** From the LoadRunner Controller menu bar, select **Diagnostics > Configuration** to open the Diagnostics Distribution dialog box.
- 2** Ensure that the **Enable the following diagnostics** check box is selected.
- 3** In the Online and Offline Diagnostics section, click **Configure**. The J2EE/.NET Diagnostics Configuration dialog box opens.

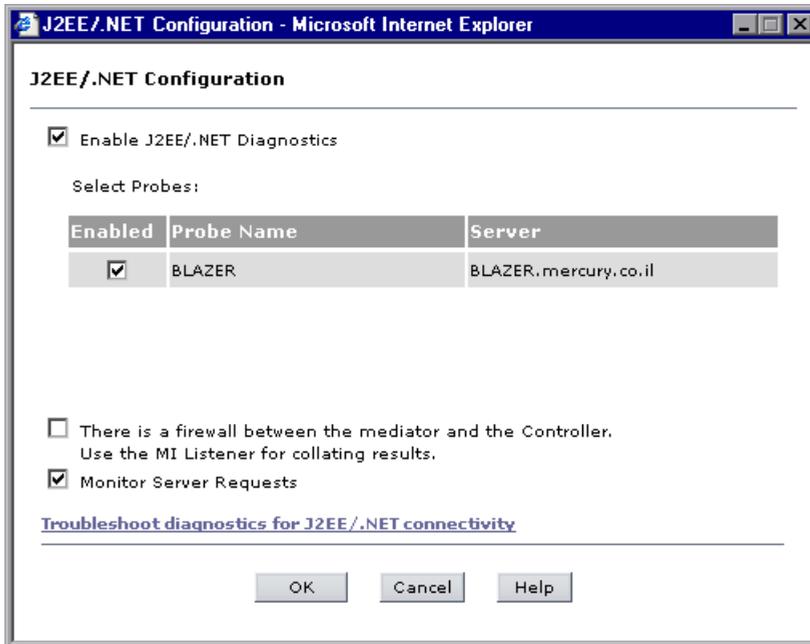


- 4** Select **There is a firewall between the mediator and the Controller**.
- 5** Specify the machine name of the MI Listener host in the **MI listener server** text box.

See Chapter 12, “Setting Up LoadRunner 8.1 to use Mercury Diagnostics,” for more information about configuring LoadRunner to use the Diagnostics components.

To configure Performance Center to work with a Diagnostics firewall:

- 1** In the Performance Center Administration Site, in the **General Settings** page, go to the **Firewall** section and ensure that you have specified the IP address of the MI Listener machine in the **Firewall Diagnostics Communicator** box.
- 2** In the Performance Center User site, choose **Load Tests > Create/Edit** from the left navigation menu to open the Load Test Configuration page.
- 3** Click an existing test that you want to configure in the **Name (# of Runs)** column or click **New Load Test**.
- 4** Click the **Diagnostics** tab, and ensure that the **Enable Diagnostics** check box is selected.
- 5** In the **Offline and Online Diagnostics** section, click **Configure** to open the J2EE/.NET Diagnostics Configuration dialog box.



6 Select **There is a firewall between the mediator and the Controller.**

See Chapter 13, “Setting Up Performance Center 8.1 to Use Mercury Diagnostics,” for more information about configuring Performance Center to use the Diagnostics components.

Part VII

Configuring Diagnostics Metrics Collectors

22

Overview of Diagnostics Metrics Collectors

This chapter describes metrics capture and how to configure the metric collectors.

This chapter describes:	On page:
About Metrics Capture	407
Configuring Metric Collector Entries	408

About Metrics Capture

The metric collectors are installed with the probe. The collectors gather system metrics from the probe host and JMX metrics from the application server. The metric collectors are configurable so that you can control which metrics are collected.

Configuring Metric Collector Entries

The metric collectors for your probe installation are defined in the metrics configuration file. The properties and entries in the metrics configuration file, `<probe_install_dir>/etc/metrics.config`, allow you to control the metric collectors.

Note: When you update the metrics configuration file, the metric collectors automatically restart so the changes will take effect.

Understanding Metric Collector Entries

Metric Collector entries instruct the metric collectors to gather specific metrics. The parameters on the left hand side of the entry control how the probe gathers the metric from the host or the JVM, and the parameters on the right hand side of the entry define how the collected metrics are processed in Diagnostics and displayed in the user interface.

The entries can have one of the following layouts:

```
<collector_name>/<metric_config>=<metric_id>|<metric_units>|<category_id>
```

or

```
<collector_name>/<metric_config>=
RATE<rate_multiplier>(<metric_id>|<metric_units>|<category_id>)
```

where:

- **<collector_name>** indicates the name of the Diagnostics metric collector. The collectors are defined in **metrics.config**.

For system metrics the value of this parameter is `system`. For JMX metrics the value of this parameter is usually defined as the name of the application server type and the version, such as `WebSphere5`.

- ▶ **<metric_config>** identifies the metric that is to be monitored on the host system or on the JVM for the application server. The format of this parameter varies depending on whether you are creating an entry for a system metric or a JMX metric. For information on formatting the **metric_config** property for the system metric collector, see “Capturing Custom System Metrics” on page 416. For information on formatting the **metric_config** property for JMX metrics, see “Capturing JMX Metrics” on page 428.

- ▶ **RATE(...)** indicates that metric values are converted to a rate (units per second) during sampling.

For example, when the **Rate** parameter is used with the metric **total servlet requests since startup**, the value of the collected metric is converted from a *count of servlet requests* to the *number of servlet requests per second*.

When **Rate** is not used, omit the parenthesis as shown in the first example above.

Note: This parameter should only be used for metrics with non-decreasing values.

- ▶ **<rate_multiplier>** is an optional parameter that indicates that the rate is to be adjusted by multiplying the it by the **<rate_multiplier>**.

For example, when the **Rate** parameter and the **rate_multiplier** are used with the metric **total gc time** (in ms), the value of the metric collected is converted from *the total time for gc* to the *percent time spent in gc*.

- ▶ **<metric_id>** indicates the name that represents the metric in the UI. The **metric_id** must be unique in the **metrics.config** file. If the value of the **metric_id** is the same as one of the default metrics, Diagnostics replaces the **metric_id** in the entry with a standard name to be used to reference the metric in the UI. If the value of the **metric_id** is not the same as one of the default metrics, then the **metric_id** is used as the name of the metric in the UI exactly as shown in the entry.

- **<metric_units>** indicates the units of measure in which the metric is reported. This is a required parameter and it must contain one of the following units of measure:
 - microseconds, milliseconds, seconds, minutes, hours, days
 - bytes, kilobytes, megabytes, gigabytes
 - percent, fraction_percent
 - count
 - load
- **<category_id>** groups a set of metrics together under the same heading in the tree in the side bar of the Metrics tab in the J2EE Profiler. This parameter has no impact on the data displayed in the Metrics Inspector in the Diagnostics views.

Note: After you have created the metric collector entry, you must add the escape character “\” before each occurrence of a back-slash '\', space ' ', or colon ':'. This is a requirement for Java properties loaded from a file.

Example

The following example shows how to create the metric collector entry for a system metric. To create an entry for a system metric called CPU on a host platform, you would enter the following:

```
system/CPU = CPU|percent
```

where:

- **system** indicates that the metric is to be collected by the system metric collector
- the first **CPU** indicates that the metric known as **CPU** on the platform is being monitored

- ▶ the second **CPU** is the name that is to be used in the UI to label the metric
- ▶ **percent** indicates the units in which the metric is measured on the host, and reported in the UI

Starting Capture for a Metric

To gather information for an additional metric, add an entry for the metric to the appropriate metric collector in the **metrics.config** file using the template described in “Understanding Metric Collector Entries” on page 408.

Modifying a Metric that is Already Being Captured

You can update both the default and the custom metrics entries in the metric collectors in the **metrics.config** file.

Stopping Capture of a Metric

To stop a metric collector from collecting a metric listed in **metrics.config**, you can either delete the metric entry or make the metric entry a comment line by adding a **#** to the beginning.

23

Configuring System Metrics Capture

This chapter describes the process for capturing system metrics and how to configure the system metric collector to capture them.

This chapter describes:	On page:
About System Metrics	413
System Metrics Captured by Default	414
Configuring System Metric Collector	415
Capturing Custom System Metrics	416
Enabling z/OS System Metrics Capture	423

About System Metrics

The system metric collector is installed with the probe. The collector gathers system level metrics, such as CPU usage and memory usage, from the probe's host. The system metric collector is configurable so that you can control which of the system metrics are collected.

Only one instance of the system metric collector is run on a given host, no matter how many instances of the probe have been started on the host. When an instance of the probe is started, it attempts to connect to the UDP port specified in the metrics properties. If a connection is established, the system metric collector instance is started. If a connection cannot be made, then a system metric collector instance has already been started on the host by another instance of the probe and a new instance cannot be started.

Each probe periodically attempts to connect to the port to make sure that a system metric collector is always running. If the probe that started the systems metric collector is stopped, one of the other instances of the probe will start a new instance of the systems metric collector when it finds that the port is available.

System Metrics Captured by Default

The following are the system metrics that the metric collector collects by default for all supported platforms (excluding z/OS):

- CPU
- MemoryUsage
- VirtualMemoryUsage
- ContextSwitchesPerSec
- DiskBytesPerSec
- DiskIOPerSec
- NetworkBytesPerSec
- NetworkIOPerSec
- PageInsPerSec
- PageOutsPerSec

You can control which of the default system metrics the system metric collector gathers and you can add other platform specific metrics so that the collector gathers the information for them as well. See *Configuring System Metric Collector* for more information. For certain platforms, such as Windows, Solaris, and Linux, you may create custom system metrics which can be gathered by the system metric collector. For details, see “Capturing Custom System Metrics”.

For information on z/OS system metrics see “Enabling z/OS System Metrics Capture” on page 423

Configuring System Metric Collector

You can configure the system metrics capture process to run in your environment, and to collect and report the system metrics that are of interest to you, by modifying the properties in the metrics configuration file, `<probe_install_dir>/etc/metrics.config`.

Note: If you update the metrics configuration file, the systems metric collector automatically restarts so that your changes can take effect.

Modifying the Default Port

The default port for the metric collectors is 35000. This value can be modified using the `system.udp.port` property if the configuration for your probe host requires that another port be used.

To modify the default port:

- 1** Locate the `system.udp.port` property in `metrics.config`.
- 2** Change the value of the `system.udp.port` property to the number of the port that you want to be used by the system metric collector. The default port is 35000.

Note: The port assigned to the system metric collector is not related to the port for the probe's Web server.

For information on configuring the entries in the metric collectors see “Configuring Metric Collector Entries” on page 408

Disabling System Metrics Collection

To disable the collection of system metrics so that they will not be collected or displayed in the UI, set the value of the `system.udp.port` property to -1.

Capturing Custom System Metrics

The Windows, Solaris, and Linux platforms allow you to create custom system metrics. The custom metrics can be monitored by the system metric collector.

The following sections provide instructions for creating the metrics and updating the entries in the system metric collector so that the custom metrics can be monitored.

Capturing Custom System Metrics on Windows Hosts

Using the features of Windows System Monitor, you can add counters to represent the performance of specific aspects of a system or service. The counters are tracked and reported in the Windows System Monitor, and can be monitored by the Diagnostics system metric collector.

To add counters using the Windows System Monitor:

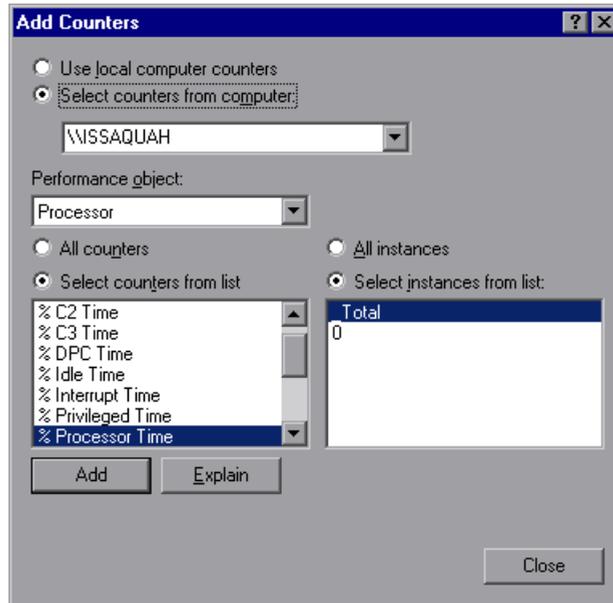
- 1** Start the Windows Performance Monitor:
 - a** Select **Start > Run** from the Start menu.
 - b** In the **Open** box on the **Run** dialog box type 'perfmon'.

The **Performance** dialog box opens showing the **System Monitor** graph with a table of the current counters beneath the graph.

2 Display the Add Counters dialog box:

- a Right-click the **System Monitor** graph and select **Add Counters...** from the pop-up menu.

Windows displays the **Add Counters** dialog box:



3 Ensure that the host computer is selected from **Select counters from computer** list.

4 In the **Performance object** list, select the object that the counter belongs to.

5 Choose **Select counters from list**, and select a counter from the list of counters that follows.

6 Choose **Select instances from list**, and select an instance from the list of instances that follows.

7 Click **Add**.

Once a counter has been added to the Systems Monitor, the system metric collector can be configured to gather the metrics for the counter. The following instructions will guide you through the steps to create an entry for the **metrics.config** based upon the following template:

```
<collector_name>/<metric_config>= <metric_id>|<metric_units>
```

This template is described in “Understanding Metric Collector Entries” on page 408.

To collect metrics for a Windows System Monitor Counter:

- 1** Open `<probe_install_dir>/etc/metrics.config`.
- 2** Create the `<metric_config>` part of the entry using the following template, type the entry for the counter:

```
\<performance_object>(<instance>)\<counter>
```

In the example shown in the preceding screen image:

- ▶ the selected Performance Object is %Processor
- ▶ the selected Instance is _Total
- ▶ the selected Counter is Processor Time

The `<metric_config>` portion of the entry that would be created for this example would be:

```
\Processor(_Total)\% Processor Time
```

- 3** Fill in the rest of the system metric entry template as shown in the following example:

```
system\Processor(_Total)\% Processor Time = ProcessorTime|percent
```

- 4 Format the initial entry by prepending a back-slash '\' before each occurrence of back-slash '\', space ' ', or colon ':' in the initial entry.

Following this step, the initial entry in the previous step becomes:

```
system\\Processor(_Total)\\% Processor Time = ProcessorTime|percent
```

This is the correctly formatted entry for **metrics.config** to enable the system metric collector to gather the metrics for a Windows System Monitor counter.

```
system\\\\"RemoteMachine"\\Processor(_TOTAL)\\% Processor Time=
Processor Time(Remote Machine)|percent
```

Note: Assuming **perfmon** is setup properly on a remote machine, you can use it to get metrics from remote machines by adding **\\MachineName** before the Performance object name as shown in the following example:

```
system\\\\"RemoteMachine"\\Processor(_TOTAL)\\% Processor
Time=Processor Time(Remote Machine)|percent
```

Capturing Custom System Metrics on Solaris Hosts

The Solaris system metrics that can be monitored by the system metric collector are found using the **kstat** command. Only a subset of the metrics found using the **kstat** command can be monitored by the system metric collector.

To collect metrics for a Solaris system metric:

- 1 Execute the **kstat** command and identify the metric that you want to monitor.

A Solaris system metric has the following format:

```
module:instance:name:statistic
```

For example:

```
vmem:35:ptms_minor:free
```

- 2 To cause the metric collector to gather the metrics for an additional system metric, add an entry for the metric to the system metric collector in the **metrics.config** file using the following template:

```
<collector_name>/<metric_config>= <metric_id>|<metric_units>
```

This template is described in “Understanding Metric Collector Entries” on page 408.

Using this template, the example from the previous step would initially appear as follows:

```
system/vmem:35:ptms_minor:free = Virtual Memory (35) Free | count
```

- 3 Format the initial entry by prepending a back-slash '\' before every back-slash '\', space ' ', or colon ':':

Following this step the initial entry in the previous step becomes:

```
system/vmem\:35\:ptms_minor\:free = Virtual\ Memory\ (35)\ Free | count
```

This is the correctly formatted entry for **metrics.config** to enable the system metric collector to gather the metrics for a Solaris systems metric.

Capturing Custom System Metrics on Linux Hosts

The Linux system metrics that can be monitored by the system metric collector are found in the **/proc** file system. To configure the system metric collector to gather custom Linux metrics you must scan the **/proc** file system to locate the desired metric, and then create the system metric collector entry for the metric in **metrics.config** based upon the where the metric information is located.

To collect metrics for a Linux system metric:

- 1 Scan the **/proc** file system to locate the metric that you would like the Diagnostics system metric collector to monitor.

To create the system metrics configuration entry in **metrics.config** for the Linux metric, you must explicitly specify where the value for the system metric is located. The location is specified using the following values:

- ▶ **File name.** The name of the file where the metric information is located, including the path from the **/proc** directory.
- ▶ **Line offset.** A count of the number of lines in the file to the line where the system metric is located. The first line is counted as line 0.
- ▶ **Word offset.** A count of the number of words that the metric value is offset into the line in the file. The first word in the line is counted as line 0. The value at the specified offset must be an unsigned integer.

For example, if you wanted the system metric collector to monitor the SwapFree system metric so that you can see it displayed in the Diagnostics views, you would scan the **/proc** directory to locate the metric, and you would discover that the metric is located in the **meminfo** file. The layout of this file is as follows:

```
MemTotal: 515548 kB
MemFree: 1552 kB
Buffers: 41616 kB
Cached: 152084 kB
SwapCached: 46064 kB
Active: 402720 kB
Inactive: 75328 kB
HighTotal: 0 kB
HighFree: 0 kB
LowTotal: 515548 kB
LowFree: 1552 kB
SwapTotal: 1048568 kB
SwapFree: 779192 kB
Dirty: 4544 kB
Writeback: 0 kB
Mapped: 300056 kB
Slab: 28764 kB
Committed_AS: 801364 kB
PageTables: 3184 kB
VmallocTotal: 499704 kB
VmallocUsed: 2184 kB
VmallocChunk: 497324 kB
HugePages_Total: 0
HugePages_Free: 0
Hugepagesize: 4096 kB
```

The location of the SwapFree metric in this file would lead to the following values:

- **File name:** meminfo
- **Line offset:** 12
- **Word offset:** 1

- 2** To gather the metrics for an additional system metric, add an entry for the metric to the system metric collector in the **metrics.config** file using the following template:

```
<collector_name>/<line>:<word>:<file>= <metric_id>|<metric_units>
```

This template is a version of the template described in “Understanding Metric Collector Entries” on page 408. The `<metric_config>` property has been replaced with the properties `<line>:<word>:<file>`.

Using this template, the example from the previous step would initially appear as follows:

```
system/12:1:meminfo = Swap Free | kilobytes
```

- 3** Format the initial entry by prepending a back-slash `'\'` before every back-slash `'\'`, space `' '`, or colon `'.'`:

Following this step the initial entry in the previous step becomes:

```
system/12\:1\:meminfo = Swap\ Free | kilobytes
```

This is the correctly formatted entry for **metrics.config** to enable the system metric collector to gather the metrics for a Solaris systems metric.

Enabling z/OS System Metrics Capture

The following system metrics can be collected for the z/OS platform:

- CPU
- DiskIOPerSec
- DiskBytesPerSec

System metrics are not captured by default, as this requires some system configuration changes.. You must perform the following configuration steps in order to enable capture of z/OS system metrics.

To enable z/OS system metrics capture:

- 1** Change the permissions for the directory `<probe_install_dir>/bin/` to recursively allow execution. This can be done using a command like the following example:

```
chmod -R 770...
```

- 2** Change the permissions for the directory `<probe_install_dir>/bin/390-zos/systemmetrics` to allow execution. This can be done using a command as illustrated in the following example:

```
chmod -R 0+x ...
```

- 3** Start the RMF Monitor III and ensure that SMF record 70-79 is collecting.
- 4** Start the RMF Data Buffer on one or more systems in the sysplex.
- 5** Check the list of system names passed to the ERBDSQRY service.
- 6** Ensure that the system is collecting SMF record 92 with subtype 5.

24

Configuring JMX Metric Capture

This chapter describes the process for capturing JMX metrics and how to configure metric collectors to capture them.

This chapter describes:	On page:
About JMX Metrics	425
Configuring WebSphere for JMX Metric Collection	426
Configuring JMX Metric Collectors	428
Capturing Custom JMX Metrics	428

About JMX Metrics

The J2EE Probe comes with pre-defined JMX metric collectors that access the JMX metrics from following application servers:

- ▶ IBM WebSphere 5.x, 6.x
- ▶ BEA WebLogic 6.x, 7.x, 8.x, and 9.x
- ▶ SAP NetWeaver 6.4

The J2EE Probe can also collect JMX data from any J2EE server that supports the JMX standard. To collect JMX metrics from a different server, refer to Chapter 24, “Configuring JMX Metric Capture.”

The J2EE Probe runs the JMX metric collectors periodically to collect the metrics from the application server. The collected metrics are displayed on the user interfaces in both Mercury Diagnostics and J2EE Diagnostics Profiler.

Configuring WebSphere for JMX Metric Collection

You may need to configure the Performance Monitoring Service on the WebSphere server in order to start receiving JMX metrics. Instructions are provided below for both WebSphere 5.x and WebSphere 6.x.

The following are the instructions for configuring WebSphere 5.x for JMX metrics collection.

To configure WebSphere 5.x server for JMX metrics collection:

- 1** Open the WebSphere Administrative Console.
- 2** In the Console navigation tree, select **Servers > Application Servers**. The console displays a table of the application servers.
- 3** Click name of the application server that you want to configure from the Application Servers Table. The console displays the **Runtime** and the **Configuration** tabs for the selected application server.
- 4** Click the **Configuration** tab.
- 5** In the **Configuration** tab:
 - Click **Performance Monitoring Service**.
 - Select the **Startup** check box under General Properties.
 - Set the **Initial Specification Level** to **standard**.
 - Click **Save**.
- 6** Click **Apply** or **OK**.
- 7** If Java 2 Security is enabled on the application server, open the server policy file. (<WebSphere 5.x Installation Directory>/profiles/server.policy) and add the following security permission to the file, to allow JMX collection:

```
grant codeBase "file:<probe_install_dir>/lib/probe-jmx.jar"  
{ permission java.security.AllPermission; }
```
- 8** Restart the application server.

The following are the instructions for configuring WebSphere 6.x for JMX metrics collection.

To configure WebSphere 6.x server for JMX metrics collection:

- 1** Open the WebSphere Administrative Console.
- 2** In the Console navigation tree, select **Servers > Application Servers**. The console displays a table of the application servers.
- 3** Click name of the application server that you want to configure from the Application Servers Table. The console displays the **Runtime** and the **Configuration** tabs for the selected application server.
- 4** Click the **Configuration** tab.
- 5** In the **Configuration** tab:
 - ▶ Under the Performance heading, click **Performance Monitoring Infrastructure (PMI)**.
 - ▶ Under the General Properties heading, select the **Enable Performance Monitoring Infrastructure (PMI)** check box.
 - ▶ Under the Currently monitored statistic set heading select **Extended**.
- 6** Click **Apply** or **OK**.
- 7** If Java 2 Security is enabled on the application server, open the server policy file. (<WebSphere 6.x Installation Directory>/work/tools/ibm-6.0/websphere/appserver/profiles/default/properties/server.policy) and add the following security permission to the file, to allow JMX collection:
**grant codeBase "file:/<probe_install_dir>/lib/probe-jmx.jar"
{ permission java.security.AllPermission; }**
- 8** Restart the application server.

Configuring JMX Metric Collectors

The JMX metric collectors are configurable so that you can control which of the JMX metrics are collected. The JMX metric collectors are defined in the `<probe_install_dir>/etc/metrics.config` file. Typically a separate collector is defined for each major version of each application server.

For information on configuring the metric collectors see “Configuring Metric Collector Entries” on page 408.

Capturing Custom JMX Metrics

The J2EE Probe is installed with a number of predefined JMX metric collectors for the application servers listed in “About JMX Metrics” on page 425. You may modify or remove any of the JMX metric entries from the predefined collectors by following the instructions in “Configuring Metric Collector Entries” on page 408. You may also create entries in the existing metric collectors and even create new collectors if there are additional JMX metrics that you would like Diagnostics to monitor.

The following sections provide instructions creating new entries in the JMX metric collectors so that additional JMX metrics can be monitored.

Capturing JMX Metrics

The metric collectors installed with the J2EE Probe include entries for many of the JMX metrics that are available for each application server. However, there may be other JMX metrics that you would like to monitor or new metrics may be exposed by the application server vendor. The following instructions will guide you through the process of creating the JMX metric entries based upon the following template:

```
<collector_name>/<metric_config>= <metric_id>|<metric_units>
```

This template is described in “Understanding Metric Collector Entries” on page 408.

To capture JMX metrics:

- 1** Open `<probe_install_dir>/etc/metrics.config`, and locate the JMX metric collector that is appropriate for the application that is being monitored by the J2EE Probe.
- 2** The `<collector_name>` parameter is the same as the rest of the entries in the collector. If you were creating an entry for WebLogic, the value of this parameter would be WebLogic.
- 3** Create the `<metric_config>` parameter using the following template:

```
<MBean object name pattern>.<attribute name>
```

For JMX metrics the `<metric_config>` parameter is a pattern that the collector uses to find a matching MBean. The pattern consists of two components, separated by the '.' character:

- ▶ **<MBean object name pattern>** is the string representation of the object name of an MBean
- ▶ **<attribute_name>** is the name of the attribute that represents the metric.

For example, for a WebLogic application server, the `<metric_config>` parameter for the throughput of all **Execute Queues** is configured as:

```
*:Type=ExecuteQueueRuntime,*.ServicedRequestTotalCount
```

For an explanation of metric patterns see “Understanding Metric Patterns” on page 430.

- 4** Fill in the rest of the JMX metric entry template as shown in the following example:

```
WebLogic/*:Type=ExecuteQueueRuntime,*.ServicedRequestTotalCount =  
RATE(Execute Queues Requests / sec|count|Execute Queues)
```

- 5** Format the initial entry by prepending a back-slash '\' before every back-slash '\', space ' ', equals (=), or colon ':':

Following this step the initial entry in the previous step becomes:

```
WebLogic/*\:Type\=ExecuteQueueRuntime,*.ServicedRequestTotalCount  
= RATE(Execute Queues Requests / sec|count|Execute Queues)
```

This is the correctly formatted entry for a JMX metric collector to enable the collector to gather a WebLogic JMX metrics.

Understanding Metric Patterns

For JMX metrics the <metric_config> parameter is a pattern that the collector uses to find a matching MBean. For example:

```
*:Type=ExecuteQueueRuntime,*.ServicedRequestTotalCount
```

In the example above, the object name is ***:Type=ExecuteQueueRuntime,***, which could actually resolve to many MBeans whose names have the **Type** component equal to **ExecuteQueueRuntime**. **ServicedRequestTotalCount** is an attribute name for which metric values will be collected by the JMX metric collector.

Note: Current implementation of the JMX collector only supports attributes that are numeric in type (i.e., long, integer, etc.).

The JMX metric collector first uses MBeanServer's query mechanism to find the matching MBeans for each object name provided in the configuration. For JMX metrics the object names are a pattern that the collector uses to find a matching MBean. For more details around the object names, see <http://java.sun.com/j2ee/1.4/docs/api/javax/management/ObjectName.html>.

Since MBean object names are patterns that can resolve into multiple MBeans, the JMX collector will validate all of the attribute names in the entry against all MBeans that match the pattern, and will aggregate the attribute values over the set of those matching MBeans. Of course, it is not always the case that the object name resolves into multiple MBeans. For example, the following object name resolves to a single MBean (on a WebLogic application server):

```
*\:Name\=weblogic.kernel.Default,Type\=ExecuteQueueRuntime,  
*.ServicedRequestTotalCount
```


Part VIII

Appendixes

A

Diagnostics Server Administration Page

This appendix describes how to access the Diagnostics Server Administration page, where you can configure Diagnostics properties. It also describes how to view the Configuration pages, and modify the properties within these pages.

This appendix describes:	On page:
Accessing the Diagnostics Server Administration Page	435
Configuring Diagnostics Using the Diagnostics Server Configuration Pages	437

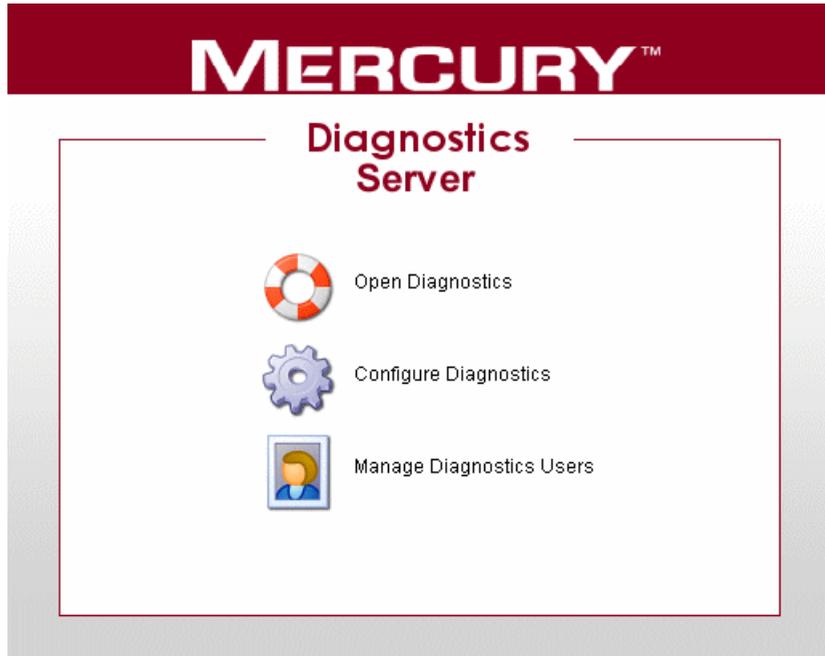
Accessing the Diagnostics Server Administration Page

You can set the user privileges, configure Diagnostics settings, and open Diagnostics directly from the administration page of the Diagnostics Server.

To use the Diagnostics Server administration page:

- 1 Open the administration page by navigating to http://<diagnostics_server_host>:2006 in your browser, or by selecting **Start > Programs > Mercury Diagnostics Server > Administration**. The port number in the URL, **2006**, is the default port for the Diagnostics Server. If you configured the Diagnostics Server to use an alternative port, use that port number in the URL.

The Diagnostics Server administration page opens in your browser.



2 Click the option for the activity that you want to perform.

- ▶ **Open Diagnostics.** Opens the Diagnostics Web pages where you can view the performance metrics collected by the probes that are reporting to the Diagnostics Server. The performance metrics are displayed in the standard Diagnostics views.

For more information about opening Diagnostics directly from the Diagnostics Server in Commander mode, see the *Mercury Diagnostics User's Guide*.

- ▶ **Configure Diagnostics.** Opens the Components page, which has a link to the Diagnostics Server Configuration page.

For more information about configuring the Diagnostics properties, see "Configuring Diagnostics Using the Diagnostics Server Configuration Pages" on page 437.

- **Manage Diagnostics Users.** Opens the User Administration page where you can add and maintain security information, and user privileges for specific users.

For more information about security and user privileges, see “Product Security & User Privileges” on page 443.

Configuring Diagnostics Using the Diagnostics Server Configuration Pages

In the Diagnostics Server Configuration pages, you set the property values that control how the Diagnostics Server communicates with the other Diagnostics components, and how it processes the data that it receives from the probes.

Note: To ensure that you are entering valid property values, it is recommended that you use the configuration pages to modify the Diagnostics Server properties, rather than editing the property files directly.

Accessing Diagnostics Server Configuration Pages

You access the Diagnostics Server Configuration pages from the Components page.

To access the Diagnostics Server configuration pages:

- 1** Open the administration page by navigating to http://<diagnostics_server_host>:2006 in your browser, or by selecting **Start > Programs > Mercury Diagnostics Server > Administration**. The port number in the URL, **2006**, is the default port for the Diagnostics Server. If you configured the Diagnostics Server to use an alternative port, use that port number in the URL.

The Diagnostics Server administration page opens in your browser.

- 2** Click **Configure Diagnostics**.

- 3 If you have not already signed into the Diagnostics Server, you are prompted for a user name and password. This must be a valid user name, and must have both **View** and **Change** privileges. For information about valid user names and privileges, see Appendix B, “Product Security & User Privileges.”

Note: Diagnostics continues to prompt for a user name and password until valid credentials are entered.

If you click **Cancel**, the following error message is displayed in your browser: **Access denied. You must specify a valid user name and password.**

If you entered a valid user name and password, but do not have the proper privileges, the following error message is displayed in your browser: **Access denied. You do not have the required permission to view this screen.**

To log on as a different user to the one you are currently logged on as, you need to close your browser and re-open it.

After you have logged on, the Diagnostics Server Components page opens:

Component Name	Component Description
registrar	Central list of all Diagnostics component deployments
security	User Management
logging	Configure log files and logging details.
monitor	Current Metrics
configuration	Configuration
license	License Management

Mercury Diagnostics Server "server-square", version 4.1.48.646

4 Click **configuration** to access the Diagnostics Server Configuration page.

Configuration	
Name	Description
Customer Information	Properties that are used to configure the customer information for the Server.
Alert Properties	Properties that are used to configure the alert settings for the Server. All changes take effect dynamically, without server restart.
Component Communications	Properties that are used to configure how this Diagnostics Component communicates with the other diagnostic components.
Memory Diagnostics	Properties that configure the way that the Server will handle memory diagnostics
Online Cache	Properties that are used to configure the Server's Online Cache.
Logging	Configure log files and logging details.
Show Advanced Options	
Mercury Diagnostics Server	

Click the link to the page whose properties you want to update. You can configure:

- Customer information
- Alert properties
- Component Communication properties
- Memory Diagnostics properties
- Online cache properties
- Logging level details

Viewing Advanced Diagnostics Server Configuration Options

The configuration options displayed on the Configuration page are the commonly configured options. The more advanced configuration options are hidden by default.

Important! Do not manipulate the advanced options without the guidance of your Mercury Customer Support representative.

To display the advanced options:

At the bottom of the page, click **Show Advanced Options**.

The list of options on the page is updated to include the advanced configuration options, and the link changes to **Hide Advanced Options**.

To hide the advanced options:

At the bottom of the page, click **Hide Advanced Options**.

The list of options on the page is updated so that the advanced configuration options are no longer visible, and the link changes to **Show Advanced Options**.

Modifying Diagnostics Server Properties

You modify the properties of the Diagnostics Server from the Configuration page.

To modify the Diagnostics Server properties:

- 1 Access the Diagnostics Server Configuration page, as described in “Accessing Diagnostics Server Configuration Pages” on page 437.
- 2 Click the link to the properties (for example, **Component Communications**) that you want to update.
- 3 Review the properties that are displayed and revise any values that should be updated.

MERCURY™			
Component Communications			
Name	Value	Description	Default Value
Commander URL	<input type="text" value="http://localhost:2006"/>	The web address for the Diagnostics Commander.	http://localhost:2006
Registered Host Name	<input type="text"/>	The hostname to register with in the Diagnostics Commander Registrar. This hostname will be used by the Commanding Server to reconnect back to the Distributed Server when used in a Load Runner environment. If no hostname is specified, the Server will register with the default system hostname.	
Diagnostics Commander Proxy Host	<input type="text"/>	The hostname or IP address of the proxy that should be used to communicate with the Diagnostics Commander. Leave blank for no proxy.	
Server/Commander Proxy Port	<input type="text"/>	The port of the proxy that should be used to communicate with the Diagnostics Commander.	80
Diagnostics Commander Proxy Protocol	<input type="text"/>	The protocol (either HTTP or HTTPS) to connect to the proxy for the Diagnostics Commander.	http
Probe Data Port	<input type="text" value="2612"/>	The Server port where it receives event data from the probe.	2612
Server Webserver Listen Address	<input type="text" value="0.0.0.0"/>	The local address the Local Server listens on for HTTP requests from the Commanding Server. The default is set to listen on all local addresses.	0.0.0.0
Server Webserver Port	<input type="text" value="2006"/>	The port where the Local Server listens on for HTTP requests from the Commanding Server.	2006

- 4 When you are satisfied with your changes, click **Submit** to save them. Otherwise click **Reset All** to cancel the changes you made.
- 5 A message appears at the top of the page to indicate that your changes were saved.

Note:

- ▶ For most properties that you update, a message is displayed reminding you to restart the Diagnostics Server. The property changes will not take effect until you restart the Diagnostics Server.

If you want make other changes to the Diagnostics Server properties, you should finish making all of your changes before restarting the Diagnostics Server.

Restarting the server will result in a small loss of data (up to 6 minutes). You should therefore schedule restarts at a time that is convenient.

- ▶ Modifying the logging level details does not require restarting the Diagnostics Server; however, it may take up to a minute for your changes to be applied.
-

B

Product Security & User Privileges

This appendix introduces Mercury Diagnostics product security and shows how to create and maintain user security permissions.

This appendix describes:	On page:
About Product Security and User Privileges	444
Understanding User Privileges	445
Accessing Diagnostics Using Default User Names	446
Assigning User Privileges	447
Tracking User Administration Activity	451
Changing Your Password	452

About Product Security and User Privileges

Product security is controlled by creating user names and granting the users privileges so that users are able to perform the functions within the application that they are responsible for.

Before you can access the Diagnostics data, you must sign in with a valid user name and password. The user name that you use to sign in must have permission to view the data on the views that you attempt to access before Diagnostics will display the information.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

Understanding User Privileges

The following privilege levels are defined for Diagnostics:

Privilege	Description
View	The user can view Diagnostics data from the UI. Access to the Configure Diagnostics menu is not allowed.
Execute	The user can make changes to the settings on the UI, such as to thresholds. Access to the Configure Diagnostics menu is not allowed.
Change	The user can access the Configure Diagnostics menu to alter component configuration, view system health, and maintain user information.
System	The user can participate in cross-server communication.
RhttpOut	The user can link to rhttp connections.

Each privilege level stands alone. There is no inheritance of privileges from one level to the next. You must grant a user all of the privilege levels that are necessary to perform the functions that they need to perform.

For example, a user must be granted both **View** and **Execute** privileges in order to be able to make changes to thresholds. A user name that has been granted only **Execute** privileges would not be useful, as it would not allow the user to see the UI that they have permission to make changes to.

Important! In Diagnostics, there is no database to store user names and privileges. If a user needs privileges on several Diagnostics components—that is, Diagnostics Servers and probes—you must create the same user name and privileges on each component.

Accessing Diagnostics Using Default User Names

The following default user names have been defined for Diagnostics:

Default User Names	Privileges	Description
user	View	Can only view the data from the UI.
superuser	View, Execute	Can view data and change thresholds from the UI.
admin	View, Change, Execute	Can view data and change thresholds from the UI. Can configure the components, and maintain user information. This is the default user.
mercury	Diagnostics Server in Commander mode: System Diagnostics Server in Mediator mode: System, View, Change	Grants other servers and probes privileges for registration, keep-alive, and data transfer. This is the default user for cross-server communication.

You can use these default user names to access Diagnostics functionality.

The passwords for the default user names are the same as the user names. For example, for the user name **admin**, the password is **admin**.

You can modify the password or privileges for the default user names to suit your needs. You can also define new user names in order to control user access to Diagnostics.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

Assigning User Privileges

Before you can view any Diagnostics data, or make any changes to the Diagnostics configuration or user privileges, you must log on to the Diagnostics Server using a user name that has valid security access with the appropriate privileges.

After logging on to the Diagnostics Server in a particular browser session, the user name remains in effect until the browser session ends. When you are finished using Diagnostics, be sure to close your browser to prevent others from accessing Diagnostics using your privileges.

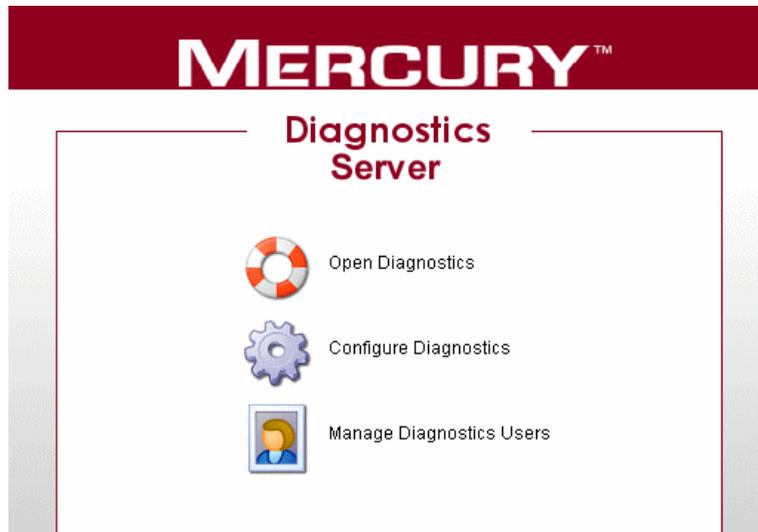
Accessing User Management

You assign user privileges from the User Administration page.

To access the User Administration page from the Server Administration page:

- 1 In your browser, navigate to http://<diagnostics_server_host>:2006, or select **Start > Programs > Mercury Diagnostics Server > Administration**. The port number in the URL, **2006**, is the default port for the Diagnostics Server. If you configured the Diagnostics Server to use an alternative port, use that port number in the URL.

The Diagnostics Server Administration page opens.



- 2 Click **Manage Diagnostics Users**.

If you have not already signed into the Diagnostics Server, you are prompted for a user name and password. This must be a valid user name.

You must have at least **View** privileges in order to view your privileges, and modify your password. To add or delete users, or update user privileges, you must have both **View** and **Change** privileges.

Note: Diagnostics continues to prompt for a user name and password until valid details are entered.

If you click **Cancel**, the following error message is displayed in your browser: **Access denied. You must specify a valid username and password.**

If you entered a valid user name and password, but do not have the proper privileges, the following error message is displayed in your browser: **Access denied. You do not have the required permission to view this screen.**

The User Administration page opens, where you can add and manage the security information for specific user names. For details, see “Managing User Details,” below.

To access the User Administration page from the Diagnostics Views:

Click **Change Password** at the top the Diagnostics views.

The Diagnostics Server Administration page opens

Managing User Details

You add new user names, update the security information for an existing user name, and delete users from the User Administration page.

All users with **View** privileges can maintain their own password. Users with both **View** and **Change** privileges can add new users and maintain the security information for any user.

Adding and Deleting Users

Users with both **View** and **Change** privileges can add new users and grant it user privileges, and delete users.

To add a user:

- 1 On the User Administration page, click **Create User**.
- 2 In the **New User Name** box, type a user name for the new user and click **OK**. The new user appears in the list of user names.
- 3 Under **Privileges**, select the appropriate privileges for the new user, as described in “Understanding User Privileges” on page 445.
- 4 Under **Change Password**, in the **Password** box, type a password for the new user, and confirm it by retyping it in the **Confirm Password** box.
- 5 In the **Password for <current user>** box, type the password of the user currently logged on.
- 6 Click **Save Changes**.

To delete a user:



- 1 On the User Administration page, in the **Password for <current user>** box, type the password of the user currently logged on.
- 2 Click the **Delete user** button corresponding to the user you want to delete.
- 3 A message box opens asking if you want to delete the selected user. Click **OK** to delete the user.

Updating Privileges of Existing Users

Users with both **View** and **Change** privileges can update the password and privileges of existing users, including the default users. A user that has only **View** privileges can modify his own password, but not his privileges.

To update an existing user's details:

- 1 On the User Administration page, under **Privileges**, select the relevant privileges for a particular user.
- 2 To modify the user's password, type the new password in the **Password** and **Confirm Password** boxes.

- 3 In the **Password for <current user>** box, type the password of the user currently logged on, and click **Save Changes**.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

Tracking User Administration Activity

Each time a user enters the Diagnostics Server User Administration page, all activity that takes place is logged in the following log file:

`<diagnostics_server_install_dir>\log\server-useradmin.log.`

The data logged in the file includes the date and time that the user entered the page, and the date and time of each action performed.

To view the log file:

- 1 Open the Diagnostics Server administration page in one of the following ways:
 - By selecting **Start > Programs > Mercury Diagnostics Server > Administration**.
 - By navigating to http://<diagnostics_server_host>:2006 in your browser. The port number in the URL, **2006**, is the default port for the Diagnostics Server. If you configured the Diagnostics Server to use an alternative port, use that port number in the URL.

The Diagnostics Server administration page opens.

- 2 Click **Configure Diagnostics**.

- 3 If you have not already signed into the Diagnostics Server, you are prompted for a user name and password. This must be a valid user name, and must have both **View** and **Change** privileges. For information about valid user names and privileges, see “Understanding User Privileges” on page 445.

Note: Diagnostics continues to prompt for a user name and password until valid credentials are entered.

If you click **Cancel**, the following error message is displayed in your browser: **Access denied. You must specify a valid user name and password.**

If you entered a valid user name and password, but do not have the proper privileges, the following error message is displayed in your browser: **Access denied. You do not have the required permission to view this screen.**

The Diagnostics Server Components page opens.

- 4 Click **logging**. The logging page opens.
- 5 Click **View Log Files**. A list of log files appears.
- 6 Click the `<diagnostics_server_install_dir>\log\server-useradmin.log` link.
The log file is displayed at the bottom the page.

Changing Your Password

Users that have been granted at least one privilege are able to change their own password.

To change your password:

- 1 Access the User Administration page as described in “Accessing User Management” on page 448.
- 2 To modify your password, type the new password in the **Password** and **Confirm Password** boxes under the **Change Password** column header.
- 3 In the **Password for <current user>** box, type your current password
- 4 Click **Save Changes**.

- 5 If the password is successfully updated, Diagnostics displays a message, **Changes have been successfully processed.**
- 6 Click **Main Menu** to return to the Diagnostics Server administration page or x out of the browser.

Note: Diagnostics provides security on both the Diagnostics Server and the Probes. To be able to drill down to a Profiler from Diagnostics, you must establish security privileges on both the Probe and the Diagnostics Server. You must maintain user names and passwords that are the same for the Probe and the Diagnostics Server independently for each component. Changing the password for a user name on one component does not change the password for the user name on other components.

C

Using the System Health Monitor

This appendix describes how to use the System Health Monitor to review the configuration of the Mercury Diagnostics components and verify that they are working properly.

This appendix describes:	On page:
Introducing the System Health Monitor	456
Accessing the System Health Monitor	456
Understanding the System Health Monitor	459
Viewing Component Status and Host Configuration Tooltips	464
Viewing Detailed Component Information	466
Viewing Troubleshooting Tips	470
Viewing Log Information for the Whole System	470
Customizing the System Health Monitor Display	471
Filtering the System Health Monitor Component Map by Customer	473
Controlling the Refresh Rate of the System Health Monitor	474
Creating and Using System Health Monitor Snapshots	475

Introducing the System Health Monitor

The System Health Monitor provides you with a map of all of the components of your Mercury Diagnostics deployment, and gives you real-time information and health status for each component. At a glance, you can determine which components are experiencing problems, the load on each component, and the amount of data flowing between components.

By drilling down to the System Health Monitor component map, you can reveal the details about the configuration, performance metrics, and processing logs for the components. You can also find troubleshooting tips for handling many of the issues that are revealed in the System Health Monitor component map. In many cases, the System Health Monitor will be your first and only stop when you need to know information about the components in your Diagnostics deployment and the machines that host them.

You can also capture a snapshot of the System Health Monitor information in an .xml file so that you can share it with others who may be able to help you diagnose the issues that you see on the map.

Accessing the System Health Monitor

You can access the System Health Monitor directly from your Web browser, from Performance Center, from the LoadRunner Controller, or from Mercury Business Availability Center.

To access the System Health Monitor from your Web browser:

- 1** Enter the following URL into your browser:

```
http://<Diagnostics Server in Commander mode host>:<Diagnostics Server in  
Commander mode port>/registrar/health
```

The default port for the Diagnostics Server in Commander mode is 2006. If you have configured the Diagnostics Server to use an alternate port, use that port number in the URL.

- 2** Enter your user name and password. For more information about user names and passwords, see Appendix B, “Product Security & User Privileges.”

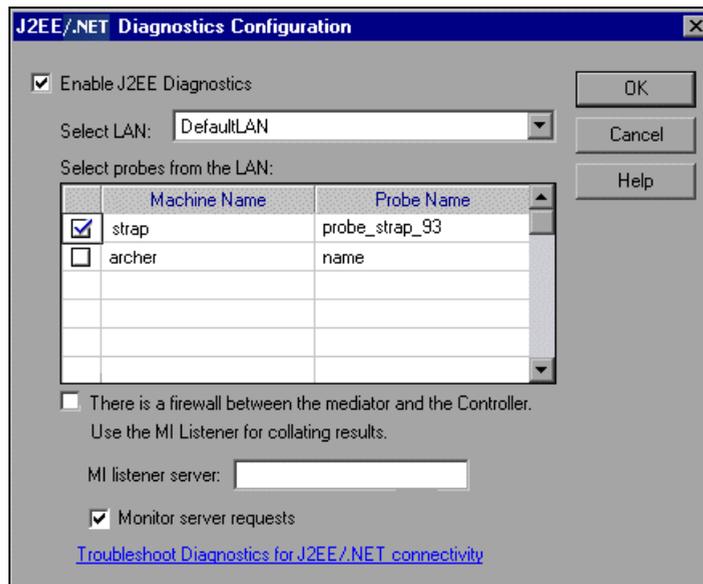
The System Health Monitor opens in your browser.

To access the System Health Monitor from Mercury Business Availability Center:

- 1** Select **Admin > Diagnostics**.
- 2** Click the **System Health** tab. The System Health information is displayed.

To access the System Health Monitor from LoadRunner Controller:

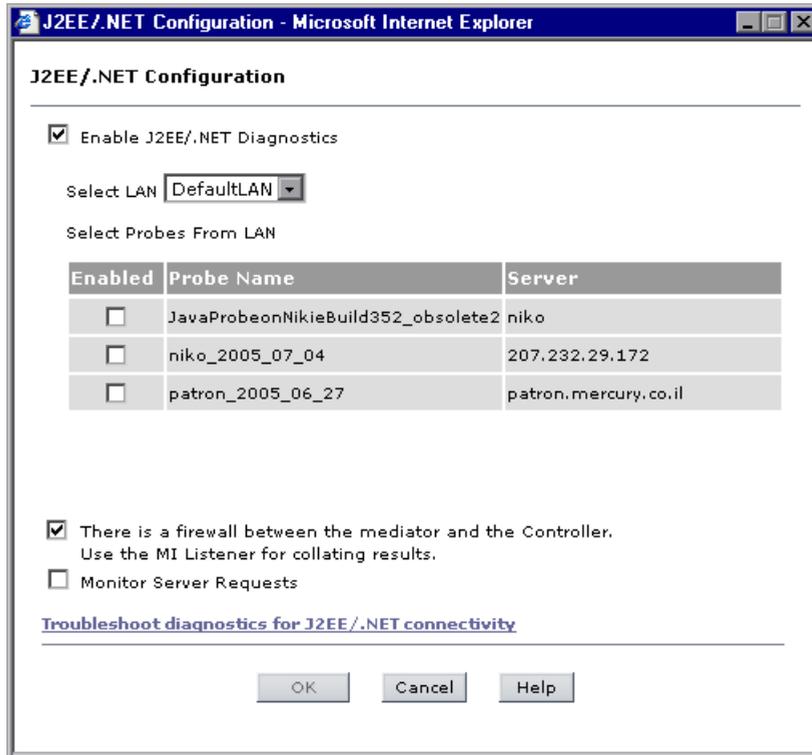
- 1** Choose **Diagnostics > Configuration** from the menu bar in the Controller Window. The Diagnostics Distribution dialog box opens.
- 2** In the Online and Offline Diagnostics section of the Diagnostics Distribution dialog box, click **Configure**. The J2EE/.NET Diagnostics Configuration dialog box opens.



- 3** Click the **Troubleshoot Diagnostics for J2EE/.NET connectivity** link. The System Health Monitor is displayed in a new browser window.

To access the System Health Monitor from Performance Center:

- 1** In the Load Tests page, click the **Diagnostics** tab, and select the **Enable Diagnostics** check box to enable Diagnostics.
- 2** In the **Offline and Online Diagnostics** section, click **Configure** to open the J2EE/.NET Configuration dialog box.



- 3** Click the **Troubleshoot Diagnostics for J2EE/.NET connectivity** link. The System Health Monitor is displayed in a new browser window.

Troubleshooting Accessing the System Health Monitor

If you are unable to access the System Health Monitor verify that:

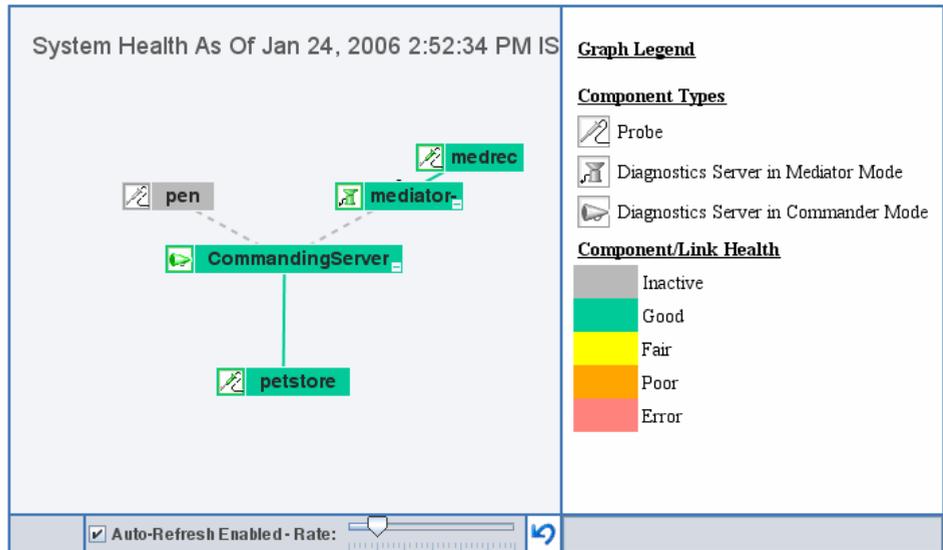
- ▶ the machine name for the Diagnostics Server host is correct.
- ▶ the port number for the Diagnostics Server is correct. The default port is 2006.
- ▶ the Diagnostics Server started successfully, and is running on the host.

Understanding the System Health Monitor

The System Health Monitor is comprised of two panes:

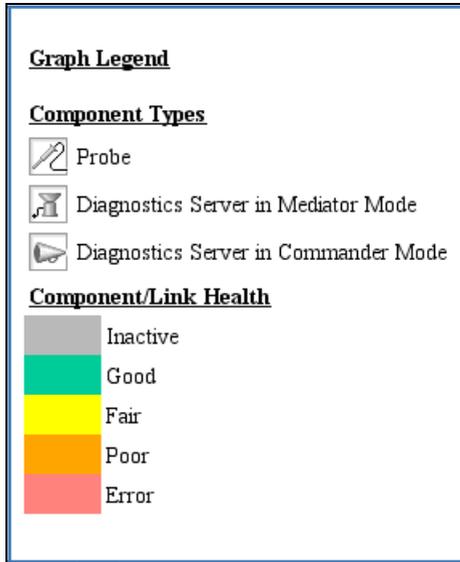
- ▶ **The graph legend** (in the right pane)
- ▶ **The System Health Monitor component map** (in the left pane)

When the System Health Monitor opens you see a deployment map of the Mercury Diagnostics components, accompanied by the graph legend.



The Graph Legend

The graph legend, to the right of the component map, helps you to interpret the information displayed in the component map.



The System Health Monitor Component Map

The System Health Monitor component map displays high-level information about your Diagnostics deployment and how well the components are performing. The System Health Monitor component map displays the following:

- ▶ Icons that represent each of the Diagnostics components that you have correctly configured as part of your Diagnostics environment.
- ▶ The color of the component icon that indicates the health of the component.
- ▶ An indicator of the amount of load that the component is processing.

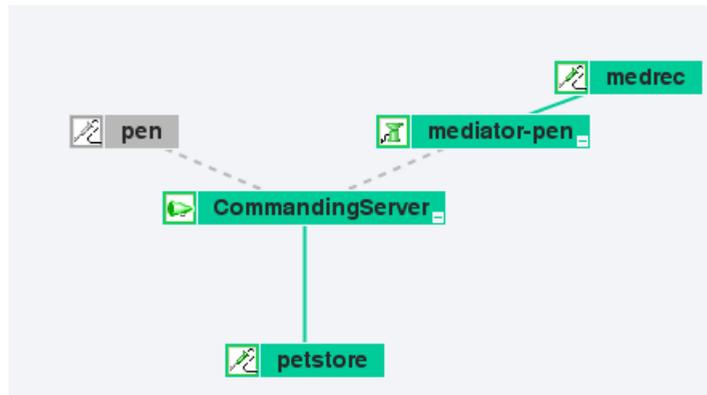
- Links between the component icons, which indicate the state of the communication links between components.

The color of the link indicates the status of the link.

- ◆ A solid line means data is flowing across the link.
- ◆ A dashed line means there is currently no data flowing.

Component Icons in the System Health Monitor Component Map

In the following image of a Diagnostics deployment, one Diagnostics Server is deployed in Commander mode and one Diagnostics Server is deployed in Mediator mode. There are Diagnostics probes reporting to both Diagnostics Servers.



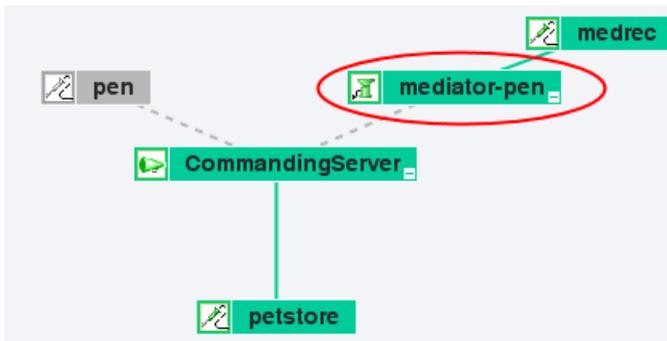
Diagnostics Server in Commander mode

When a Diagnostics Server is deployed in Commander mode, it has both commanding and mediating responsibilities. A Diagnostics Server in Commander mode is represented in the System Health Monitor component map by an icon labelled **CommandingServer**.



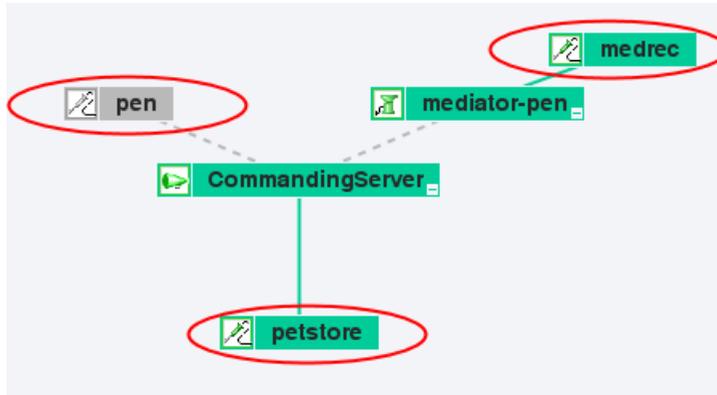
Diagnostics Server in Mediator mode

When a Diagnostics Server is deployed in Mediator mode, it only has mediating responsibilities. A Diagnostics Server in Mediator mode is represented by an icon labelled **mediator-<name of host>**.



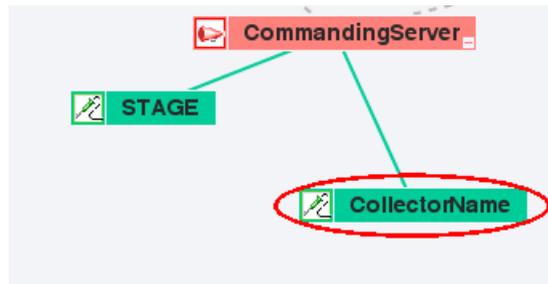
Diagnostics Probes

The System Health Monitor component map displays all of the probes that are reporting to each particular Diagnostics Server. A probe is represented by an icon labelled <name of probe>.



Diagnostics Collector

The System Health Monitor component map displays all of the Diagnostics Collectors that are reporting to each particular Diagnostics Server. A Collector is represented by an icon labelled <CollectorName>.



Viewing Component Status and Host Configuration Tooltips

You can view tooltips for each component on the System Health Monitor component map. The tooltips provide you with basic status and host configuration information.

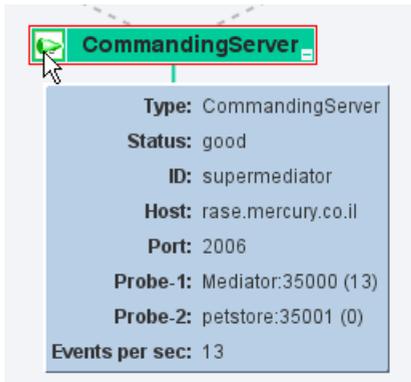
To view component status and host configuration tooltips:

Hold the mouse pointer over the component until a tooltip is displayed.

The examples that follow, show tooltips displaying component status and host configuration information for each different component.

The information in the tooltip for the Diagnostics Server components includes a list of the probes that have been configured to work with the Diagnostics Server.

► Diagnostics Server in Commander mode



► Diagnostics Server in Mediator mode

The screenshot shows a tooltip for a component named 'mediator-MERLOT'. The tooltip is a light blue box with a white border, containing the following information:

- Type:** mediator
- Status:** good
- ID:** mediator-MERLOT
- Host:** MERLOT.mercury.co.il
- Port:** 2612
- Probe events per sec:** 461
- Probe-1:** _vti_bin.NET.1133653801015632 (0)
- Probe-2:** _layouts.NET.1133653810328132 (0)
- Probe-3:** 1Default.NET.1133653803468757 (0)
- Probe-4:** 902862423Default.NET.1133653459094378 (0)
- Probe-5:** _vti_bin.NET.1133653647781257 (0)
- Probe-6:** Mediator:35000 (461)

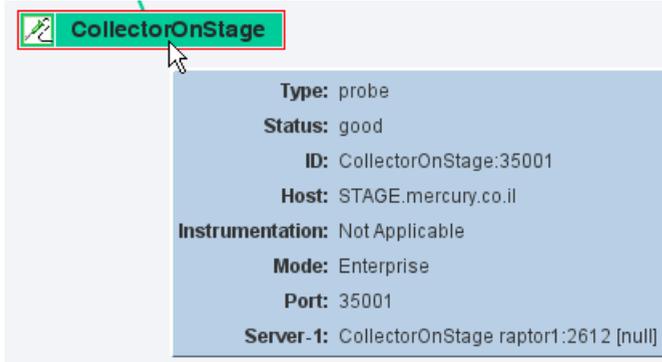
► Diagnostics Probe

The information in the tooltip for a probe component includes the name of the Diagnostics Server that the probe is working with (**Server** field), and the mode that the probe has been configured to work in (**Mode** field). The tooltip also directs you to the relevant instrumentation points file.

The screenshot shows a tooltip for a component named 'PetStoreAbsent'. The tooltip is a light blue box with a white border, containing the following information:

- Type:** probe
- Status:** good
- ID:** PetStoreAbsent:35001
- Events per sec:** 0
- Host:** ABSENT.mercury.co.il
- Instrumentation:** auto_detect.points
- Mode:** Enterprise, PRO
- Port:** 35001
- Server-1:** PetStoreAbsent drug:2612 [null]

► Diagnostics Collector

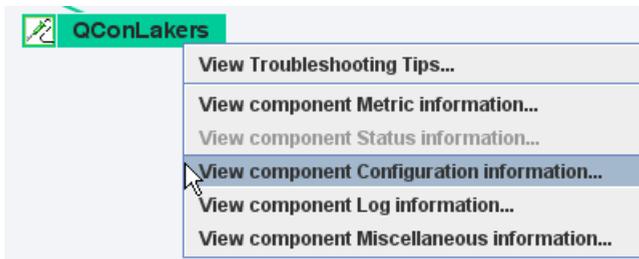


Viewing Detailed Component Information

You can view detailed information about the configuration and performance of each component in the System Health Monitor component map. The information is displayed in the Component Monitor detail table.

To view detailed component information:

Right-click the component and choose the relevant item from the menu.



Alternatively, double-click the component icon and scroll to the relevant section.

You can view the following information:

- ▶ **Component Metric information:** Lists the processing metrics for the selected component.
- ▶ **Component Configuration information:** Lists the component configuration for the selected component.
- ▶ **Component Log information:** Lists the log messages for the selected component.
- ▶ **Component Status Information:** Provides further information about the performance status of the component, where relevant.
- ▶ **Component Miscellaneous information:** Provides miscellaneous information about the component.

Component Metric Information

The following image shows a section of the Component Metric information for the Diagnostics Server in Commander mode:

Component Monitor Update As Of Dec 7, 2005 2:21:12 PM GMT+02:00	
⊗ Metric	
Events since launch	<unknown>
Last update (commander time)	Wed Dec 07 04:22:35 PST 2005
Up time	1 day(s) 0 hour(s) 4 minute(s) 53 second(s)
WDE bytes relayed	166836
WDE bytes relayed since launch	92396096
WDE samples failed	75
WDE samples failed since launch	58870
WDE samples relayed	171
WDE samples relayed since launch	94676

Component Configuration Information

The following image shows the Component Configuration information for the .NET Probe:

Component Monitor Update As Of Dec 7, 2005 2:27:12 PM GMT+02:00	
⊞ Configuration	
Host	ABSENT.mercury.co.il
IP Address	10.168.4.112
Install dir	D:\dotnet_1702\
Instrumentation	ASP.NET.points,MSPetShop.points
LAN ID	DotNetRealProbes
Lightweight Memory Diagnostics	ON
Mode	AM,PRO
Port	35001
System	ABSENT (Microsoft Windows NT 5.2.3790.0)
VM	.NET 1.1.4322.2032
Version	4.0.114.102

Component Log information

The following image shows the Component Log information for the Diagnostics Server in Mediator mode:

Component Monitor Update As Of Dec 7, 2005 2:34:02 PM GMT+02:00	
⊗ Log	
Wed Dec 07 04:06:30 PST 2005	WARNING: real user fragment 31289557 timed out
Wed Dec 07 03:09:02 PST 2005	WARNING: real user fragment 32930633 timed out
Wed Dec 07 03:08:49 PST 2005	WARNING: real user fragment 32687513 timed out
Wed Dec 07 01:08:38 PST 2005	WARNING: real user fragment 8834978 timed out
Tue Dec 06 23:08:41 PST 2005	WARNING: real user fragment 12640068 timed out
Tue Dec 06 05:51:44 PST 2005	SEVERE: out of order com.mercury.diagnostics.server.persistence.impl.I
Tue Dec 06 05:51:44 PST 2005	SEVERE: out of order com.mercury.diagnostics.server.persistence.impl.I
Tue Dec 06 05:46:01 PST 2005	SEVERE: out of order com.mercurv.diaenostics.server.persistence.impl.I

Component Status Information

Below is an example of Component Status information for the J2EE Probe. This information is available when the status of the component drops below **Good**.

Component Monitor Update As Of Dec 7, 2005 3:13:20 PM GMT+02:00	
⊗ Status	
Error	Probe reporting communication problems with mediator: myDomain6 aol:2612 [null]

Component Miscellaneous Information

The following image shows the Component Miscellaneous information for the J2EE Probe:

Component Monitor Update As Of Dec 7, 2005 2:42:42 PM GMT+02:00	
⊗ Miscellaneous	
ProbeProtocolVersion	3

Viewing Troubleshooting Tips

Troubleshooting information is available for each component in the System Health Monitor component map.

To view tips for troubleshooting Diagnostics components:

Right-click the relevant component and choose **View Troubleshooting Tips**. Troubleshooting information for the selected component is displayed.

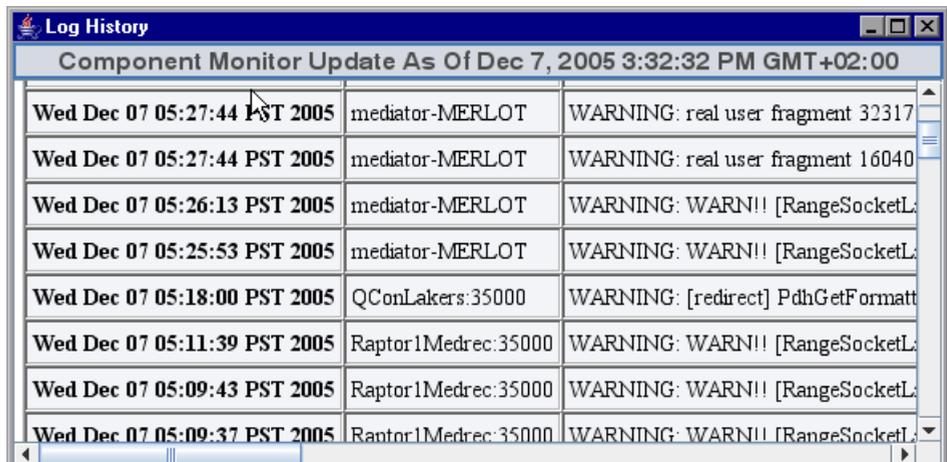
Scan the information displayed for the symptoms that you are investigating, to see if there is a recommended solution documented in the troubleshooting tips.

Viewing Log Information for the Whole System

You can view the log messages for all of the components in the System Health Monitor on one screen.

To view log information for all the components:

Right-click anywhere in the System Health Monitor component map (except on a component icon) and choose **View Log History**. The log for all component activity is displayed.



The screenshot shows a window titled "Log History" with a subtitle "Component Monitor Update As Of Dec 7, 2005 3:32:32 PM GMT+02:00". The window contains a table with three columns: Time, Component Name, and Message. The table lists several warning messages from various components like mediator-MERLOT, QConLakers, and Raptor1Medrec.

Time	Component Name	Message
Wed Dec 07 05:27:44 PST 2005	mediator-MERLOT	WARNING: real user fragment 32317
Wed Dec 07 05:27:44 PST 2005	mediator-MERLOT	WARNING: real user fragment 16040
Wed Dec 07 05:26:13 PST 2005	mediator-MERLOT	WARNING: WARN!! [RangeSocketL
Wed Dec 07 05:25:53 PST 2005	mediator-MERLOT	WARNING: WARN!! [RangeSocketL
Wed Dec 07 05:18:00 PST 2005	QConLakers:35000	WARNING: [redirect] PdhGetFormatt
Wed Dec 07 05:11:39 PST 2005	Raptor1Medrec:35000	WARNING: WARN!! [RangeSocketL
Wed Dec 07 05:09:43 PST 2005	Raptor1Medrec:35000	WARNING: WARN!! [RangeSocketL
Wed Dec 07 05:09:37 PST 2005	Raptor1Medrec:35000	WARNING: WARN!! [RangeSocketL

Customizing the System Health Monitor Display

You can customize the way information is displayed on the System Health Monitor.

Manipulating the Appearance of the System Health Monitor Component Map

You can manipulate the System Health Monitor component map to display any area of the tree, to display expanded branches or collapsed branches, and to change the overall panorama of the image.

To manipulate the System Health Monitor component map display:

- ▶ Click a component to move it to the center of the map.
- ▶ Click anywhere in the component map to change the emphasis of the map. The components will center around the point where the mouse was clicked.
- ▶ Click and drag anywhere in the component map to move and rotate the components around that point.
- ▶ Expand and collapse the branches of a sub-tree by clicking the expand (+) or collapse (-) symbols, displayed on the lower-right corner of the component icon.
- ▶ Increase or reduce the gap between each branch by holding down the ALT button on the keyboard and dragging in the component map.

Displaying the Graph Legend

You can choose whether or not to display the graph legend on the System Health Monitor.

To show the graph legend:

If the graph legend is not displayed, right-click anywhere in the System Health Monitor component map (except on a component icon) and choose **Show Graph Legend**.

To hide the graph legend:

If the graph legend is displayed, right-click anywhere in the System Health Monitor component map (except on a component icon) and choose **Hide Graph Legend**.

Displaying Load

You can choose the way in which load is displayed on the System Health Monitor. You can choose between displaying the load indicator as a circle that appears behind the component icon, or as a bar scale that appears below the component icon.

To display load as a circle:

When load is displayed as a bar scale, right-click anywhere in the System Health Monitor component map (except on a component icon) and choose **Show Load Using Circles**.

In the following example, the circle behind the component represents the amount of load.



To display load as a bar scale:

When load is displayed as circles, right-click anywhere in the System Health Monitor component map (except on a component icon) and choose **Show Load Using Scale**.

In the following example, the bar scale below the component represents the amount of load.

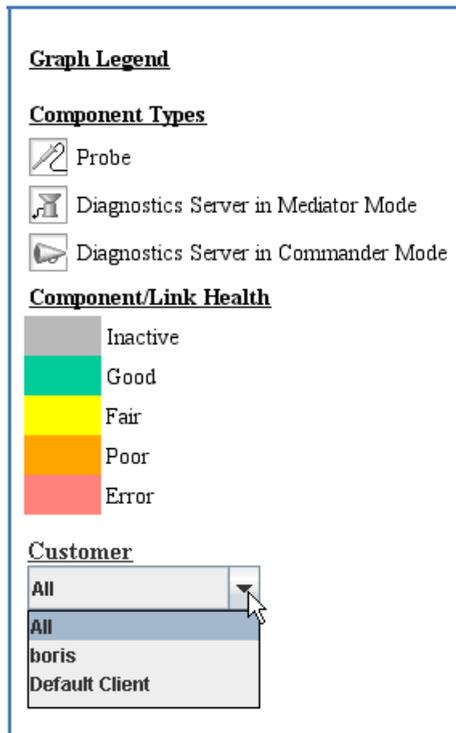


Filtering the System Health Monitor Component Map by Customer

When you are running Diagnostics in an Mercury Managed Services environment with multiple customers, you can filter the information that is displayed in the System Health Monitor so that only the information for a selected customer is displayed in the graph.

By default, the System Health Monitor will display the components for all of the customers.

When more than one customer is represented in the Diagnostics deployment displayed in the System Health Monitor, a list of customers is displayed in the Graph Legend box.



Select the customer whose components you would like to see displayed in the System Health Monitor, or select **All** to display the components for all customers at once.

Controlling the Refresh Rate of the System Health Monitor

By default, the System Health Monitor is configured to automatically refresh the information displayed. You can change the auto-refresh frequency, or completely disable the auto-refresh feature. You can also request a manual refresh at any time.

The controls at the bottom of the System Health Monitor that are used to control the refresh of the information displayed, are shown below.



To disable the automatic refresh feature:

Clear the **Auto-Refresh Enabled** check box.

To enable the automatic refresh feature:

Select the **Auto-Refresh Enabled** check box.

To adjust the auto-refresh rate:

Slide the **Rate** slide control to the right to decrease the auto-refresh frequency and to the left to increase the frequency.

To refresh the display manually:

Click the **Refresh** button located to the right of the slide bar.

Creating and Using System Health Monitor Snapshots

You can take snapshots of the System Health Monitor to share the information with others. You can also import snapshots that others have taken of their System Health Monitors to learn information about their systems.

Exporting a Snapshot of the System Health Monitor

You can export the information displayed on System Health Monitor to an .xml formatted file so that you can share the information with others. This is especially useful when you need help diagnosing a problem. The information in the System Health .xml file can be viewed in .xml format or can be imported into any working copy of the System Health Monitor.

Note: Remember! People who have access to the Diagnostics Server in Commander mode host can view the System Health Monitor directly using their Web browser if you give them the URL.

To export a System Health Monitor snapshot as an .xml file:

- 1 Enter the following URL in your Web browser:

`http://<Diagnostics Server in Commander mode host>:<Diagnostics Server in Commander mode port>/registrar/xml`

The default port is 2006. If you have configured the Diagnostics Server to use an alternate port, use that port number in the URL

- 2 Using the **Save** menu option in your Web browser, save the Web page to a file. Be sure to give the file a name that will help you remember why you saved the snapshot. For example, to remember that you took a snapshot of a System Health Monitor with a poor performing Diagnostics Server in Mediator mode, you may want to name the file:

`sys_health_mediator_yyyymmdd.xml`

Importing a Snapshot of a System Health Monitor

If someone gives you an .xml snapshot of a System Health Monitor, you can view it from an existing System Health Monitor.

To import a System Health Monitor snapshot:

- 1 Copy the System Health snapshot file to a directory on the Diagnostics Server in Commander mode host machine.
- 2 Enter the following URL in your Web browser:

```
http://<Diagnostics Server in Commander mode host>:<Diagnostics Server in  
Commander mode port>/registrar/health?xml=<XML_file>
```

where <XML_file> is the path to the .xml file on the Diagnostics Server in Commander mode host machine.

For example, if you copied **sys_health_mediator_yyyymmdd.xml** to the hard drive (C:) on the Diagnostics Server in Commander mode host machine, whose name is **jake**, the URL might look like this:

```
http://jake:2006/registrar/health?xml=c:\sys_health_mediator_yyyymmdd.xml
```

The <Diagnostics Server in Commander mode port>, by default, should be 2006. If you have configured the Diagnostics Server to use an alternate port, use that port number in the URL.

When the System Health Monitor is displayed in your Web browser, the information is from the imported .xml file.

D

Diagnostics Data Management

This appendix provides a detailed description of how Mercury Diagnostics data is managed and stored.

This appendix describes:	On page:
About Diagnostics Data	478
Custom View Data	478
Performance History Data	480
Data Storage Size & Data Retention	483
Pre-Installation Data Management Considerations	485
Backing Up Diagnostics Data	486
Handling Diagnostics Data when Upgrading Diagnostics	489

About Diagnostics Data

There are two main types of Diagnostics data. The first consists of the custom views that each user has created. The second consists of the Diagnostics data collected by the probes and aggregated by the Diagnostics Servers. This data is stored in a time-series database on the Diagnostics Servers. Each Diagnostics Server stores the data that is collected by the probes that report to it. In addition, the Diagnostics Server in Commander mode stores the virtual transactions' data both for AD and AM. The organization and maintenance of the data files that make up the data base are described in the following discussion.

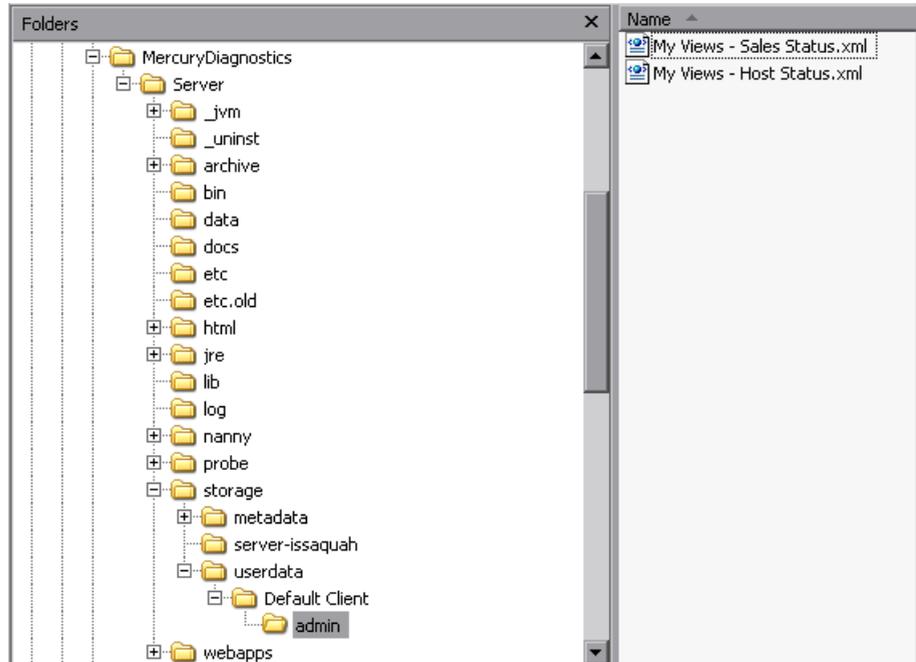
Custom View Data

Diagnostics users can create and save customized views as described in the chapter, “Customizing Diagnostics Views,” in the *Mercury Diagnostics User's Guide*. Diagnostics stores the customized views as XML files on the host for the Diagnostics Server in Commander mode.

Custom View Data Organization

The user defined custom views are stored as XML files in the `<diagnostics_server_install_dir>/storage/userdata` directory on the host for the Diagnostics Server in Commander mode. The custom view files are relatively small with a maximum size of approximately 100 K each.

Each user that has defined a custom view has their own custom view sub-directory in the **userdata** directory. For example, if the 'admin' user created two custom views, Sales Status and Host Status, the two views would be stored as separate xml files in the `<diagnostics_server_install_dir>/storage/userdata/Default Client/admin` directory on the Diagnostics Server in Commander mode as shown in the following example.



Performance History Data

Diagnostics stores the historical performance data in a time series database on the Diagnostics Server in Mediator mode. If the Diagnostics Server has numerous probes reporting to it, the stored historical performance data can grow to many gigabytes of data. Although the amount of data collected for each application can vary in size, it is recommended that you plan for approximately 3 GB of data for each virtual machine that you are monitoring. For more information, see “Data Storage Size & Data Retention” on page 483.

Performance History Data Organization

The Diagnostics performance history data is in the `<diagnostics_server_install_dir>/archive/mediator-<host_name>/persistence/<customer_name>_` directory of the Diagnostics Server in Mediator mode where:

- ▶ `<host_name>` is the name of the host for the Diagnostics Server in Mediator mode
- ▶ `<customer_name>` is the customer name that you entered when you installed the Diagnostics Server in Mediator mode. Note that the name of this directory is the customer name with an appended underscore.

Note: Unless you are an MMS customer, the customer name should always be Default Client.

Performance History Data File Types

The Diagnostics performance history data is stored in five types of files:

Symbol Table Files

The symbol tables contain string-to-integer mappings for small and fast data encoding of the other data files. For example, `/login.do` might be encoded as `1347854`.

The symbol tables are stored in two files in the `<diagnostics_server_install_dir>/archive/mediator-<host_name>/symboltable` directory:

Per-time-Period Performance Summary Files

In the Diagnostics views, data can be viewed in yearly, quarterly, monthly, weekly, daily, 6-hourly, 1-hour, 20-minute, and 5-minute periods. Each viewable time period is stored in separate *summary* files. For example, over the course of one week, the following summary files would be created:

- 1 weekly summary file
- 7 daily summary files
- 28 six-hourly summary files
- 168 hourly summary files
- 504 twenty-minute summary files
- 2,016 five-minute summary files

Each summary file is named according to the minute (in GMT) at which Diagnostics started storing summary data in it. For example, a summary file that contained data beginning at midnight (GMT) on November 23, 2006 would be named **2006_11_23_0.summary**.

The summary files are stored in sub-directories based upon the duration of the data that they contain. The subdirectories are **Weeks/7d**, **Days/1d**, **Days/6h**, **Hours/1h**, **Hours/20m**, **Hours/5m**. For example, a summary file for one of the daily time periods would be stored in the sub-directory `<diagnostics_server_install_dir>/archive/mediator-<host_name>/persistence/<customer_name>_ /Days/1d`.

The data from the summary files are used to create the graphs and tables shown in the Diagnostics views.

Per-Major-Time-Period Performance Trend Files

Diagnostics stores the charted trend data for each major time period in *trend* files. The major time periods are months, weeks, days, and hours. The trend files are named according to the minute (in GMT) at which the trended data that they contain was captured. For example, a file that contained hourly trend data starting at midnight November 23, 2006 would be named as follows: **Hours/1h/ 2006_11_23_0.trend**.

The trend files are accessed when the user selects an element to be charted on the UI. To retrieve a trend for a minor time period, such as 5 minutes, only a small portion of the data in a major time period trend file is read.

Per-Major-Time-Period Instance Tree Files

The *instance tree* files are similar to the trend files; there is a corresponding dump of the collected instance call trees for each major time period. For example, **Hours/1h/2005_11_23_0.tree**.

Instance tree files are accessed when the user drills down to an instance call tree on the Diagnostics UI.

Work Journal File

Diagnostics stores a work journal in `<diagnostics_server_install_dir>/archive/mediator-<host_name>/persistence/<customer_name>_<run_name>/aggregationMarks`. The *journal* contains the name of summary files that have been rolled up into the summary trend and instance tree files for larger time periods.

Data Storage Size & Data Retention

Diagnostics uses data compression and data retention strategies that allow it to optimize its use of disk storage.

Data File Compression

Diagnostics performance data for an application is stored in multiple files that represent different time periods. However, Diagnostics only *writes* new diagnostics data to the files that represent the *current* time period. This means that most of the Diagnostics data files for an application can be stored as read-only files

Since the historical data files are quite large they are automatically compressed after each time period is complete in order to save disk space. Compressed files have the same files names with a **.zip** extension.

Data File Retention

In order to make optimum use of disk space, historical performance data is stored in tiers with the data in each tier retained based upon the diagnostics data retention policy. The data in the tiers with lower resolution data points is kept for longer periods of time to assist with such activities as capacity planning. The data in the tiers with high resolution data points is kept for shorter periods of time to assist with such activities as performance diagnostics. For this retention policy, measurements have shown that Diagnostics uses approximately 3 GB for each probed virtual machine.

The Diagnostics data retention policy is illustrated in the table below:

Directory	Data is kept for ...	Viewable Periods	Trend Resolution
Years	5 Years	Year, Quarter	1 Day
Months	24 Months	Month	6 Hours
Weeks	52 Weeks	Week	1 Hour
Days	93 Days	Day & 6 Hours	5 minutes
Hours	72 Hours	Hour, 20 minutes & 5 minutes	5 seconds

Data File Retention and Trended Data in the Diagnostics Views

When Diagnostics displays trended metrics in the Diagnostics views, it attempts to show approximately sixty data points for each viewable period so that the trend that is presented will be meaningful and easy to understand. To arrive at the data points needed for each viewable period, the Diagnostics Server averages the data from the appropriate tier in the data files. For example, when you are looking at trended data for the last hour, one data point per minute is shown in the graph. These data points would have been created by averaging raw data points for twelve 5-second time periods.

Pre-Installation Data Management Considerations

When preparing to install and configure a large Diagnostics Server, you should consider the following performance tuning recommendations:

- ▶ For maximum performance, the Diagnostics Server should be installed onto an empty or recently defragmented disk. The archive directory should be stored on that same disk. Alternatively, the archive directory could be mounted on an empty or recently defragmented disk.

Note: It is recommended that the disk used for diagnostics time series database that is stored in the `<diagnostics_server_install_dir>/archive` not be used for other disk activity (i.e. don't mount the `<diagnostics_server_install_dir>/archive` directory on the same disk that is used for your system files, temporary files etc...)

- ▶ To reduce fragmentation over time and increase system performance a separate disk (or partition) dedicated to the archive directory is recommended.
- ▶ Intensive background disk processes (such as disk de-fragmentation or virus scans) should be disabled on the disk where the archive directory is stored.
- ▶ Network file systems such as NFS or Samba should not be used.

Note: The better the raw performance of the disk that you dedicate to the archive directory of the Diagnostics Server, the more load the Diagnostics Server can handle. Ask your system administrator to make sure the disk mounted for the archive directory is a high-performance disk or array.

Backing Up Diagnostics Data

It is recommended that you back up the Diagnostics data regularly so that it can be restored in the case of a disk or system failure.

Backing up Data Remotely

Remote backup is possible by downloading the Diagnostics data files over HTTP to a local directory, and then backing up that directory using your normal backup procedures.

The Diagnostics Server supports the HTTP If-Modified-Since and Request-Range headers ("rget") to allow standard HTTP mirroring software to download or incrementally update these files.

Diagnostics is installed with a remote backup script stored at `<diagnostics_server_install_dir/bin/remote-backup.sh`. The remote backup script is a Unix script that uses the wget utility (<http://www.gnu.org/software/wget/wget.html>) to download over HTTP incrementally. On Windows, Cygwin (<http://www.cygwin.com/>) can be used to run this script.

The following table lists the `remote-backup.sh` parameters:

Parameter	Description
-h	The host (or IP address) to download from
-o	The directory to store the backup in
-u	The HTTP username to use Default: admin
-P	The HTTP password to use Default: admin
-P	The HTTP port number to use (optional) Default: 2006
-r	The ID of the Diagnostics Server in Mediator mode being read from (for rhttp backups)(optional)

For example, to backup a Diagnostics Server running on the dragonfly machine into the **dragonfly-backup** directory:

```
% mkdir dragonfly-backup
% bin/remote-backup.sh -u admin -p secret -h dragonfly -o dragonfly-backup
```

The data is backed up in the following directories:

Data	Backup Directory
Server configuration	etc/
User custom views	storage/userdata
Raw performance history data	archive/.../persistence/
Symbol table	archive/.../symboltable/

Restoring Data After a Failure

The files in the backup directory are stored in the structure used by the Diagnostics Server.

To restore the time series database from the backup:

- 1** Install a clean Diagnostics Server. The Diagnostics Server is started automatically after the installation completes.
- 2** Shut down the D4.1 Diagnostics Server.
- 3** Ensure that the Diagnostics Server has been shut down by verifying that there are no java/javaw processes in your process list. On Windows systems, you can use the Task Manager to do this and on UNIX systems, you can use ps.
- 4** Delete the <diagnostics_server_install_dir>/archive directory from Diagnostics Server.
- 5** Copy the database backup to replace the <diagnostics_server_install_dir>/archive.

- 6 If the host name for the Diagnostics Server has changed since the backup was taken you must update the directory name that is based upon the Diagnostics Server host name to reflect the new host name.

Rename `<diagnostics_server_install_dir>/archive/mediator-<host-name>` so that `<host-name>` reflects the new Diagnostics Server host name. For example, if host name in the backup was `oldhost` and the new host name is `newhost` you would change

`<diagnostics_server_install_dir>/archive/mediator-oldhost` to
`<diagnostics_server_install_dir>/archive/mediator-newhost`

Index Regeneration

When a restored Diagnostics Server is first started, the indexed data, which was not backed up, must be regenerated. Index regeneration is started automatically in the background and could take several hours to complete. While the indexes are regenerated the Diagnostics server is able to receive events from probes; but some historical data cannot be displayed in the Diagnostics views until the restoration is complete.

Known Limitation

In Mercury Diagnostics version 4.1, binary data is written in the native byte order. This means that a Diagnostics data backup from a Solaris machine (Big Endian) cannot be restored and used on a x86 machine (Windows/Linux -- Little Endian).

Note: The backup script supplied with the Diagnostics Server backs up the data in a specific order. Failing to back up files in the correct order causes the restored backup to be unusable. It is therefore recommended to always use the supplied script to create data backups.

Handling Diagnostics Data when Upgrading Diagnostics

When you upgrade your Diagnostics Server from D4.0 to D4.1 you must follow the following data upgrade instructions in order to be able to continue to use the custom views and the Diagnostics data that was stored on the old Diagnostics Server.

Upgrading Performance History Data

The following instructions enable you to use the performance history data from the D4.0 version of the Diagnostics Server in the D4.1 Diagnostics Server.

To move your Diagnostics data from a D4.0 Diagnostics Server:

- 1** Shut down the D4.0 Diagnostics Server.
- 2** Ensure that the Diagnostics Server has been shut down by verifying that there are no java/javaw processes in your process list. On Windows systems, you can use the Task Manager to do this and on UNIX systems, you can use ps.
- 3** Create a backup copy of the <diagnostics_server_install_dir>/archive directory and the <diagnostics_server_install_dir>/storage/userdata directory in a location where they will not be impacted by the uninstall of the D4.0 Diagnostics Server or the installation of the D4.1 Diagnostics Server.
- 4** Uninstall the D4.0 Diagnostics Server as instructed in “Uninstalling the Diagnostics Server” on page 57.
- 5** Install the D4.1 Diagnostics Server. The D4.1 Diagnostics Server is started automatically after the installation has finished.
- 6** Shut down the D4.1 Diagnostics Server.
- 7** Ensure that the Diagnostics Server has been shut down by verifying that there are no java/javaw processes in your process list. On Windows systems, you can use the Task Manager to do this and on UNIX systems, you can use ps.
- 8** Delete the <diagnostics_server_install_dir>/archive directory for the D4.1 Diagnostics Server.

- 9 Copy the backup copy of the D4.0 `<diagnostics_server_install_dir>/archive` that you created into the D4.1 `<diagnostics_server_install_dir>/archive`.
- 10 If the host name for the D4.1 Diagnostics Server is different than the host name for the D4.0 Diagnostics Server, you must update the directory name that is based upon the Diagnostics Server host name to reflect the new host name.

Rename `<diagnostics_server_install_dir>/archive/mediator-<host-name>` so that `<host-name>` reflects the D4.1 Diagnostics Server host name. For example, if your D4.0 host name was `oldhost` and the D4.1 host name is `newhost` you would change

`<diagnostics_server_install_dir>/archive/mediator-oldhost` to
`<diagnostics_server_install_dir>/archive/mediator-newhost`

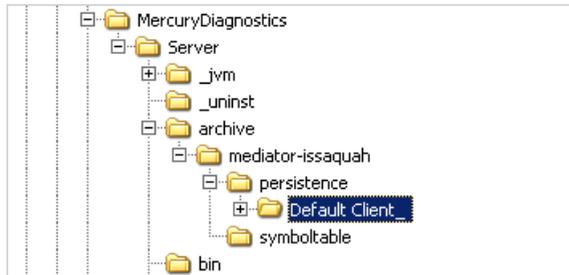
- 11 Move all of the contents of the `<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence` directory to a safe back up directory called `PERCISTENCE_BACKUP`.
- 12 Create the customer name directory inside the `<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence` directory.

The customer name directory is new for D4.1. Create a new directory that is named based on the customer name that you defined when you installed the Diagnostics Server. The new directory name must be the customer name with an underscore appended to it as follows:

`<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence/CustomerName_`.

Note: Unless you are an MMS customer, the customer name should always be `Default Client`.

When you have created the new directory, the directory structure must look as follows:



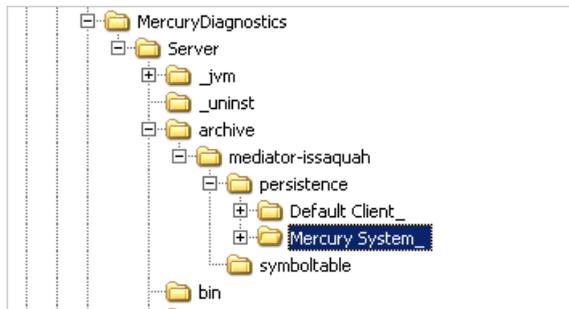
In this example the customer name was Default Client.

- 13** Create the Mercury System directory inside the `<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence` directory.

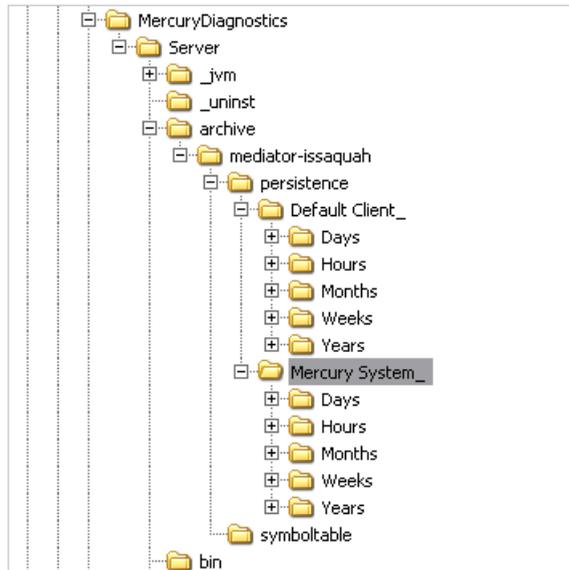
The Mercury System directory is new for D4.1. Create a new directory that is named Mercury System_. Be sure to include the underscore at the end of the name. The new directory structure will be as follows:

`<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence/Mercury System_.`

When you have created the new directory, the directory structure must look as follows:



- 14** Copy the contents of the PERSISTENCE_BACKUP that you created from `<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence` in the earlier step into both of the new directories that you just created. Copy the directories in PERSISTENCE_BACKUP to `<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence/CustomerName_` and to `<diagnostics_server_install_dir>/archive/mediator-<host-name>/persistence/Mercury System_`. Each of the new directories should contain the same information as shown in the following image:



- 15** Delete the `<diagnostics_server_install_dir>/storage/userdata/<customer_name>` directory for the D4.1 Diagnostics Server.
- 16** Copy the backup copy of the D4.0 `<diagnostics_server_install_dir>/storage/userdata/<customer_name>` that you created and paste it to the D4.1 `<diagnostics_server_install_dir>/storage/userdata/<customer_name>`.

Note: When you open the custom views that were created in Diagnostics D4.0 for the first time in Diagnostics D4.1, Diagnostics will upgrade the view for any changes that are necessary because of changes that were made to the functionality of Diagnostics. When Diagnostics changes your custom views a message is displayed to let you know that your custom view has been modified.

- 17** Start D4.1 Diagnostics Server. For instructions see “Starting and Stopping the Diagnostics Server” on page 52.

E

Diagnostics Component Configuration and Communication Diagrams

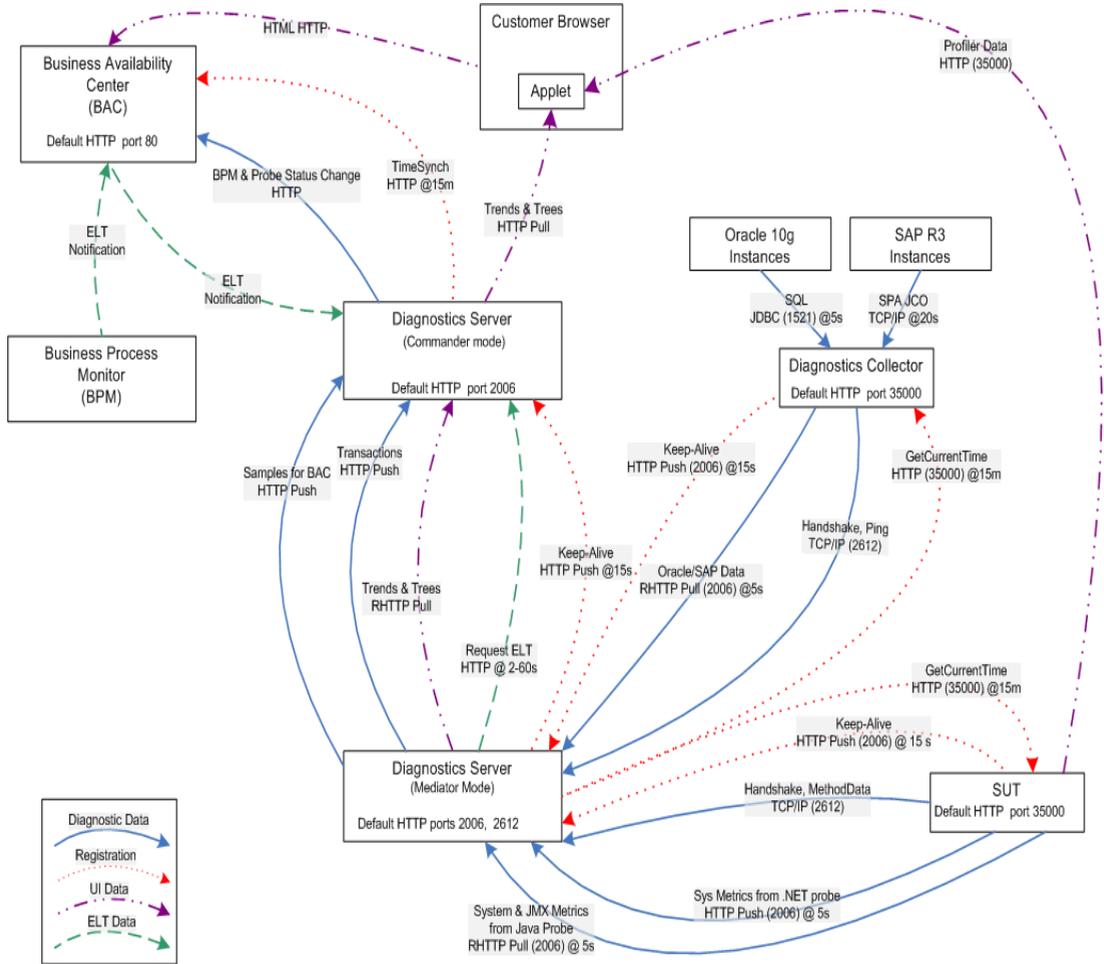
This appendix provides you with a diagram to assist you as you install and configure the Diagnostics components.

This appendix describes:	On page:
Diagnostics Component Configuration and Communication	496

Note: The configuration information described in this appendix is intended for experienced users with in-depth knowledge of Mercury Diagnostics. The diagrams are intended to provide a high-level view, not provide an in-depth knowledge of the working of the components.

Diagnostics Component Configuration and Communication

Integration with Mercury Business Availability Center



F

Diagnostics Upgrade Paths for Mercury Products

This appendix provides you with the recommended paths for upgrading your current Mercury products to use the components included in Mercury Diagnostics 4.1.

This appendix describes:	On page:
Overview of Product Upgrade Paths	497
General Upgrade Recommendations	498
Product Upgrade Paths	499

Overview of Product Upgrade Paths

The following instructions have been provided to help you understand how to upgrade Mercury products that you are currently using so that you can take advantage of the new features of Mercury Diagnostics 4.1.

You should consult with Mercury Customer Support if you have any questions about the information provided here.

General Upgrade Recommendations

The following recommendations are generally applicable when upgrading from earlier versions of Diagnostics or from versions of Products that did not work with Diagnostics.

- ▶ Upgrade decisions should be made on a case-by-case basis—each product deployment is unique, and the recommended upgrade path for one deployment may not apply to another deployment that appears to be similar.
- ▶ Before you install the latest version of a component on the host that was used for the earlier version of the component, make sure that the host can satisfy the system requirements for the new component.
- ▶ If you have decided to upgrade a probe, you should also upgrade the Diagnostics Server in Mediator mode to make sure that you can take full advantage of the latest features of the probe.
- ▶ You should contact Mercury Customer Support when you need to upgrade from an earlier version of Mercury's Application Management products, Topaz, Mercury Business Availability Center, LoadRunner, or Performance Center.
- ▶ Diagnostics D4.0 SP1 is an AM release. The upgrades listed below for LoadRunner provide the instructions for upgrading to the D3.6 version of Diagnostics with the D4.0 SP1 probes. It is also possible to install D4.0 SP1 Diagnostics Standalone and switch the product mode of the D4.0 SP1 probes to AD. This would enable you to see the new features of D4.0 SP1 from the Standalone UI. The new features would not be visible from the Diagnostics UI in LoadRunner. There is no upgrade for the D3.6 Mediator or Commander to D4.0 SP1.

Product Upgrade Paths

The table below lists the upgrade path for each product:

Current Product	Current Mercury Diagnostics Version	Upgrade Path
LoadRunner 7.8 with Transaction Breakdown	None	<ul style="list-style-type: none"> • Upgrade LR to 8.1 • Uninstall All Diagnostics Components • Install All D3.6 Diagnostics Components except for the probes. • Install D4.0 SP1 Diagnostics probes
LoadRunner 8.0	None	<ul style="list-style-type: none"> • Upgrade LR to 8.1 • Install all D3.6 Diagnostics Components except for the probes. • Install D4.0 SP1 Diagnostics probes.
LoadRunner 8.0	D3.0, D3.5	<ul style="list-style-type: none"> • Upgrade LR to 8.1 • Uninstall Diagnostics Server, Database, Mediator • Install D3.6 Commander and Mediator • Upgrade probes to D4.0 SP1 Diagnostics probes.
LoadRunner 8.0	D3.6	<ul style="list-style-type: none"> • Upgrade LR to 8.1 • Upgrade probes to D4.0 SP1 Diagnostics probes.
LoadRunner 8.1	None	<ul style="list-style-type: none"> • Install D3.6 Commander and Mediator • Install D4.0 SP1 Diagnostics probes.

Current Product	Current Mercury Diagnostics Version	Upgrade Path
LoadRunner 8.1	D3.6	<ul style="list-style-type: none"> • Upgrade LR to 8.1 • Upgrade probes to D4.0 SP1 Diagnostics probes.
Performance Center 8.1	None	<ul style="list-style-type: none"> • Install D3.6 Commander and Mediator • Install D4.0 SP1 Diagnostics probes.
Performance Center 8.1	D3.6	<ul style="list-style-type: none"> • Upgrade LR to 8.1 • Upgrade probes to D4.0 SP1 Diagnostics probes.
Mercury Business Availability Center 4.5 FP1 & FP2 Standard	None	<ul style="list-style-type: none"> • Uninstall Mercury Business Availability Center 4.5 • Install Mercury Business Availability Center 6.x • Install all D4.0 SP1 Diagnostics Components <p>Note: There is no data upgrade path.</p>
Mercury Business Availability Center 4.5 FP2 Special Edition	R0, R1	<ul style="list-style-type: none"> • Uninstall Mercury Business Availability Center 4.5 • Uninstall Diagnostics • Install Mercury Business Availability Center 6.x • Install all D4.0 SP1 Diagnostics Components <p>Note: Requires 1 week PSO engagement to tune probe & Diagnostics Server</p>

Current Product	Current Mercury Diagnostics Version	Upgrade Path
Mercury Business Availability Center 4.5 FP2 Special Edition	D3.0	<ul style="list-style-type: none"> • Uninstall Mercury Business Availability Center 4.5 • Uninstall D3.0 Diagnostics Components • Install Mercury Business Availability Center 6.x • Install all D4.0 SP1 Diagnostics Components <p>Note: Requires 1 week PSO engagement to tune probe & Diagnostics Server</p>
Mercury Business Availability Center 5.x	D3.0	<ul style="list-style-type: none"> • Uninstall Mercury Business Availability Center 5.x • Uninstall D3.0 Diagnostics Components • Install Mercury Business Availability Center 6.x • Install all D4.0 SP1 Diagnostics Components <p>Note: Requires 1 week PSO engagement to tune probe & Diagnostics Server</p>
Deep Diagnostics 3.x	D3.x	<ul style="list-style-type: none"> • Upgrade J2EE Probes to D4.0 SP1 using Manual Upgrade Process

G

Troubleshooting Mercury Diagnostics

This appendix provides troubleshooting tips for problems that may occur when using Mercury Diagnostics.

This appendix describes:	On page:
Component Installation Interrupted on a Solaris Machine	503
J2EE Probe Fails to Operate Properly	504
J2EE Probe Throttling Seems Excessive	504
Error During WAS Startup with Mercury Diagnostics Profiler for J2EE	505

Component Installation Interrupted on a Solaris Machine

If a component installer on a Solaris machine is interrupted before it has finished installing the component, there is no option for automatically uninstalling or reinstalling the component. You must manually clean up the partial installation of the component before you can start the installation again.

To manually clean up after an interrupted installation:

- 1** Clean the installation directory.
- 2** Delete `~/vpd.properties` and `~/vpd.patches`.
- 3** Delete the Solaris directories: `/var/sadm/pkg/IS*` and `/var/sadm/pkg/MERQ`.

J2EE Probe Fails to Operate Properly

If the J2EE Probe does not operate properly, check whether the **ClassLoader.class** file located in the folder `<probe_instal_dir>\classes\boot\java\lang\` was created during the installation process.

If the file was not created, run the JRE Instrumenter to create it. See Chapter 8, “Running the JRE Instrumenter.”

J2EE Probe Throttling Seems Excessive

By reviewing the probe logs, you can determine when throttling is being started and stopped. If throttling is being engaged more than you would like you should check to make sure that the probe thread configuration is synchronized with the number of threads that the application server has been configured to allow. If the application server has more threads than the probe can handle, throttling is engaged sooner and may not be disengaged.

For information on tuning the probe to reduce throttling see “Controlling Probe Throttling” on page 329.

Error During WAS Startup with Mercury Diagnostics Profiler for J2EE

Symptoms:

Class Loader errors occur when starting WAS with the Mercury Diagnostics Profiler for J2EE.

Reason:

Additional classes need to be excluded from the instrumentation.

Solution:

- 1 Open the property file, `<probe_install_dir>\etc\inst.properties`
- 2 Update the `classes.to.exclude` property to exclude `!com\.ibm\..*` by appending the class to the end of the existing values.

```
classes.to.exclude=!iaik\.security\.*;!c8e\.*;!org\.jboss\.net\.protocol\.file\.Handler;!org\.jboss\.net\.protocol\.file\.FileURLConnection;!.*ByCGLIB.*;!com\.ibm\..*
```


H

Using UNIX Commands

When running an installation on a UNIX platform, you can usually follow the instructions that appear on the screen. The on-screen instructions can be confusing at times.

If something is unclear, use the following guidelines:

- ▶ To select an option from a list of options, type the number corresponding to the option and press `ENTER`. Then type `0` and press `ENTER` again to confirm your choice.
- ▶ When selecting multiple options, for each selection type the corresponding number and press `ENTER`. Only when you have finished selecting all the options that you want to select, type `0` and press `ENTER` again to confirm your choices.
- ▶ If you have selected an option and want to clear it, retype the corresponding number, or type the number of another option, and press `ENTER`. Then type `0` and press `ENTER` again to confirm your choice.
- ▶ When entering information at a prompt:
 - To accept a default value that is displayed at the prompt, press `ENTER`.
 - Type the information, and press `ENTER` to continue.
- ▶ To continue with the next step of an installation, type `1` to select **Next**, and press `ENTER`.
- ▶ To go back to previous prompts to make changes, type `2` to select **Previous**, and press `ENTER`.
- ▶ To cancel an installation, type `3` to select **Cancel**, and press `ENTER`.
- ▶ To redisplay a prompt, type `4` to select **Redisplay**, and press `ENTER`.

- Appendixes

Index

A

- advanced configuration
 - Diagnostics Server 439
- advanced Mediator assignment 257
- AM product mode, Probe settings 309
- application server
 - configuring *See* application server configuration
- application server configuration 171
 - generic 206
 - JBoss 201
 - multiple instances 312
 - Oracle 9i 198
 - SAP NetWeaver 204
 - WebLogic 6.1 189
 - WebLogic 7 191
 - WebLogic 8 193
 - WebLogic 8.1 193
 - WebSphere 173
 - WebSphere 4.0 174
 - WebSphere 5.x ,6.x 179
- application server, advanced configuration 305
 - Deep Diagnostics allocation capture 326
- application servers, supported 7
- Automatic Method Trimming
 - controlling 326
 - depth 328
 - latency 326

B

- Business Availability Center, *See* Mercury Business Availibility Center

C

- capture points files 343
 - AD- and AM-only point entries 350
 - Deep-Diagnostics-only point entries 351
 - madatory point entries 346
 - optional point entries 346
- ClassLoader class, recreating 504
- Collector 6
- communication diagrams 495
- component configuration 495
- configuring application servers
 - WebLogic 188
- configuring application servers, *See* application server configuration

D

- data flow diagram, Diagnostics 8
- data management, *See* Diagnostics data management
- default Mediator assignment 257
- deployment configuration 9
- deployment, recommended configuration 9
- depth trimming 328
- Diagnostics components 5
 - host requirements 14
 - synchronizing time between 242
- Diagnostics data flow 8
- Diagnostics data management
 - backing up data 486, 489
 - custom screen data 478
 - data sizes & data retention 483
 - performance considerations 485
 - performance history data 480, 481
- Diagnostics Distribution dialog box 457

Index

Diagnostics layers

- .NET layers 355

- J2EE layers 353

- Portal layers 356

Diagnostics Probe for .NET, *See* Probe, .NET

Diagnostics Probe for J2EE, *See* J2EE Probe

Diagnostics Server 6, 22

- about 6

- adjusting heap size 247

- administration 435

- changing default port 250

- configuration pages 258, 437

- configuring 55

- configuring for large installation 247

- configuring for multi-homed

 - environments 252

- configuring LoadRunner offline file 259

- configuring time synchronization 244

- configuring, advanced 241, 435

- HTTPS communication 375

- installing (Unix) 41

- installing (Windows) 30, 59

- jetty.xml file, modifying 252

- jetty.xml file, sample 254

- LoadRunner / Performance Center assignments 257

- modifying properties 441

- overriding default host name 250

- reducing memory usage 256

- setting event host name 252

- starting and stopping 52

- synchronizing time between

 - Diagnostics components 242

 - System Health Monitor 464, 465

 - system requirements 19

 - troubleshooting 58

 - uninstalling 56

 - verifying installation 54

 - viewing advanced configuration options 439

Diagnostics setup menu 30, 92

disable Probes 403

documentation updates xvii

E

event host

- setting name for 252

F

filter

- System Health Monitor 473

H

heap size 247

host requirements, Diagnostics components 14

HTTP proxy communication 385

- .NET Probe 387

- Diagnostics Server in Mediator Mode 386

- J2EE Probe 387

HTTPS communication 374

- enabling for Diagnostics Server in Commanding mode 375

- enabling for Mercury Business

 - Availability Center server 382

I

IIS, restarting 270

installation

- Diagnostics Server (Unix) 41

- Diagnostics Server (Windows) 30, 59

- interruption during 503

- J2EE Probe 111

- planning 21

- Probe for .NET 92

- Probe for .NET, verifying 106

- recommended order 20

- requirements, *See* installation requirements

installation requirements

- .NET Diagnostics Profiler 17

- .NET Probe 17, 18

- Diagnostics Server 14

- J2EE Diagnostics Profiler 16

- J2EE Probe 16

J

- J2EE Diagnostics Profiler
 - disabling 322
 - Probe settings, PRO product mode 308
 - security 337
 - user management 337
 - user permissions 337
 - WAS startup error 505
- J2EE Probe
 - administration page 334
 - configuring and installing, about 112, 154
 - controlling log messages 323
 - excessive throttling 504
 - failed operation 504
 - installing 111
 - installing as a Profiler 116
 - installing on z/OS 146
 - installing using generic installer 149
 - installing with a Mercury Diagnostics Server 120
 - overriding default Probe host name 334
 - permissions 337
 - security 337
 - silent installation 150
 - system requirements 16
 - throttling, controlling 329
 - uninstalling 164
 - upgrading 161
 - user management 337
 - user permissions 337
 - using with Deep Diagnostics 160
- J2EE Probe advanced configuration 305
 - adding a Mercury product 307
 - AM product mode–Application Management Probe settings 309
 - Automatic Method Trimming 326
 - log messages 323
 - PRO product mode–Diagnostics Profiler for J2EE Probe settings 308
 - proxy server 332
 - removing a Mercury product 312
 - reverse HTTP for Probe in MMS 332

- specifying properties as java system properties 332
 - throttling 329
- J2EE/.NET Diagnostics Configuration dialog box 401, 402, 457, 458
- java executable (Windows)
 - changing 146
- java system properties in Probe 332
- JBoss, application server configuration 201
- JDK/JRE executable 116, 124
- jetty.xml file
 - modifying 252
 - sample 254
- JMX metrics
 - accessing 425
 - collecting 428
 - understanding patterns 430
- JRE Instrumenter
 - about 165
 - running 166

L

- latency trimming 326
- light-weight memory Diagnostics (LWMD) 278
- Linux custom metrics 421
- LoadRunner
 - AddIn, installing 227
 - Diagnostics Server details, specifying 232
 - Diagnostics, configuring scenarios to use 234
 - Diagnostics, setting up 231
- LoadRunner Offline File
 - advanced Diagnostics Server configuration 259
 - estimating size of 259
 - reducing size of 260
- log messages 323
- LWMD *See* light-weight memory Diagnostics (LWMD)

M

Mediator assignment
 advanced 257
 default 257
 memory usage, reducing 256
 Mercury Business Availability Center
 assigning permissions for Diagnostics
 users 224
 changing the Diagnostics server
 details 218
 Diagnostics configuration 219
 Diagnostics monitoring, enabling in
 Business Process Monitor (BPM)
 related views 222
 Diagnostics server details, specifying
 214
 Diagnostics view 220
 setting up to use Diagnostics 212
 Mercury Business Availability Center server,
 HTTPS communication 382
 Mercury Customer Support Web site xvii
 Mercury Diagnostics Probe for .NET, *See*
 Probe, .NET
 Mercury Diagnostics Probe for J2EE, *See* J2EE
 Probe
 Mercury Diagnostics Probes 5
 Mercury Home Page xvii
 Mercury Managed Services (MMS)
 configuring reverse HTTP for Probe
 332
 enabling reverse HTTP for Probe 333
 Mercury product, adding to Probe 307
 metrics collector
 modifying default port 415
 system metrics 415
 metrics, Diagnostics
 System Health Monitor 467
 MMS, *See* Mercury Managed Services (MMS)

N

NET (.NET) Probe, *See* Probe, .NET
 NET Diagnostics Profiler
 enabling 272

O

Oracle 9i, application server configuration
 198
 order of installation, recommended 20

P

Performance Center
 Diagnostics server details, specifying
 236
 load tests, configuring to use
 Diagnostics 238
 setting up to use Diagnostics 235
 permissions, *See* user permissions
 preinstallation considerations 19
 PRO product mode, Probe settings 308
 Probe 5
 disable 403
 System Health Monitor 465
 Probe instrumentation
 .NET layers 355
 capture points files 343
 J2EE layers 353
 portal layers 356
 understanding Diagnostics layers 352
 Probe, .NET 91
 configuring 108
 elements and attributes 281
 enabling 271
 installation, verifying 106
 installing 92
 installing as a Profiler 97
 installing with a Mercury Diagnostics
 Server 100
 system requirements 17, 18
 uninstalling 109
 version, determining 109
 Probe, .NET advanced configuration
 ASP.NET applications 264
 classes/methods 265
 customizing instrumentation for
 ASP.NET applications 267
 depth trimming 276
 disabling logging 280
 elements and attributes 281
 latency trimming and throttling 272

- light-weight memory Diagnostics (LWMD) 278
- overriding default Probe host name 280
- restarting IIS 270
- Probe, J2EE, *See* J2EE Probe
- product security 444
- product upgrade paths 497, 499
 - general recommendations 498
- proxy server
 - configuring for Probe 332

R

- refresh rate
 - System Health Monitor 474
- requirements
 - Diagnostics Server 19
- reverse HTTP
 - configuring for Probe in MMS 332
 - disabling for Probe in MMS 333
 - enabling for Probe in MMS 333

S

- SAP NetWeaver, application server
 - configuration 204
- security permissions 445
- Server, Diagnostics *See* Diagnostics Server
- setup menu 30, 92
- snapshots
 - System Health Monitor 475
- Solaris custom metrics 420
- solution templates 359
 - understanding data 360
 - viewing information 361
- System Health Monitor 455
 - accessing 456
 - accessing from browser 456
 - Commander properties 465
 - component configuration
 - information 468
 - component icons 461
 - component information, viewing
 - detailed 466
 - component log information 469

- component map 460
- component metric information 467
- component status information 469
- customizing display 471
- Diagnostics Server in Commander
 - mode properties 464
- Diagnostics Server in Mediator mode
 - properties 465
- displaying legend 471
- displaying load 472
- filter by customer 473
- icons 460
- legend 460
- miscellaneous information 469
- Probe properties 465
- refreshing data manually 474
- refreshing display 474
- snapshots, exporting 475
- snapshots, importing 476
- system log information 470
- tooltips, component status and host
 - configuration 464
 - viewing log history 470
 - viewing troubleshooting tips 470
- system metrics 407, 413
 - capture 407
 - customizing 408
 - customizing in Windows 416
 - customizing on Linux host 421
 - customizing on Solaris host 420
 - default 414
 - metrics collector 408, 415
 - metrics collector entries 408
 - modifying captured metrics 411
 - stopping capture 411
- system requirements
 - Diagnostics Probe for .NET host 17
 - Diagnostics Probe for J2EE host 16
 - Diagnostics Server host 14

T

- throttling 329
 - J2EE Probe 329
 - tuning 331
 - turning off 330

Index

time synchronization 242
troubleshooting Diagnostics 503
troubleshooting Diagnostics Servers 58

U

uninstalling Diagnostics Server 56
UNIX commands 507
updates, documentation xvii
upgrade paths 497, 499

- general recommendations 498

upgrading earlier versions of Diagnostics 18
user permissions 444

- accessing Diagnostics, default users 446
- assigning and managing 447
- assigning for Diagnostics users in Mercury Business Availability Center 224
- managing user details 339, 448, 449

W

WebLogic 6.1, application server configuration 189
WebLogic 7, application server configuration 191
WebLogic 8, application server configuration 193
WebLogic 8.1, application server configuration 193
WebLogic, application server configuration 188
WebSphere 4.0, application server configuration 174
WebSphere 5.x and 6.x, application server configuration 179
WebSphere, application server configuration 173
Windows custom metrics 416

Z

z/OS

- installing the Probe 146