



# Sun GlassFish Enterprise Server 2.1 Troubleshooting Guide



Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054  
U.S.A.

Part No: 820-4339-10  
January 2009

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

# Contents

---

<b>Preface</b> .....	9
<b>1 Overview</b> .....	13
Planning Ahead .....	13
▼ To back up and restore a server configuration .....	14
Identifying the Problem .....	14
▼ Verifying configuration information .....	14
Seeking a Solution .....	16
Verify System Configuration .....	16
Evaluate Messages .....	16
Examine Log Files .....	17
See if the Problem has been Solved Before .....	20
Search the Product Documentation .....	20
Search the Knowledge Base .....	20
Search or Participate in the Online Forum .....	21
Contact Support .....	21
<b>2 Common Problems</b> .....	23
Node Agent Does not Sync Corrupt/Deleted domain.xml .....	24
Solution .....	24
asadmin start-domain Command Fails .....	24
Error: CLI143 There is more than one domain... .....	24
Error: Could Not Start Domain .....	25
Automatically Restarting a Domain or Node Agent on Reboot .....	25
Restarting Automatically on UNIX Platforms .....	26
Restarting Automatically on the Microsoft Windows Platform .....	26
Cannot Find Log Files .....	29

Accessing Local Server Fails ( <a href="http://localhost:8080">http://localhost:8080</a> ) .....	30
Did the Server Start? .....	30
Was the Server Started at the Expected Port? .....	30
Accessing Remote Server Fails .....	31
Is the Server Available Locally? .....	31
Is the Proxy Setting Causing a Problem? .....	32
Cannot Access the Administration Console .....	32
Is the Application Server Running? .....	32
Is the Administration Console Running on the Expected Port? .....	33
Is the Security Manager Disabled? .....	33
Cannot Access a Server Application .....	34
Is the Application Server Running? .....	34
Was Application Deployment Successful? .....	34
Invalid User or Password When Using Don't Prompt Option .....	34
Administrator User Name or Password Not Known .....	35
Server Will Not Start on Windows (Port Conflict) .....	36
Is Another Application Running on the Server's Port? .....	36
Has an Ungraceful Shutdown Occurred on a Previously Running Server? .....	36
Port Conflicts Debugging Multiple Instances on the Same Server .....	37
Description .....	37
Solution .....	37
Two Server Instances Bind to Same Port on Windows .....	37
Description .....	37
Solution .....	38
Error: System cannot find the path specified .....	38
Description .....	38
Solution 1 .....	38
Solution 2 .....	38
Solution 3 .....	38
Application Generates Error persistence.support.JDODataStoreException .....	39
Description .....	39
Solution .....	39
Using asadmin set Command May Produce Unexpected Results .....	39
Description .....	39
Cannot Undeploy Or Redeploy Application With Open Streams to Jar Files (Windows Only) .....	40

Description .....	40
Solution .....	40
Cannot Reinstall Enterprise Server After Manual Deletion of Directories .....	40
Description .....	40
Solution 1 .....	41
Solution 2 .....	41
Cannot Produce a JVM Thread Dump After Server Crash .....	41
Description .....	41
Solution .....	41
Unable to Start Server with 3.5 GB JVM heap .....	41
Unable to Set Up Cluster on OpenSuse Linux .....	42
▼ Changing your JVM installation .....	42
<b>3 HADB Problems .....</b>	<b>43</b>
HADB Database Creation Fails .....	43
Problems Related to Shared Memory .....	44
Too Few Semaphores .....	45
Database Nodes Cannot Be Reached and the Database Does Not Function .....	46
The Management Agents Could Not Establish a Domain .....	46
hadbm create or hadbm addnodes Command Hangs .....	47
ma (Management Agent Process) Crashes .....	47
Server Responds Slowly After Idle Period .....	47
Description .....	47
Solution .....	47
Requests Are Not Succeeding .....	48
Is the Load Balancer Timeout Correct? .....	48
Are the System Clocks Synchronized? .....	48
Is the Enterprise Server Communicating With HADB? .....	49
Session Persistence Problems .....	49
The create-session-store Command Failed .....	49
Configuring Instance-Level Session Persistence Did Not Work .....	50
Session Data Seems To Be Corrupted .....	50
HADB Performance Problems .....	51
Is There a Shortage of CPU or Memory Resources, or Too Much Swapping? .....	51
Is There Disk Contention? .....	52

Is There a Shortage of HADB Data Device Space? .....	53
Is There a Shortage of Other HADB Resources? .....	53
High Load Problems .....	54
Is the Tuple Log Out Of Space? .....	54
Is the node-internal Log Full? .....	55
Are There Enough Locks? .....	55
Can You Fix the Problem by Doing Some Performance Tuning? .....	56
Connection Problem Caused by Lack of Semaphore Resources .....	57
Description .....	57
Solution .....	57
Improving CPU Utilization .....	57
Description .....	57
Solution .....	57
HADB Administration Problems .....	58
hadbm Command Fails: The agents could not be reached .....	58
hadbm Command Fails: command not found .....	59
hadbm Command Fails: JAVA_HOME not defined .....	59
hadbm createdomain fails, but two split domains are created .....	60
create Fails: path does not exist on a host .....	60
Database Does Not Start .....	61
clear Command Failed .....	62
create-session-store Failed .....	62
hadbm Command Hangs .....	63
Cannot Restart the HADB .....	64
Shared Memory Segment Key Already Exists (Windows only) .....	65
Description .....	65
Solution .....	65
Failure in configure-ha-cluster .....	65
Description .....	65
Solution 1 .....	66
Solution2 .....	66
Unable to Run configure-ha-cluster .....	66
Description .....	66
Solution .....	66
hadbm set Command Fails .....	66
Description .....	66

---

Solution .....	67
Failure in configure-ha-cluster: Creating an HADB Instance Fails .....	67
Description .....	67
Solution .....	67
<b>4 Security Problems .....</b>	<b>69</b>
java.security.AccessControlException: Access Denied Error .....	69
Description .....	69
Solution .....	69
javax.ejb.AccessLocalException: Client Not Authorized Error .....	70
Description .....	70
Solution .....	70
Authentication is Not Working With the Solaris Realm .....	70
Mutual Authentication Not Working With the Application Client .....	71
Description .....	71
Solution .....	71
<b>5 Frequently Asked Questions .....</b>	<b>73</b>
What Happens When No Server Side Realm is Configured? .....	73
Can I Use a PKCS12 Certificate for My Client Certificate? .....	73
Can I See the TLS/SSL Handshake Information for an SSL Client? .....	74
Can I Change the Keystore Password? .....	74
How Do I Maintain a Session in JAX-RPC? .....	74
How Do I Access the Naming Service From a Standalone Java Client? .....	75
▼ To access the naming service from an application client .....	75
Are RMI-IIOP Stubs Needed to Access Remote EJBs? .....	76
How Do I Change the Log Level for an Application Logger? .....	77
<b>Index .....</b>	<b>79</b>





# Preface

---

Troubleshooting solutions for the Sun GlassFish™ Enterprise Server product.

This preface contains information about and conventions for the entire Sun GlassFish Enterprise Server documentation set.

## Sun GlassFish Enterprise Server Documentation Set

TABLE P-1 Books in the Enterprise Server Documentation Set

Book Title	Description
<i>Documentation Center</i>	Enterprise Server documentation topics organized by task and subject.
<i>Release Notes</i>	Late-breaking information about the software and the documentation. Includes a comprehensive, table-based summary of the supported hardware, operating system, Java™ Development Kit (JDK™), and database drivers.
<i>Quick Start Guide</i>	How to get started with the Enterprise Server product.
<i>Installation Guide</i>	Installing the software and its components.
<i>Application Deployment Guide</i>	Deployment of applications and application components to the Enterprise Server. Includes information about deployment descriptors.
<i>Developer's Guide</i>	Creating and implementing Java Platform, Enterprise Edition (Java EE platform) applications intended to run on the Enterprise Server that follow the open Java standards model for Java EE components and APIs. Includes information about developer tools, security, debugging, and creating lifecycle modules.
<i>Java EE 5 Tutorial</i>	Using Java EE 5 platform technologies and APIs to develop Java EE applications.
<i>Java WSIT Tutorial</i>	Developing web applications using the Web Service Interoperability Technologies (WSIT). Describes how, when, and why to use the WSIT technologies and the features and options that each technology supports.
<i>Administration Guide</i>	System administration for the Enterprise Server, including configuration, monitoring, security, resource management, and web services management.

TABLE P-1 Books in the Enterprise Server Documentation Set (Continued)

Book Title	Description
<i>High Availability Administration Guide</i>	Setting up clusters, working with node agents, and using load balancers.
<i>Administration Reference</i>	Editing the Enterprise Server configuration file, <code>domain.xml</code> .
<i>Performance Tuning Guide</i>	Tuning the Enterprise Server to improve performance.
<i>Reference Manual</i>	Utility commands available with the Enterprise Server; written in man page style. Includes the <code>asadmin</code> command line interface.

## Default Paths and File Names

The following table describes the default paths and file names that are used in this book.

TABLE P-2 Default Paths and File Names

Placeholder	Description	Default Value
<i>as-install</i>	Represents the base installation directory for Enterprise Server.	Solaris™ and Linux installations, non-root user: <i>user's-home-directory/SUNWappserver</i>  Solaris and Linux installations, root user: <i>/opt/SUNWappserver</i>  Windows, all installations: <i>SystemDrive:\Sun\AppServer</i>
<i>domain-root-dir</i>	Represents the directory containing all domains.	All installations: <i>as-install/domains/</i>
<i>domain-dir</i>	Represents the directory for a domain.  In configuration files, you might see <i>domain-dir</i> represented as follows:  <code>\${com.sun.aas.instanceRoot}</code>	<i>domain-root-dir/domain-dir</i>
<i>instance-dir</i>	Represents the directory for a server instance.	<i>domain-dir/instance-dir</i>
<i>samples-dir</i>	Represents the directory containing sample applications.	<i>as-install/samples</i>
<i>docs-dir</i>	Represents the directory containing documentation.	<i>as-install/docs</i>

## Typographic Conventions

The following table describes the typographic changes that are used in this book.

TABLE P-3 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
<b>AaBbCc123</b>	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>AaBbCc123</i>	A placeholder to be replaced with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized (note that some emphasized items appear bold online)	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file.

## Symbol Conventions

The following table explains symbols that might be used in this book.

TABLE P-4 Symbol Conventions

Symbol	Description	Example	Meaning
[ ]	Contains optional arguments and command options.	<code>ls [-l]</code>	The <code>-l</code> option is not required.
{   }	Contains a set of choices for a required command option.	<code>-d {y n}</code>	The <code>-d</code> option requires that you use either the <code>y</code> argument or the <code>n</code> argument.
`\${ }`	Indicates a variable reference.	<code>\${com.sun.javaRoot}</code>	References the value of the <code>com.sun.javaRoot</code> variable.
-	Joins simultaneous multiple keystrokes.	Control-A	Press the Control key while you press the A key.
+	Joins consecutive multiple keystrokes.	Ctrl+A+N	Press the Control key, release it, and then press the subsequent keys.

TABLE P-4 Symbol Conventions (Continued)

Symbol	Description	Example	Meaning
→	Indicates menu item selection in a graphical user interface.	File → New → Templates	From the File menu, choose New. From the New submenu, choose Templates.

## Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (<http://www.sun.com/documentation/>)
- Support (<http://www.sun.com/support/>)
- Training (<http://www.sun.com/training/>)

## Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

---

**Note** – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

---

## Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions.

To share your comments, go to <http://docs.sun.com> and click Feedback. In the online form, provide the document title and part number. The part number is a seven-digit or nine-digit number that can be found on the title page of the book or at the top of the document.

# Overview

---

This chapter provides a description of the tools, methods, and information sources available for troubleshooting the Sun GlassFish Enterprise Server. Guidelines for evaluating and investigating a problem are included.

- “Planning Ahead” on page 13
- “Identifying the Problem” on page 14
- “Seeking a Solution” on page 16

## Planning Ahead

As applications get deployed, undeployed, and redeployed, and as you experiment with different server configuration settings, there may be times when your server gets into a confused or unstable state. In such cases, it is useful to have a previously saved working configuration on which to fall back. This is not problem solving, per se, but rather a way to avoid problems in the first place.

The Enterprise Server `asadmin` command includes a `backup-domain` option that backs up the domain(s) you specify. Use this option to take periodic “snapshots” of your server configuration. Then, if necessary, use the `restore-domain` option to restore one or more domains to a known working state.

Refer to the *Enterprise Server Administration Guide* for complete instructions on using the `asadmin backup-domain` and `restore-domain` options. Briefly, however, for the purposes of this *Troubleshooting Guide*, use the following procedure to backup and restore a server configuration:

## ▼ To back up and restore a server configuration

### 1 Start the Enterprise Server.

```
install_dir/bin/asadmin start-domain domain_name
```

### 2 Stop the domain.

```
install_dir/bin/asadmin stop-domain domain_name
```

### 3 Back up the domain.

```
install_dir/bin/asadmin backup-domain domain_name
```

Backed up directories are stored by default in the *install\_dir/backups* directory.

### 4 Make changes to the Enterprise Server configuration and/or domain(s), as desired.

### 5 If necessary, restore the server and/or domain configuration to the state saved in Step 3, above.

```
install_dir/bin/asadmin restore-domain --filename backup_file domain_name
```

## Identifying the Problem

Java EE application servers are typically deployed in complex and highly sophisticated operating environments. The Sun GlassFish Enterprise Server covers a broad range of technologies, including Java, Java servlets, XML, JSP, JDBC data sources, EJB technology, and more. Other products and tools associated with the Enterprise Server are LDAP, Web Server, , deployment and migration tools, and so on. Understanding and diagnosing complex issues involving so many disparate components requires thorough knowledge and a careful diagnostic process.

Gathering any or all of the following information will make it easier to classify a problem and search for solutions. Note that operating system utilities, such as `pkginfo` and `showrev` on Solaris and `rpm` on Linux, are helpful in gathering system information.

## ▼ Verifying configuration information

- 1 What are the exact version numbers of the operating system and products installed?
- 2 Have any patches been applied? If so, specify product and operating system patch numbers.
- 3 How is the system configured?
- 4 What system resources does the system have (memory, disk, swap space, and so on)?

- 5 How many application servers, web servers, and directory servers are installed?
- 6 How is the web server connected to Enterprise Server? On the same machine or not?
- 7 How is the Enterprise Server connected to the directory server?
- 8 Are application servers in a cluster or not?
- 9 Was any upgrade done? If so, what were source and target versions?
- 10 Was a migration done? If so, what were source and target versions?
- 11 Have any new applications been deployed?
- 12 Is SSL enabled or not?
- 13 What versions of the HADB and the backend database are being used?
- 14 What JDBC driver is being used to access the database?
- 15 What JDK version is being used?
- 16 What are the JVM heap, stack, and garbage collection-related parameters set to?
- 17 What are the JVM options?
- 18 What third-party technologies are being used in the installation?
- 19 **Are the interoperating component versions in compliance with the compatibility matrix specified in the release notes?**

After gathering this information:

- Collect web server error and access log data (web server instance-specific).
- Collect any Enterprise Server stack traces. Note that a fresh set of logs associated with the specific problem should be run. This avoids scanning gigabytes of irrelevant log information.
- Determine the sequence of events that occurred when the problem first appeared, including any steps that may already have been taken to resolve the problem.

## Seeking a Solution

After identifying the problem and formulating a preliminary hypothesis of what may be wrong, you are ready to do some investigation.

The following topics are addressed in this section:

- “Verify System Configuration” on page 16
- “Evaluate Messages” on page 16
- “Examine Log Files” on page 17
- “See if the Problem has been Solved Before” on page 20
- “Search the Product Documentation” on page 20
- “Search the Knowledge Base” on page 20
- “Search or Participate in the Online Forum” on page 21
- “Contact Support” on page 21

### Verify System Configuration

Sometimes the most obvious solutions are overlooked, and so the first step is to verify the system configuration. Refer to the *Sun GlassFish Enterprise Server Release Notes* for the most up-to-date system requirements and dependencies.

### Evaluate Messages

Messages generally include information about the attempted action, the outcome of the action, and, if applicable, the cause of jeopardy or failure.

#### Types of Messages

The log files contain the following general types of message entries:

- **Error** – These messages mark critical failures that cause status to be reported as Failed. Error messages generally provide detailed information about the nature and the cause of the problem that occurred.
- **Warning** – These messages mark non-critical failures. Warning messages generally contain information about the cause and the nature of the failure, and also provide possible remedies.
- **Information** – These messages mark normal completion of particular tasks.

#### Error Messages

A problem is often accompanied by an error message that prevents the application from proceeding.



- In some cases, the message is very clear about what is wrong and what needs to be done, if anything, to fix it. For example, if you start a domain using the `asadmin start-domain` command, then inadvertently issue the same command again after the domain has started, the following message is displayed:

```
userD:\Sun\studio5_se\appserver8\bin>asadmin start-domain
Domain already started : domain1
Domain domain1 Started.
```

In this case, the message gives clear guidance and the problem can be disregarded.

- Sometimes an error message gives only general information about the problem or solution, or suggests multiple possibilities. For example:

```
[16/Jun/2003:22:20:50] SEVERE ( 2204): WEB0200: Configuration error
in web module [JAXBProjectStudio] (while initializing virtual server
[server1]) com.iplanet.ias.config.ConfigException: Failed to load
deployment descriptor for: JAXBProjectStudio cause:
java.io.FileNotFoundException:
```

In this case, the problem is not obvious, or there might be multiple things wrong. You might have to consider various possibilities and perhaps a number of solutions. If the proposed fix is time consuming or costly, take steps to ensure that the fix is likely to be correct before actually doing anything.

- Some error messages are either not helpful or provide little guidance; for example:

```
[23/Jun/2003:16:50:45] WARNING ( 1972): for host 127.0.0.1 trying to GET
/SupplierServiceClient1/SupplierServiceClient1_SOAP.html, send-file
reports: HTTP4144: error sending D:/Sun/studio5_se/appserver8/domains/
domain1/server1/applications/j2ee-modules/SupplierServiceClient1_1/
SupplierServiceClient1_SOAP.html (Overlapped I/O operation is in progress.)
status=1:5
```

In this case, there is very little information to go on. It is especially important to identify the exact situation that caused the error, and what the symptoms are before proceeding.

For descriptions of all the Enterprise Server error messages, refer to [TBDlink](#).

## Examine Log Files

A number of the Enterprise Server subsystems create log files and log their events to these files. The primary purpose of these log files is to provide troubleshooting information.

In addition to the message text, a logged message provides the following information:

- Date and time of the event
- Log level for the event — Enterprise Server-specified log level ID or name

- Process identifier (PID) — PID of the Enterprise Server process
- (optional) Virtual server identifier (VSID) — VSID that generated the message
- Message identifier (MID) — subsystem and a four digit integer
- Message data

The specific logs associated with each Enterprise Server problem area are discussed in the associated chapters of this manual.

## Log Levels

The Enterprise Server has many log levels that can be set in the Administration GUI (FINEST, FINER, FINE, CONFIG, INFO, WARNING, SEVERE, ALERT, and FATAL). All messages are logged when the log level is set to FINEST and only serious error messages appear if the log level is set to FATAL.

Note that the more detailed log levels (FINEST, FINER, FINE) can generate high volumes of log information for certain events, which may make it appear at first glance that there is an error condition when in fact there is not.

All messages with a log level less than the default level of INFO (FINEST, FINER, FINE, and CONFIG) provide information related to debugging and must be specifically enabled. Instructions for doing this are contained in the *Sun GlassFish Enterprise Server Administrator's Guide*.

In addition to the standard JDK log levels, the Enterprise Server has added log levels designed to map more intuitively to the Enterprise Server log file (`server.log`) and to tightly integrate with Solaris. The log levels ALERT and FATAL are specific to the Enterprise Server and are not implemented in the JDK1.4 logging API.

---

**Note** – For information on the event log mechanism used in the Microsoft Windows operating environment, refer to the Windows help system index using the keywords Event Logging. If you choose to send logs to the Windows `server.log` file, only messages with a log level of INFO, WARNING, SEVERE, ALERT, or FATAL are logged to the Windows Event Log.

---

## Log Options

The Administration GUI provides the following two logging options:

- **Option 1** — Log `stdout` (`System.out.print`) content to the event log
- **Option 2** — Log `stderr` (`System.err.print`) content to the event log

When these options are set, `stdout` and `stderr` messages are written to the `server.log` file. (The event log is a `syslog` daemon on Solaris and Event Log on Microsoft Windows.)

If the above options are not set:

- Anything written to `stdout` or `stderr` (that is, using `System.out` or `System.err`) will not appear in the logs.
- Messages logged with the JDK logger will appear in the logs.
- Messages written to `stdout` or `stderr` appear with the `INFO` level, but do not have a message ID.

## Client Side Logging

The Application Client Container (ACC) has its own log service and can only log to a local file. The ACC typically runs in its own process, on a different host from the Enterprise Server. It has its own logging infrastructure and its own log file. The `sun-acc.xml` file contains the ACC configuration. Refer to the *Sun GlassFish Enterprise Server Application Server Developer's Guide to Clients* for more information.

## Obtaining a Thread Dump

This section explains how to obtain a thread dump for Enterprise Server. By default, the server dumps a core file and restarts with the `-Xrs java-option` flag in the `server.xml` file.

### On UNIX

The following procedure describes how to obtain a server thread dump on UNIX.

#### ▼ To obtain a server thread dump on UNIX

- 1 **Verify that the `server.xml` file for the affected server instance does not include the `-Xrs java-option` flag. Remove the `-Xrs java-option` flag if it exists.**
- 2 **If the option is changed, restart the server instance.**
- 3 **Use the `ps` command to determine the `java` and/or `appservDAS` processes under which the application server is running.**
- 4 **Run the following command on the application server instance:**

```
kill -3 pid
```

The `kill` command redirects the thread dump to the `server.log` file for server the instance.

### On Windows

The following procedure describes how to obtain a server thread dump on Windows.

## ▼ To obtain a server thread dump on Windows

- 1 **Verify that the `server.xml` file for your server instance does not include the `-Xrs java-option` flag. Remove the `-Xrs java-option` flag if it exists.**
- 2 **If the option was changed, restart your Enterprise Server.**
- 3 **Type `ctrl-brk` in the Enterprise Server window. The thread dump will be redirected to the `server.log` file for the instance.**

## See if the Problem has been Solved Before

A good initial step is to scan this Troubleshooting Guide to see if the problem is addressed here. If so, select the appropriate solution. Many of the solutions contain references to other documents in the Enterprise Server document collection for additional details, explanations, or examples.

## Search the Product Documentation

Start by reading the Release Notes for the version of the product you are troubleshooting.

The documentation for this Enterprise Server product release is available at TBDlink.

## Search the Knowledge Base

The Knowledge Base is a collection of articles on product issues that provide information helpful for troubleshooting. To access the Knowledge Base:

### ▼ To search the Knowledge Base

- 1 **Go to [SunSolve](http://sunsolve.sun.com) (<http://sunsolve.sun.com>).**
- 2 **Under SunSolve Collections, click the Search Collections link.**
- 3 **Select the checkbox for the collection(s) to search.**
- 4 **Click Next.**
- 5 **Enter the search criteria.**
- 6 **Click Go.**

## Search or Participate in the Online Forum

Browse directly in any of the online forums, or log in and register to start posting messages. The Enterprise Server online forum is available at: <http://forum.java.sun.com/index.jspa>

## Contact Support

When necessary, gather together the information you have acquired and contact technical support at <http://www.sun.com/service/contacting>.



## Common Problems

---

This section covers the most common problems you may encounter when using the Enterprise Server:

- “`asadmin start-domain` Command Fails” on page 24
- “Automatically Restarting a Domain or Node Agent on Reboot” on page 25
- “Cannot Find Log Files” on page 29
- “Accessing Local Server Fails (`http://localhost:8080`)” on page 30
- “Accessing Remote Server Fails” on page 31
- “Cannot Access the Administration Console” on page 32
- “Cannot Access a Server Application” on page 34
- “Server Will Not Start on Windows (Port Conflict)” on page 36
- “Port Conflicts Debugging Multiple Instances on the Same Server” on page 37
- “Two Server Instances Bind to Same Port on Windows” on page 37
- “Error: System cannot find the path specified” on page 38
- “Application Generates Error `persistence.support.JD0DataStoreException`” on page 39
- “Using `asadmin set` Command May Produce Unexpected Results” on page 39
- “Cannot Undeploy Or Redeploy Application With Open Streams to Jar Files (Windows Only)” on page 40
- “Cannot Reinstall Enterprise Server After Manual Deletion of Directories” on page 40
- “Cannot Produce a JVM Thread Dump After Server Crash” on page 41
- “Unable to Start Server with 3.5 GB JVM heap” on page 41
- “Unable to Set Up Cluster on OpenSuse Linux” on page 42
- “Node Agent Does not Sync Corrupt/Deleted `domain.xml`” on page 24

## Node Agent Does not Sync Corrupt/Deleted domain.xml

Node agents started with the command `asadmin start-node-agent --syncinstances=true` does not sync a corrupted `domain.xml` file in the instance's config directory. The `start-node-agent` command fails because it cannot sync the `domain.xml` file in the DAS with the one in the instance's config directory. Manually edited `domain.xml` file are not synchronized because manual edits do not change the timestamp of the file.

### Solution

Perform a dummy operation in the `domain.xml` file that has been manually edited, so that when you run the `start-node-agent` or `start-instance` command, the sync operation of `domain.xml` succeeds.

## asadmin start-domain Command Fails

The command `asadmin start-domain` fails with one of the following errors:

- “Error: CLI143 There is more than one domain...” on page 24
- “Error: Could Not Start Domain” on page 25

### Error: CLI143 There is more than one domain...

#### Description

When issued with no arguments, the command `asadmin start-domain` fails with the error:

```
CLI143 There is more than one domain in C:\Sun\AppServer\domains.  
Please use operand to specify the domain.  
CLI156 Could not start the domain null.
```

This error occurs when there is more than one domain in the domains directory, none of them is named `domain1`, and no domain is specified with the `start-domain` command.

#### Solution

Specify the domain when issuing the `start-domain` command:

```
asadmin start-domain domain1
```



## Error: Could Not Start Domain

### Description

This message comes from Enterprise Server. The full message looks like either:

```
Could not start the domain.  
There are no domains.
```

*or*

```
Could not start the domain.  
No default domain. Need to enter a domain.
```

This error occurs when Enterprise Server 8 is installed on the same system, and its `asadmin` command (at `/usr/sbin`) is found on the PATH before the `asadmin` command for Enterprise Server at `install_dir/bin`. The situation is especially likely to occur on Solaris/Linux systems when `.` is not specified as part of the PATH variable. Without `.` in the PATH, the `asadmin` command in `/usr/sbin` is found first, even when the current directory is `install_dir/bin`.

### Solution

Make sure `install_dir/bin` is in the PATH ahead of `/usr/sbin`, or that `.` is in the PATH ahead of `/usr/sbin` if you access `asadmin` by changing directories to `install_dir/bin`. Alternatively, if you do change to `install_dir/bin` to run `asadmin`, be sure to include `./` in the command name; for example:

```
cd install_dir/bin  
./asadmin
```

## Automatically Restarting a Domain or Node Agent on Reboot

If your domain or node agent is stopped unexpectedly (for example, if you need to restart your machine), you can configure your system to automatically restart the domain or node agent on reboot.

- “Restarting Automatically on UNIX Platforms” on page 26
- “Restarting Automatically on the Microsoft Windows Platform” on page 26

## Restarting Automatically on UNIX Platforms

### Restarting a Domain

To restart your domain on a UNIX platform, add a line containing the appropriate `asadmin start-domain` command to your `/etc/inittab` file. If you use `/etc/rc.local`, or your system's equivalent, place a the desired `asadmin` command in `/etc/rc.local`.

For example, to restart `domain1` for an Troubleshooting Guide installed in the `/opt/SUNWappserver` directory, using a password file called `password.txt`, add the following line to `/etc/inittab` or `/etc/rc.local`:

```
das:3:respawn:/opt/SUNWappserver/bin/asadmin start-domain --user admin
--passwordfile /opt/SUNWappserver/password.txt domain1
```

Be sure to put the text on one line. The first three letters are a unique designator for the process and can be altered.

### Restarting a Node Agent

To restart a node agent, the syntax is similar. For example, to restart `agent1` for an Troubleshooting Guide installed in the `/opt/SUNWappserver` directory using a password file called `password.txt`, add the following line to `/etc/inittab` or `/etc/rc.local`:

```
das:3:respawn:/opt/SUNWappserver/bin/asadmin start-node-agent --user admin
--passwordfile /opt/SUNWappserver/password.txt agent1
```

## Restarting Automatically on the Microsoft Windows Platform

To restart automatically on Microsoft Windows, create a Windows Service. Use the `appservService.exe` and `appservAgentService.exe` executables shipped with Sun GlassFish Enterprise Server in conjunction with the Service Control command (`sc.exe`) provided by Microsoft.

- The `sc.exe` command comes with Windows XP and is either located in the `C:\windows\system32` directory or `C:\winnt\system32` directory.
- As of this writing, the Windows 2000 `sc.exe` is available for download at: <ftp://ftp.microsoft.com/reskit/win2000/sc.zip>. For more information on using `sc.exe`, see:

[http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dndllpro/html/msdn\\_scmslite.asp](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dndllpro/html/msdn_scmslite.asp).

Use `appservService.exe` and `appservAgentService.exe` as follows:

```
C:\winnt\system32\sc.exe create service_name binPath= \\fully_qualified_
path_to_appservService.exe \\fully_qualified_path_to_asadmin.bat
start_command\"
\\fully_qualified_path_to_asadmin.bat stop_command\" start= auto
DisplayName= "display_name"
```

## Starting and Stopping a Domain

To create a service called SunJavaSystemAppServer DOMAIN1 that starts and stops domain1 using password file C:\Sun\AppServer\password.txt, run the following command:

```
C:\windows\system32\sc.exe create domain1 binPath= "C:\Sun\AppServer\
lib\appservService.exe \\C:\Sun\AppServer\bin\asadmin.bat start-domain
--user admin --passwordfile C:\Sun\AppServer\password.txt domain1\"
\\C:\Sun\AppServer\bin\asadmin.bat stop-domain domain1\"
start=auto DisplayName= "SunJavaSystemAppServer DOMAIN1"
```

## Starting and Stopping a Node Agent

To create a service that starts and stops the node agent agent1, run the following command:

```
C:\windows\system32\sc.exe create agent1 binPath= "C:\Sun\AppServer\
lib\appservAgentService.exe \\C:\Sun\AppServer\bin\asadmin.bat
start-node-agent --user admin --passwordfile C:\Sun\AppServer\
password.txt agent1\" \\C:\Sun\AppServer\bin\asadmin.bat
stop-node-agent agent1\" start=auto DisplayName="SJESAS_SE8.1 AGENT1"
```

---

**Note** – The start and stop commands entered as part of the binPath= parameter must have the correct syntax. To test, run the commands from the command prompt. If the commands do not properly start or stop the domain or node agent, the service will not work correctly.

Also, do not use a mixture of asadmin start and stop commands and service start and stops. Mixing the two can cause the server status to be out of sync. For example, the service might not show that the component has started even though the component is not running. To avoid this situation, always use the sc.exe command to start and stop the component when using services.

---

## Security for Automatic Restarts

Handle the password and master password required when starting in one of the following ways:

- On Microsoft Windows, configure the service to ask the user for the password.
  - In the Services Control Panel, double-click the service you created.
  - In the Properties window, click the Log On tab.
  - Check “Allow service to interact with desktop” to prompt for the required passwords when starting the component.

You have to log in to see the prompts, and entries are not echoed back as you type them. This method is the most secure way to use the services option, but user interaction is required before the service becomes available.

If the “interact with desktop” option is not set, the service stays in a “start-pending” state and appears to hang. Kill the service process to recover from this state.

- On Windows or UNIX, create a domain using the `--savemasterpassword=true` option and create a password file to store the admin password. When starting the component, use the `--passwordfile` option to point to the file that contains the password. The admin password can also be added by using the `--password` option with the `asadmin start` command. Be aware that this method is less secure because the admin password is stored in clear text.

For example:

- Create a domain with a saved master password. Using this syntax, you are prompted for the admin password and master password:

```
asadmin create-domain --adminport 4848 --adminuser admin
--savemasterpassword=true --instanceport 8080 domain1
```

- On Windows, create a service using a password file to populate the admin password:

```
C:\windows\system32\sc.exe create domain1 binPath=
"C:\Sun\AppServer\lib\appservService.exe
"C:\Sun\AppServer\bin\asadmin.bat start-domain --user admin
--passwordfile C:\Sun\AppServer\password.txt domain1\"
"C:\Sun\AppServer\bin\asadmin.bat stop-domain domain1\"
start= auto DisplayName= "SJESAS_PE8.1 DOMAIN1"
```

The path to the password file `password.txt` is `C:\Sun\AppServer\password.txt`. It contains the password in the following format:

```
AS_ADMIN_password=password
```

For example, for a password `adminadmin`:

```
AS_ADMIN_password=adminadmin
```

- On UNIX, use the `--passwordfile` option in the line you add to the `inittab` file:

```
das:3:respawn:/opt/SUNWappserver/bin/asadmin start-domain
--user admin --passwordfile /opt/SUNWappserver/password.txt domain1
```

The path to the password file `password.txt` is `/opt/SUNWappserver/password.txt`. It contains the password in the following format:

```
AS_ADMIN_password=password
```

For example, for a password `adminadmin`:

```
AS_ADMIN_password=adminadmin
```

- Creating a service using a password that is populated from a command line option:

```
C:\windows\system32\sc.exe create domain1 binPath= "C:\Sun\AppServer\lib\appservService.exe \"C:\Sun\AppServer\bin\asadmin.bat start-domain --user admin --password adminadmin domain1\" \"C:\Sun\AppServer\bin\asadmin.bat stop-domain domain1\" start=auto
DisplayName="SJESAS_PE8.1 DOMAIN1"
```

## Cannot Find Log Files

The following Enterprise Server logs are useful for troubleshooting installation problems:

- **Server log file** — For troubleshooting server configuration and deployment problems
- **HTTP server access logs** — For troubleshooting HTTP server problems, and for tracing the activity of HTTP requests entering the Enterprise Server instances

Both the installation and uninstallation programs create log files and log all installation and uninstallation events to these files. The primary purpose of these log files is to provide troubleshooting information.

In addition to installation program messages and log files, operating system utilities such as `pkginfo` and `showrev` on Solaris and `rpm` on Linux can be used to gather system information.

Log file entries include information about the attempted action, the outcome of the action, and, if applicable, the cause of failure. The log files contain the following types of message entries:

- **INFO** — These messages mark normal completion of a particular installation tasks.
- **WARNING** — These messages mark non-critical failures. Warning messages generally contain information about the cause and the nature of the failure, and also provide possible remedies.
- **ERROR** — These messages mark critical failures that cause installation or uninstallation status to be reported as Failed. Error messages generally provide detailed information about the nature and the cause of the problem that occurred.

The domain-specific logs are located in `install_dir/domains/domain1/logs/`. Log files for the server installation in general are located as follows:

- **Solaris, root user installation/uninstallation:**

```
/var/sadm/install/logs
```

- **Solaris, non-root installation/uninstallation:**

```
/var/tmp
```

- **Linux installation/uninstallation:**

`/var/tmp`

## Accessing Local Server Fails (<http://localhost:8080>)

Things to check for this error include the following:

- “Did the Server Start?” on page 30
- “Was the Server Started at the Expected Port?” on page 30

### Did the Server Start?

#### Description

If the console window is still open, the expected message is:

```
Domain domain Started
```

where *domain* is the name of the default domain. This indicates that the default domain was started successfully.

If the console window is already closed, check for messages in the log file:

```
install_dir/domains/domain1/logs/server.log
```

If startup was successful, the expected message is similar to that on the console, and appears at the end of the log file:

```
[INFO][...][...][date&time][Application server startup complete .]
```

### Was the Server Started at the Expected Port?

#### Description

The server might be running at a different port number than expected, either because it was intentionally installed there, or because another server was already running on the default port when the server was installed.

#### ▼ To determine which port number the server is actually using

##### 1 Examine the server's configuration file:

```
install_dir/domains/domain1/config/domain.xml
```

##### 2 Find the `http-listener` element.

### 3 Inspect the value of the port attribute.

Be sure to enter the correct port number when invoking the server.

---

**Note** – The server’s default port number is 8080, however, there are a number of ways in which the expected value can change:

- A different port number was specified during installation.
  - A previous installation exists.
  - If the specified port number is already taken by another application when the server is started, the port number rolls forward to the next available number. For example, if a server is already running on the default 8080 port, the new Application Server instance uses port number 8081. If two servers are running, the port number rolls to 8082, and so on.
- 

## Accessing Remote Server Fails

When attempting to open the start page of the Application Server, the initial screen does not appear.

Things to check include the following:

- “[Is the Server Available Locally?](#)” on page 31
- “[Is the Proxy Setting Causing a Problem?](#)” on page 32

## Is the Server Available Locally?

### Description

If the server cannot be accessed from the web, but it is running locally, then the server is actually running.

### Solution

Verify that the server is running locally.

### ▼ To verify that the server is running locally

- 1 Log on to the machine where the server is running.
- 2 Go to the local web page. For example, if 8080 is the default port, go to:

`http://localhost:8080/`

If the start page does appear, there is a problem with the web connection that prevents accessing the server remotely. If the start page does not appear, see “[Did the Server Start?](#)” on page 30

## Is the Proxy Setting Causing a Problem?

### Description

The server should be accessible directly from the host on which it is running (`localhost`); for example, using the default port, `8080`:

```
http://localhost:8080/
```

### Solution

A server instance running on `localhost` may not be accessible if the server host machine is connected to the web through a proxy. To solve this problem, do one of the following:

- Set the browser to bypass the proxy server when accessing `localhost`. Refer to the browser's help system for information on how to do this.
- Use the fully-qualified host name or IP address of your system; for example:

```
http://myhost.mydomain.com:8080/
```

---

**Note** – To find the host name and domain for the `localhost` machine:

- **On Microsoft Windows** — On the desktop, right-click My Computer and select Properties from the pop-up menu. A System Properties dialog is displayed. Click Network Identification to see the computer name.
  - **On Solaris or Linux** — Type `hostname` at the command prompt.
- 

## Cannot Access the Administration Console

The Administration Console provides an interface for administrative functions. If the Administration Console is not accessible, it might be for one of several reasons.

- [“Is the Application Server Running?” on page 32](#)
- [“Is the Administration Console Running on the Expected Port?” on page 33](#)
- [“Is the Security Manager Disabled?” on page 33](#)

## Is the Application Server Running?

### Description

The server must be running before the Administration Console can be accessed.



## Solution

Review the information in [“Did the Server Start?”](#) on page 30 to determine if the server is running.

## Is the Administration Console Running on the Expected Port?

### Description

The default port number for the EE and SE Administration Console is 4849; for the PE Administration Console it is 4848. Also note that the URL for the EE and SE console requires secure HTTP (`https://servername:4849.`), whereas the PE console uses standard HTTP (`http://servername:4848`). However, it could be running on a different port number than expected, either because it was intentionally installed there, or because that port was taken when the server was started.

### Solution

Refer to [“Was the Server Started at the Expected Port?”](#) on page 30 for guidelines on verifying the port on which the Administration Console is running, and be sure to enter the correct port number and HTTP protocol when invoking the Administration Console.

## Is the Security Manager Disabled?

### Description

Since there is no configuration interface in the Enterprise Server for disabling the Security Manager, it can only be disabled by directly modifying the `domain.xml` configuration file in such a way that the following line is removed:

```
<jvm-option>-Djava.security.policy=yourPolicy</jvm-option>
```

### Solution

The `-Djava.security.policy=yourPolicy` option must be present in the `domain.xml` file to access the Administration Console.

## Cannot Access a Server Application

If a particular application cannot be accessed through the Enterprise Server, some things to check include the following:

- [“Is the Application Server Running?”](#) on page 34
- [“Was Application Deployment Successful?”](#) on page 34
- [“Invalid User or Password When Using Don’t Prompt Option”](#) on page 34
- [“Administrator User Name or Password Not Known”](#) on page 35

### Is the Application Server Running?

#### Description

If the Enterprise Server is not running, applications will not be accessible.

#### Solution

Review the information in [“Did the Server Start?”](#) on page 30 to determine if the server is running. The server must be running before a server application can be accessed.

### Was Application Deployment Successful?

#### Description

An application must be successfully deployed before it can be accessed.

#### Solution

Check the server's log file:

```
install_dir/domains/domain1/server.log
```

### Invalid User or Password When Using Don’t Prompt Option

You are getting the error, Invalid user or password, but you installed the system with the Don’t Prompt option, so the password should be supplied automatically.

#### Description

The correct password may not have been specified during installation, or it may not be passed when the domain is started.

## Solution

Check the password in the `.asadminprefs` file. On UNIX/Linux systems, it is in the home directory for the user under which the server was installed. On Windows, it is in `C:\Documents and Settings\username`. The contents look something like this:

```
AS_ADMIN_USER=admin
AS_ADMIN_PASSWORD=administrator
```

## Administrator User Name or Password Not Known

If you have forgotten the administrator user name, you can find it by inspecting the `.adminprefs` file, as described in the section above, or by inspecting the `install_dir/domains/domain1/config/keyfile`, where `domain1` is the default domain. For a different domain, substitute its name in the path.

If you have forgotten the administrator password, you will need to create a new user name-password pair by removing the user name and password, creating new ones, and restarting the server. (You will not be able to read the password, because it is encrypted in the keyfile.)

### ▼ To remove the user name and password completely

1 Stop the server, if it is currently running.

2 Change to the appropriate `WEB-INF` directory; for example:

```
install_dir/lib/install/applications/adminapp/adminapp_war/WEB-INF
```

3 Comment out the entire `<security-constraint>` element in the `web.xml` file.

Do not delete the element, as you will be reenabling it later. This action disables security for command-line operations.

---

**Note** – The commands will still expect a value for `--username` (or `-u`) and `--password` (or `-w`). But these can be dummy values, since the server side does not impose any security.

---

4 Start the server.

At this point, the server does not have command-line security.

5 Run the following command:

```
asadmin create-file-user --user <dummy> --password <dummy>
--userpassword <new_secret> --groups asadmin <new_user_id>
```

This command creates the following new entry:

```
<install_dir\>/domains/domain1/config/keyfile
```

- 6 **Uncomment the `<security-constraint\>` element in `web.xml` file.**
- 7 **Restart the server for the new user name-password to take effect.**

---

**Note** – When the server is started, any remote command-line operations will need `new_user_id` and `new_secret` as user name and password.

---

## Server Will Not Start on Windows (Port Conflict)

If a message similar to the following is displayed when starting the Enterprise Server on Microsoft Windows, a server port conflict has occurred:

```
Address already in use
```

This error occurs when another application is running on the Application Server port (default 8080), or because a previous instance of the Enterprise Server did not shut down cleanly.

Other things to check include the following:

- [“Is Another Application Running on the Server's Port?”](#) on page 36
- [“Has an Ungraceful Shutdown Occurred on a Previously Running Server?”](#) on page 36

## Is Another Application Running on the Server's Port?

If another application is using the server's port, stop the other application, then restart the Application Server.

---

**Note** – The installer attempts to avoid port conflicts by choosing the next available port when the default port is in use—but that only works if application using the default port was running when the Enterprise Server was installed.

---

## Has an Ungraceful Shutdown Occurred on a Previously Running Server?

Use the `asadmin stop-domain` command to stop the server, or explicitly kill the Java process and then restart the Enterprise Server.

# Port Conflicts Debugging Multiple Instances on the Same Server

## Description

Port conflict errors can occur when debugging multiple instances on the same server that are part of the same cluster.

## Solution

Modify the `domain.xml` file to remove the `address` attribute from the `-Xrunjdpw` option in the `java-config` element for the cluster. This results in the JVM choosing a random debug port for the instance. The port number chosen for the instance is displayed in the server log when it is started; for example:

- **Before:**

```
debug-options="-Xdebug -Xrunjdpw:transport=dt_socket,server=y,
suspend=n,address=9009"
```

- **After:**

```
debug-options="-Xdebug -Xrunjdpw:transport=dt_socket,server=y,
suspend=n"
```

# Two Server Instances Bind to Same Port on Windows

## Description

This problem only occurs on Windows 2000/XP systems with the Enterprise Server software, and is due to a known Windows security flaw rather than a problem with the Enterprise Server itself.

The problem occurs when two or more instances of the Enterprise Server are created using the same port number for the `instanceport` option; for example:

```
asadmin create-domain -adminport 5001 <options> -instanceport 6001
<domain>
asadmin create-domain -adminport 5002 <options> -instanceport 6001
<domain>
```

When the two domains are started on a UNIX/Linux system, a port conflict error is thrown and the second instance fails to start. However, when the two domains are started on Windows 2000/XP, no error is thrown, both server instances start, but only the first instance is accessible at the specified port. When that first server instance is subsequently shut down, the second instance then becomes accessible. Moreover, when both instance are running, the Windows `netstat` command shows the duplicate listeners as active, but only the first listener can respond to requests.

## Solution

Be sure to use unique port numbers for all server instances on Windows systems.

## Error: System cannot find the path specified

### Description

This error message occurs when attempting to start the server after deleting the J2SE directory that was specified during installation. This situation generally occurs after being informed during the install that the J2SE platform needs an upgrade, and the upgrade takes place after the Enterprise Server installation.

### Solution 1

To use the new J2SE for all domains, change the `AS_JAVA` variable in `asenv.conf` (Solaris/Linux), or `asenv.bat` (Windows).

### Solution 2

The J2SE version can be changed on a per-domain basis by modifying the `java-home` attribute for the `java-config` element in the domain's `domain.xml` file.

```
<java-config ...  
  java-home="path"  
  ... \>
```

### Solution 3

A more time-intensive solution is to uninstall and then reinstall the server.

## Application Generates Error persistence.support.JDODataStoreException

### JDODataStoreException

#### Description

A `com.sun.jdo.api.persistence.support.JDODataStoreException` is generated by an application, with a nested `java.sql.SQLException` indicating a duplicate primary key.

Even if the application is checking for a `CreateException`, it does not see one. The Enterprise JavaBeans specification requires a `CreateException` to be thrown only if two beans with the same primary key are created in the same transaction, so a `CreateException` is not thrown on transaction rollback if two entity beans with duplicate primary keys are created.

#### Solution

If an application creates an entity bean with a duplicate primary key, check to see if the primary key exists by calling `findByPrimaryKey` before calling `create`.

## Using `asadmin set` Command May Produce Unexpected Results

#### Description

Unexpected results are returned when setting variables in a command, such as:

```
asadmin set name=${a-b}
```

In this case, `name` is set to `b`, not `{a-b}` because the shell syntax `${a=b}` is interpreted as “if the variable `a` is unset, substitute the value `b`, otherwise substitute the value of `a`.” This is standard shell behavior. For example, consider the following:

```
asadmin set default-config.http-service.http-listener.http-listener-1.port=
${http-listener-1-port}
```

In this case, `default-config.http-service.http-listener.http-listener-1.port` is set to `listener-1-port`, which is invalid.

## Cannot Undeploy Or Redeploy Application With Open Streams to Jar Files (Windows Only)

### Description

On Windows systems, after running an application, subsequent attempts to undeploy it or redeploy it throw exceptions about the server being unable to delete a file or rename a directory.

On Windows systems, an application may use `getClass().getResource` or `getResourceAsStream` methods to locate a resource inside the application, particularly in jar files that are in the application or accessible to it. If the streams remain open, subsequent attempts to redeploy or undeploy the application can fail. In addition, the Java runtime by default caches streams to jar files for performance reasons.

### Solution

Be sure to close streams opened by your applications. Also, if an application needs to be redeployed or undeployed repeatedly, and also needs to obtain a resource from a jar file using `getResource` or `getResourceAsStream`, consider using `getClass().getResource` which returns a URL object, then invoke the `url.setUseCaches` method to turn off caching for that jar file, and then use `url.getInputStream()` to obtain the stream.

Although turning off caching for accesses to the jar file can slow down performance, this approach does allow the application to be undeployed or redeployed. Note also that if the `getClass().getResourceAsStream` method is used instead, then the jar file in which the resource is located will be cached (this is the default Java runtime setting) and remain open until the server is stopped.

## Cannot Reinstall Enterprise Server After Manual Deletion of Directories

### Description

If the Enterprise Server directories are deleted manually rather than by means of the included `uninstall` the program, subsequent attempts to reinstall the Enterprise Server in the same directory fail. This is because the installation directory information stored in `/tmp/product registry` file remains even though the program directories have been removed.



## Solution 1

Remove Enterprise Server directory information from the `<location>` property entries in the `/tmp/productregistry` file; for example, change:

```
<location>/opt/SUNWappserver/jdk</location>
```

to

```
<location></location>
```

## Solution 2

Reinstall the Enterprise Server in a different directory.

# Cannot Produce a JVM Thread Dump After Server Crash

## Description

If the Enterprise Server crashes, the server dumps a core file and, by default, restarts with the `-Xrs` flag, which prevents the dump of a JVM thread dump.

## Solution

### ▼ To enable a JVM thread dump

- 1 Comment out the `-Xrs` flag in the `server.xml` file for the Enterprise Server.
- 2 Kill the server process (`kill -3` on UNIX; `Ctrl+Break` on Windows).

# Unable to Start Server with 3.5 GB JVM heap

If you are using a 3.5 GB JVM heap, you may intermittently see one of the following errors when you try to start the server:

- `OutOfMemory` exception
- Server hangs

Consider reducing the thread stack size to 128 KB (`-Xss128k`).

## Unable to Set Up Cluster on OpenSuse Linux

The following commands fail when you are running on OpenSuse Linux (10.3):

- `asadmin create-domain`
- `asadmin create-instance`
- `asadmin node-agent`

According to this bug report ([https://bugzilla.novell.com/show\\_bug.cgi?id=331680](https://bugzilla.novell.com/show_bug.cgi?id=331680)), the JDK distributed with OpenSuse causes problems with JXTA (which is the underlying technology of clustering). One way to solve this problem would be to change your JVM.

### ▼ Changing your JVM installation

- 1 Download and install the latest JVM.
- 2 **Modify `install-dir/config/asenv.conf` and change the value of the `AS_JAVA` environment variable to point to the new JVM location.**

# HADB Problems

---

This section covers problems you may encounter when using the Enterprise Server product with the High Availability Database (HADB) module, HADB Management Client, or the load balancer plugin. These components can be installed separately or with the rest of the Enterprise Server.

Topics in this chapter include:

- “HADB Database Creation Fails” on page 43
- “Server Responds Slowly After Idle Period” on page 47
- “Requests Are Not Succeeding” on page 48
- “Session Persistence Problems” on page 49
- “HADB Performance Problems” on page 51
- “High Load Problems” on page 54
- “Connection Problem Caused by Lack of Semaphore Resources” on page 57
- “Improving CPU Utilization” on page 57
- “HADB Administration Problems” on page 58
- “Shared Memory Segment Key Already Exists (Windows only)” on page 65
- “Failure in `configure-ha-cluster`” on page 65
- “Unable to Run `configure-ha-cluster`” on page 66
- “`hadbm set` Command Fails” on page 66
- “Failure in `configure-ha-cluster`: Creating an HADB Instance Fails” on page 67

## HADB Database Creation Fails

The database creation may fail due to the following reasons:

```
failed to start database : HADB Database creation failed
```

To determine the cause of the problem, use the Log Viewer and/or inspect the `install_dir/hadb/4/log` directory. Some possible errors are:

- “Problems Related to Shared Memory” on page 44

- “Too Few Semaphores” on page 45
- “Database Nodes Cannot Be Reached and the Database Does Not Function” on page 46
- “The Management Agents Could Not Establish a Domain” on page 46
- “hadbm create or hadbm addnodes Command Hangs” on page 47
- “ma (Management Agent Process) Crashes” on page 47

## Problems Related to Shared Memory

### Description

This problem may occur due to any of the following reasons:

#### Cause 1

Shared memory is not configured or the configuration is not working.

#### Solution 1

Follow the instructions described in the *Sun GlassFish Enterprise Server Installation Guide*. Remember to reboot the system after configuring shared memory settings.

#### Cause 2

The physical memory is not enough to satisfy the node requirements. You may see the following error message:

```
HADB-S-05512: Attaching shared memory segment with key <xx\> failed,  
OS status=12 OS message: Not enough space.
```

#### Solution 2

Verify that shared memory is configured and the configuration is working, as mentioned above.

For production systems, reduce the number of nodes on the host or increase the physical memory on the host.

For test/development systems, reduce the shared memory usage by setting the `LogBufferSize` and `DataBufferPoolSize` to a value lower than the default values of 48 and 200MB, respectively. The allowed minimum for these variables are 32 and 64MB, respectively.

#### Cause 3

The size of a shared memory segment has exceeded the allowed maximum size.

```
HADB-S-05510: Getting shared memory segment with key <xx\> failed,  
OS status=22. OS message: Invalid argument.
```

### Solution 3

Verify that shared memory is configured and the configuration is working, as mentioned above.

If shared memory is configured correctly, check whether you have specified any shared memory segment size (`LogBufferSize` or `DataBufferPoolSize`) larger than the system-configured maximum value set in the operating system configuration files (`shmsys:shminfo_shmmax` in `/etc/system` on Solaris).

### Cause 4

There is already a shared memory segment created with the specified identifier:

```
HADB-S-05515: Shared memory segment with key <segment_key\> exists already.
```

### Solution 4

List the shared memory segments and check. The `ipcs` can be used to list the segments in UNIX. Windows uses memory mapped files for shared memory. HADB uses the `getTempPath` system call to get the system-defined temporary directory where these files, named as `f_segmentid`, are stored.

Check whether there is already another running database or any other program using the shared memory segment with this identifier. If so, create a database with another port base. If there are no running databases or other programs using this segment, free the segment with `hadbm delete unused_database`.

Check whether the segments are freed. If they are still there, remove them (use `ipcrm` in UNIX and `delete $TMP/f_*` in Windows). The file name consists of the `f_` prefix followed by the `segment_key` translated into hexadecimal. For example, if the error message indicates that segment key `15201` still exists, the temp file would be named `f_3B61`.

## Too Few Semaphores

### Description

```
HADB-E-05521: Operation on semaphore with key "46025" failed, OS status=28 :  
No space left on device
```

This can be caused when the number of semaphores is too low. Since the semaphores are provided as a global resource by the operating system, the configuration depends on all processes running on the host, not only the HADB. This can occur either while starting the HADB, or during runtime.

## Solution

Configure the semaphore settings by editing the `/etc/system` file. Instructions and guidelines are contained in the Configuring Shared Memory and Semaphores section of the Preparing for HADB Setup chapter of the *Sun GlassFish Enterprise Server Installation Guide*.

## Database Nodes Cannot Be Reached and the Database Does Not Function

### Solution

The IP addresses of the involved hosts should be fixed. HADB uses the fixed IP addresses present at database creation, so you cannot use dynamic IP addresses (DHCP) for production systems.

## The Management Agents Could Not Establish a Domain

### Description

The HADB management system is dependent on UDP Multicast messages on multicast address 228.8.8.8. If these messages cannot get through, the `createdomain` command fails with the following message:

```
The management agents could not establish a domain, please check that the
hosts can communicate with UDP multicast.
```

Possible causes include:

- The agents are running on hosts with several network interfaces on different subnets.
- There is a switch on the network that does not forward multicast messages.
- There is router on the network that does not route multicast messages with the address 228.8.8.8.
- Multicast messages are disabled in the operating system (for example, on Solaris 10).

### Solution 1

If the hosts have several network interfaces on different subnets, the management agent must be configured to use one of the subnets. Set the `ma.server.internal.interfaces` attribute.

### Solution 2

Configure the needed network infrastructure to support multicast messages.

## hadbm create or hadbm addnodes **Command Hangs**

### Description

Some hosts in the host list given to `hadbm create` or `addnodes` have multiple network interfaces, while others have only one, and the `hadbm create/addnodes` command hangs.

### Solution

For the hosts having multiple network interfaces, specify the dotted IP address of the network interface (for example, `129.241.111.23`) to be used by `hadb` when issuing `hadbm create/addnodes`. If the host name is used instead of IP address, the first interface registered on the host will be used, and there is no guarantee that the nodes will be able to communicate.

## ma (Management Agent Process) Crashes

### Description

The `ma` (Management Agent process) crashes for various reasons.

### Solution

Display diagnostic information by using `hadbm listdomain`. Typically, the remedy is to restart the failed agent. If that does not help, restart all agents in turn.

## Server Responds Slowly After Idle Period

### Description

The server takes a long time to service a request after a long period of idleness, and the sever log shows “lost connection” messages of the form:

```
java.io.IOException: ..HA Store: Lost connection to the server.
```

In such cases, the server needs to recreate the JDBC pool for HADB.

### Solution

Change the timeout value. The default HADB connection timeout value is `1800` seconds. If the application server does not send any request over a JDBC connection during this period, HADB closes the connection, and the application server needs to re-establish it. To change the timeout value, use the `hadbm set SessionTimeout=` command.

---

**Note** – Make sure the HADB connection timeout is greater than the JDBC connection pool timeout. If the JDBC connection timeout is more than the HADB connection time out, the connection will be closed from the HADB side, but will remain in the application server connection pool. When the application then tries to use the connection, the application server will have to recreate the connection, which incurs significant overhead.

---

## Requests Are Not Succeeding

The following problems are addressed in this section:

- [“Is the Load Balancer Timeout Correct?” on page 48](#)
- [“Are the System Clocks Synchronized?” on page 48](#)
- [“Is the Enterprise Server Communicating With HADB?” on page 49](#)

### Is the Load Balancer Timeout Correct?

#### Description

When configuring the `response-timeout-in-seconds` property in the `loadbalancer.xml` file, you must take into account the maximum timeouts for all the applications that are running. If the response timeout is set to a very low value, numerous in-flight requests will fail because the load balancer will not wait long enough for the Enterprise Server to respond to the request.

Conversely, setting the response timeout to an inordinately large value will result in requests being queued to an instance that has stopped responding, resulting in numerous failed requests.

#### Solution

Set the `response-timeout-in-seconds` value to the maximum response time of all the applications.

### Are the System Clocks Synchronized?

#### Description

When a session is stored in HADB, it includes some time information, including the last time the session was accessed and the last time it was modified. If the clocks are not synchronized, then when an instance fails and another instance takes over (on another machine), that instance may think the session was expired when it was not, or worse yet, that the session was last accessed in the future!



---

**Note** – In a non-colocated configuration, it is important to synchronize the clocks on that machines that are hosting HADB nodes. For more information, see the Installation Guide chapter, “Preparing for HADB Setup.”

---

## Solution

Verify that clocks are synchronized for all systems in the cluster.

## Is the Enterprise Server Communicating With HADB?

### Description

HADB may be created and running, but if the persistence store has not yet been created, the Enterprise Server will not be able to communicate with the HADB. This situation is accompanied by the following message:

```
WARNING (7715): ConnectionUtilgetConnectionsFromPool failed using
connection URL: connection URL
```

### Solution

Create the session store in the HADB with a command like the following:

```
asadmin create-session-store --storeurl connection URL --storeuser haadmin
--storepassword hapasswd --dbssystempassword super123
```

## Session Persistence Problems

The following problems are addressed in this section:

- [“The create-session-store Command Failed”](#) on page 49
- [“Configuring Instance-Level Session Persistence Did Not Work”](#) on page 50
- [“Session Data Seems To Be Corrupted”](#) on page 50

## The create-session-store Command Failed

### Description

The `asadmin create-session-store` command cannot run across firewalls. Therefore, for the `create-session-store` command to work, the application server instance and the HADB must be on the same side of a firewall.

The `create-session-store` command communicates with the HADB and not with the application server instance.

## Solution

Locate the HADB and the application server instance on the same side of a firewall.

## Configuring Instance-Level Session Persistence Did Not Work

The application-level session persistence configuration always takes precedence over instance-level session persistence configuration. Even if you change the instance-level session persistence configuration after an application has been deployed, the settings for the application still override the settings for the application server instance.

## Session Data Seems To Be Corrupted

### Description

Session data may be corrupted if the system log reports errors under the following circumstances:

- During session persistence
- When the session state is read during session activation
- When the session state is read after session failover

If the data has been corrupted, there are three possible solutions for bringing the session store back to a consistent state, as described below.

### Solution 1

Use the `asadmin clear-session-store` command to clear the session store.

### Solution 2

If clearing the session store does not work, reinitialize the data space on all the nodes and clear the data in the HADB using the `hadbm clear` command.

### Solution 3

If clearing the HADB does not work, delete and then recreate the database.

For solutions 2 and 3, above, after clearing the HADB, recreate the session store to reestablish the database schema.

# HADB Performance Problems

Performance is affected when the transactions to HADB get delayed or aborted. This situation is generally caused by a shortage of system resources. Any wait beyond five seconds causes the transactions to abort. Any node failures also cause the active transaction on that node at crash time to abort. Any double failures (failure of both mirrors) will make the HADB unavailable. The causes of the failures can generally be found in the HADB history files.

To isolate the problem, consider the following:

- [“Is There a Shortage of CPU or Memory Resources, or Too Much Swapping?”](#) on page 51
- [“Is There Disk Contention?”](#) on page 52
- [“Is There a Shortage of HADB Data Device Space?”](#) on page 53
- [“Is There a Shortage of Other HADB Resources?”](#) on page 53

## Is There a Shortage of CPU or Memory Resources, or Too Much Swapping?

### Description

Node restarts or double failures due to “Process blocked for  $x$  sec, max block time is 2.500000 sec.” In this case,  $x$  is the length of time the process was blocked, and it was greater than 2.5 seconds.

The HADB Node Supervisor Process (NSUP/`clu_nsup_srv`) tracks the time elapsed since the last time it did some monitoring work. If that time duration exceeds a specified maximum (2500ms by default), NSUP concludes that it was blocked too long and restarts the node.

NSUP being blocked for more than 2.5 seconds cause the node to restart. If mirror nodes are placed on the same host, the likelihood of double failure is high. Simultaneous occurrence of the blocking on the mirror hosts may also lead to double failures.

The situation is especially likely to arise when there are other processes—for example, in a colocated configuration—in the system that compete for CPU, or memory which produces extensive swapping and multiple page faults as processes are rescheduled.

NSUP being blocked can also be caused by negative system clock adjustments.

### Solution

Ensure that HADB nodes get enough system resources. Ensure also that the time synchronization daemon does not make large (not higher than 2 seconds) jumps.

## Is There Disk Contention?

### Description

A disk contention can have a negative impact on user data read/writes to the disk devices, as well as on HADB writing to history files. Severe disk contention may delay or abort user transactions. Delay in history file writing may cause node restarts and, in the worst case, lead to double failures.

The disk contention can be identified by monitoring the disk I/O from the OS, for the disks used for data, log devices and history files. History file write delays are written to the HADB history files. This can be identified by “NSUP BEWARE *timestamp* Last flush took too long (x msec).”

This warning shows that disk I/O took too long. If the delay exceeds ten seconds, the node supervisor restarts the trans process with the error message:

```
Child process trans0 10938 does not respond.  
Child died - restarting nsup.  
Psup::stop: stopping all processes.
```

This message indicates that a trans (clu\_trans\_srv) process has been too busy doing other things (for example, waiting to write to the history file) to reply to the node supervisor’s request checking the heartbeat of the trans process. This causes the nsup to believe that the trans has died and then restarts it.

When the operating system is overloaded with too many processes (many HADB nodes co-located with other processes), the scheduling of I/O operations may be delayed. When the HADB related I/O work is delayed, HADB nodes write the following in the history files, “HADB warning: Schedule of async <read,write\> operation took ...”

This problem is observed especially in Red Hat AS 2.1 when multiple HADB nodes are placed on the same host and all the nodes use the same disk to place their devices.

### Solution

Use one disk per node to place the devices used by that node. If the node has more than one data devices and the disk contention is observed, move one data device to another disk. The same applies to the history file as well.

Make sure that all data and log devices and all history files reside on local disks (not NFS-mounted or other remotely mounted disks).

If the monitoring tools still show contention on the HADB disks, the data buffer pool size can be increased.

## Is There a Shortage of HADB Data Device Space?

### Description

One possible reason for transaction failure is running out of data device space. If this situation occurs, HADB will write warnings to the history file, and abort the transaction which tried to insert and/or update data.

Typical messages are:

```
HIGH LOAD: about to run out of device space, ...
```

```
HIGH LOAD: about to run out of device space on mirror node, ...
```

The general rule of thumb is that the data devices must have room for at least four times the volume of the user data. Please refer to the *Tuning Guide* for additional explanation.

### Solution 1

Increase the size of the data devices using the following command:

```
hadbm set DeviceSize=size
```

This solution requires that there is space available on the physical disks which are used for the HADB data devices on all nodes.

HADBM automatically restarts each node of the database.

### Solution 2

Stop and delete the HADB, and create a new instance with more nodes and/or larger data devices and/or several data devices per node. Unfortunately, using this solution will erase all persistent data and the schemas created by the Application Server. See the *Administrator's Guide* for more information about this procedure.

## Is There a Shortage of Other HADB Resources?

When an HADB node is started, it will allocate:

- Several shared memory segments of fixed size
- Internal data structures of fixed size

If an HADB node runs out of resources it will delay and/or abort transactions. Resource usage information is shipped between mirror nodes, so that a node can delay or abort an operation which is likely to fail on its mirror node.

Transactions that are delayed repeatedly may time out and return an error message to the client. If they do not time out, the situation will be visible to the client only as decreased performance during the periods in which the system is short on resources.

These problems frequently occur in “High Load” situations. For details, see [“High Load Problems” on page 54](#)

## High Load Problems

High load scenarios are recognizable by the following symptoms:

- User requests do not succeed
- The database gives multiple timeout and “transaction aborted” messages
- Frequent “HIGH LOAD” warnings in the history file
- Sporadic failures

If a high load problem is suspected, consider the following:

- [“Is the Tuple Log Out Of Space?” on page 54](#)
- [“Is the node-internal Log Full?” on page 55](#)
- [“Are There Enough Locks?” on page 55](#)
- [“Can You Fix the Problem by Doing Some Performance Tuning?” on page 56](#)

---

**Note** – Frequently, all of these problems can be solved by making more CPU horsepower available.

---

## Is the Tuple Log Out Of Space?

All user operations (delete, insert, update) are logged in the tuple log and executed. There tuple log may fill up because:

- Execution slows due to CPU or disk I/O contention
- The mirror node is slow in receiving the log records, which can happen as a result of:
  - Network contention, so the log records do not reach the mirror node
  - CPU and disk contention at the mirror node, which keeps it from processing the received log records quickly enough (“log throw due to . . .” messages in the history files).

If the tuple log is out of space, the history files contain messages showing HIGH LOAD on the tuple log.

### Solution 1

Check CPU usage, as described in [“Improving CPU Utilization” on page 57](#)

## Solution 2

If CPU utilization is not a problem, check the disk I/O. If the disk shows contention, avoid page faults when log records are being processed by increasing the data buffer size with `hadbm set DataBufferPoolSize=...`. If there is disk contention, follow the solutions suggested in [“Is There Disk Contention?”](#) on page 52

## Solution 3

Look for evidence of network contention, and resolve bottlenecks.

## Solution 4

Increase the tuple log buffer using `hadbm set LogBufferSize=...`

## Is the node-internal Log Full?

Too many node-internal operations are scheduled but not processed due to CPU or disk I/O problems.

If the node-internal log is out of space, the history files contain messages showing HIGH LOAD on the node internal log.

## Solution 1

Check CPU usage, as described in [“Improving CPU Utilization”](#) on page 57

## Solution 2

If CPU utilization is not a problem, check the disk I/O. If the disk shows contention, avoid page faults when log records are being processed by increasing the data buffer size with `hadbm set DataBufferPoolSize=...`. If there is disk contention, follow the solutions suggested in [“Is There Disk Contention?”](#) on page 52

## Are There Enough Locks?

Some extra symptoms that identify this condition are:

- Error code 2080 or 2096 delivered to the client.
- `hadbm resourceinfo --locks` shows locks allocated, and all are in use all the time

## Solution 1: Split Long Transactions

A transaction running on a node is not allowed to use more than 25% of the number of locks allocated on that node. Read transactions running at the “repeatable read” isolation level and the update/insert/delete transactions hold the locks until the transaction terminates. Therefore, it is recommended to split long transactions into small batch of separate transactions.

## Solution 2: Increase the number of locks

Use `hadbm set NumberOfLocks=` to increase the number of locks.

## Can You Fix the Problem by Doing Some Performance Tuning?

In most situations, reducing load or increasing the availability of resources will improve host performance. Some of the more common steps to take are:

- Run the nodes on hosts with better hardware characteristics (more internal memory, higher processor speed, more processors).
- Add physical disks and use several data devices, not more than one device on each physical disk.
- Add more nodes, on new hosts, and refragment the data to utilize the new nodes.
- Change configuration variables to allocate larger memory segments or internal data structures.

In addition, the following resources can be adjusted to improve “HIGH LOAD” problems, as described in the Performance and Tuning Guide:

TABLE 3-1 HADB Performance Tuning Properties

Resource	Property
Size of Database Buffer	<code>hadbm attribute DataBufferPoolSize</code>
Size of Tuple Log Buffer	<code>hadbm attribute LogBufferSize</code>
Size of Node Internal Log Buffer	<code>hadbm attribute InternalLogBufferSize</code>
Number of Database Locks	<code>hadbm attribute NumberOfLocks</code>



# Connection Problem Caused by Lack of Semaphore Resources

## Description

This problem is accompanied by a message in the history file:

```
HADB-E-05521: Operation on semaphore with key "46025" failed,  
OS status=28 : No space left on device where:
```

You must configure more semaphore unso structures on the host computer. See the “Preparing for HADB Setup” chapter in the *Sun GlassFish Enterprise Server Installation Guide* for information on how to configure `semnu` on your operating system.

## Solution

Stop the affected HADB node, reconfigure and reboot the affected host, restart the HADB node. HADB will be available during the process.

# Improving CPU Utilization

## Description

Available CPU cycles and I/O capacity can impose severe restrictions on performance. Resolving and preventing such issues is necessary to optimize system performance (in addition to configuring the HADB optimally.)

## Solution

If there are additional CPUs on the host that are not exploited, add new nodes to the same host. Otherwise add new machines and add new nodes on them.

If the machine has enough memory, increase the `DataBufferPoolSize`, and increase other internal buffers that may be putting warnings into the log files. Otherwise, add new machines and add new nodes on them.

For more information on this subject, consult the Performance and Tuning Guide.

# HADB Administration Problems

The `hadbm` command and its many subcommands and options are provided for administering the high-availability database (HADB). The `hadbm` command is located in the `install_dir/SUNWhadb/4/bin` directory.

Refer to the chapter on Configuring the High Availability Database in the *Sun GlassFish Enterprise Server Administrator's Guide* for a full explanation of this command. Specifics on the various `hadbm` subcommands are explained in the `hadbm` man pages.

The following problems are addressed in this section:

- “`hadbm` Command Fails: The agents could not be reached” on page 58
- “`hadbm` Command Fails: command not found” on page 59
- “`hadbm` Command Fails: `JAVA_HOME` not defined” on page 59
- “`hadbm` `createdomain` fails, but two split domains are created” on page 60
- “`create` Fails: path does not exist on a host” on page 60
- “Database Does Not Start” on page 61
- “`clear` Command Failed” on page 62
- “`create-session-store` Failed” on page 62
- “`hadbm` Command Hangs” on page 63
- “Cannot Restart the HADB” on page 64

## `hadbm` **Command Fails:** The agents could not be reached

### Description

The command fails with the error:

```
The agents <url\> could not be reached.
```

The hosts in the URL could be unreachable either because the hosts are down, because the communication pathway has not been established, because the port number in the URL is wrong, or because the management agents are down.

### Solution

Verify that the URL is correct. If the URL is correct, verify that the hosts are up and running and are ready to accept communications; for example:

```
ping hostname1ping hostname2...
```

## hadbm **Command Fails:** command not found

### Description

The hadbm command can be run from the current directory, or you can set the search PATH to access the hadb commands from anywhere, which is much more convenient. The error, “hadbm: Command not found,” indicates that neither of these conditions has been met.

### Solution 1

cd to the directory that contains the hadbm command and run it from there:

```
cd install_dir/SUNWhadb/4/bin/  
./hadbm
```

### Solution 2

Use the full path to invoke the hadbm command:

```
install_dir/SUNWhadb/4/bin/hadbm
```

### Solution 3

You can use the hadbm command from anywhere by setting the PATH variable. Instructions for setting the PATH variable are contained in the “Preparing for HADB Setup” chapter of the *Sun GlassFish Enterprise Server Installation Guide*.

To verify that the PATH settings are correct, run the following commands:

```
which asadmin  
which hadbm
```

These commands should echo the paths to the utilities.

## hadbm **Command Fails:** JAVA\_HOME not defined

### Description

The message “hadbm: <path\>: Invalid Java home location” indicates that the JAVA\_HOME environment variable has not been set properly.

### Solution

If multiple Java versions are installed on the system, ensure that the JAVA\_HOME environment variable points to the correct Java version (1.4.1\_03 or above for Enterprise Edition).

Instructions for setting the `PATH` variable are contained in the “Preparing for HADB Setup” chapter of the *Sun GlassFish Enterprise Server Installation Guide*.

## hadbm createdomain **fails**, but two split domains are created

### Description

If running the HADB management agent on a host with multiple network interfaces, the `createdomain` command may fail if not all network interfaces are on the same subnet:

```
hadbm:Error 22020: The management agents could not establish a domain,
please check that the hosts can communicate with UDP multicast.
```

By default, the management agents use the “first” interface for UDP multicasts (“first” as returned by `java.net.NetworkInterface.getNetworkInterfaces()`).

### Solution

The best solution is to tell the management agent which subnet to use by setting `ma.server.mainternal.interfaces` in the configuration file; for example:

```
ma.server.mainternal.interfaces=10.11.100.0
```

Alternatively, one may configure the router between the subnets to route multicast packets. By default, the management agent uses multicast address `228.8.8.8`.

## create **Fails**: path does not exist on a host

### Description

After issuing the `hadbm create` command, an error similar to the following appears on the console:

```
./hadbm create ...
...
hadbm: Error 22022: Path path does not exist on host host
```

This error message can also appear when new nodes are added without the specified paths do not exist on the machines.

## Solution

Log in to the host and create paths for the HADB devices and HADB history files. Run `hadbm create` and specify the `--devicepath` and `--historypath` options to the paths created. Also make sure that the user running the management agent on the host has read and write access to these directories.

---

**Note** – HADB executables cannot be installed on different paths on different hosts.

---

## Database Does Not Start

The create or start command fails with the console error message:

```
hadbm: Error 22095: Database could not be started...
```

Consider the following possibilities:

### Was there a shared memory get segment failure?

#### Description

Start may fail if the resources (shared memory, disk space) allocated for that node are taken by some other processes, after the node is stopped.

#### Solution

Refer to [“Problems Related to Shared Memory”](#) on page 44 for suggestions on resolving this issue.

### Do the History Files Contain Errors?

#### Description

If the problem still persists, inspect the HADB history files. Some of the more likely error messages to look for are:

- Could not verify node address
  - This message occurs when another process is using the port that an HADB server is trying to process. It can occur in several situations:
    - The `portBase` is used by another process running on this host machine.
      - Set the `PortBase` attribute to another value using the following command:

```
hadbm set portbase=value
```

- An attempt to stop the HADB node for maintenance failed.  
Try again to stop the node with the `hadbm` command. If that fails, `kill` the OS process `clu_nsup_srv` for this node *without* the `-9` option. The `nsup` process should then stop its `hadb` child process. If the parent process `nsup` does not exist, `kill` all the child processes using `kill -9`.
- The HADB node was stopped for maintenance and an `inetd` process restarted the HADB node before you intended to start it.  
Make sure that `inetd` does not start the HADB node before stopping it.
- `hadbm` *command* fails with internal error: "The database could not be started"  
Check the following:
  - Shared memory is all correct on all machines in the HADB configuration.
  - No other HADB databases are running on the machines, or any other processes that could be using the same port numbers.
  - All necessary directories exist and have write permissions.
  - There is enough space in directory where devices are going to be written.

## Solutions

After verifying that none of the above errors have occurred, try the following remedies, in order:

- Delete the database and retry.
- Delete the database, reboot, and retry.
- Delete database, reinstall the HADB software, and retry.
- [“Contact Support” on page 21](#).

For more information, refer to the *Error Message Reference*.

## clear Command Failed

The `clear` command reinitializes the database device files residing on disks. This may fail due to problems with disk or disk access. Check whether any error message from `hadbm` indicates this. If not, look into the `ma .log` files and check whether `devinit` has generated any error messages.

## create-session-store Failed

The `asadmin create-session-store` command could fail for one of these reasons:

- [“Invalid user name or password” on page 63](#)
- [“SQLException: No suitable driver” on page 63](#)

## Invalid user name or password

This error occurs when the `--dbssystempassword` supplied to the `create-session-store` command is not the same password as the one given at the time of database creation.

### Solution 1

Try the command again with the correct password.

### Solution 2

If you cannot remember the `dbssystem` password, you need to clear the database using `hadbm clear` and provide a new `dbssystem` password.

## SQLException: No suitable driver

The `create-session-store` produces the error: `SessionStoreException: java.sql.SQLException: No suitable driver.`

### Solution 1

This error can occur when `asadmin` is not able to find `hadbjdbc4.jar` from the `AS_HADB` path defined in `asenv.conf` in the Enterprise Server `config` directory.

The solution is to change `AS_HADB` to point to the location of the HADB installation.

Here is a sample `AS_HADB` entry from an `asenv.conf` file:

```
AS_HADB=/export/home0/hercules/0815/SUNwhadb/4.4.0-8
```

### Solution 2

This error can also occur if you provide the incorrect value for `--storeUrl`. To solve this problem, obtain the correct URL using `hadbm get jdbcURL`.

## hadbm Command Hangs

If the management agent with which the `hadbm` communicates dies before the operation finishes, then the `hadbm` process may hang. Check whether the all the agents are running.

## Cannot Restart the HADB

### Description

HADB restart does not work after a double node failure. Additional recovery actions are needed before HADB can be restarted.

Symptoms of a double node failure include:

- `hadbm status` shows that the HADB status is non-operational.
- The node status shows that the nodes are in `Starting` or `Recovering` state. Even after stopping and then restarting each of the nodes, they remain in the `Starting` state. Eventually, the node status changes to `Stopped`.

This problem occurs when mirror HADB host machines have failed or been rebooted, typically after a power outage, or when a machine is rebooted without first stopping the HADB (in a single-machine installation), or when a pair of mirror machines from both Data Redundancy Units (DRUs) are rebooted.

HADB cannot heal itself automatically in such “double failure” situations because the part of the data that resided on the pair nodes is lost. In such cases, the `hadbm start` command does not succeed, and the `hadbm status` command shows that HADB is in a non-operational state.

For more information on the DRUs and HADB configuration, see “Administering the High Availability Database” in the *Administration Guide*, and the *Deployment Guide*.

---

**Tip** – If the HADB exhibits strange behavior (for example consistent timeout problems), and you want to check whether a restart cures the problem, use the `hadbm restart` command.

When the HADB is restarted in this manner, HADB services remain available. Conversely, if HADB is started and stopped in separate operations using `hadbm stop` and `hadbm start`, HADB services are unavailable while HADB is stopped.

---

### Solution

Verify that the node states show `Starting/Recovering`, then reset the database. Follow the instructions under “Recovering from Session Data Corruption” in the “Administering the High Availability Database” chapter of the *Administration Guide*.



# Shared Memory Segment Key Already Exists (Windows only)

## Description

The hadbm process returns the following error:

```
HADB-S-05515: Shared memory segment with key "NNNNN" exists already
```

This can happen during HADB instance creation following a controlled stop without deleting a previously created instance that is using the same portbase. The problem may also be the result of a failed HADB instance deletion for any reason.

## Solution

Delete all stopped hadb instances to make sure all HADB resources are free before attempting to reuse them.

If the problem persists, manually remove the HADB Shared Memory segments by deleting the HADB files in \$TMP/f\_.\*.

# Failure in configure-ha-cluster

## Description

Creation of an HADB domain comprising some host names appears to succeed, and the listdomain command confirms it:

```
$hadbm listdomain -w admin
Hostname Enabled? Running? Release Interfaces
host1 Yes      Yes      V4-4-1-3 128.139.113.110
host2 Yes      Yes      V4-4-1-3 128.139.113.111
```

The database is then created with the hadbm create command, and the appropriate host names including the full domain names are used as parameters for the --hosts option:

```
$ hadbm create --hosts=host1.xyz.abc.com,host2.xyz.abc.com...
```

The following error is then returned:

```
hadbm:Error 22176: The host host1.xyz.abc.com is not registered in the
HADB management domain. Use hadbm createdomain to set up the management
domain or hadbm extenddomain to include new hosts in an existing domain.
```

## Solution 1

Use the names that `listdomain` displays; for example:

```
hadbm create --hosts=host1,host2...
```

## Solution2

Use decimal IP addresses (DDN); for example:

```
hadbm create --hosts=128.139.113.110,128.139.113.111
```

# Unable to Run configure-ha-cluster

## Description

Two different installations of HADB are configured: one on server hosts, and another on the `hadbm` client host(s), each running different versions of HADB. The management agents are started with one HADB versions, and then `hadbm create` is run with the other version. The following error is returned:

```
HADBMGMT007:hadbm create command failed. Return value: 1 STDOUT:  
STDERR: hadbm:Error 22170: The software package V4.4.x could not be  
found at path <pacakgepath\>/4.4-x on host <hostname\>.  
CLI137 Command configure-ha-cluster failed.
```

## Solution

Use the same HADB version for the management agents and all `hadbm` clients.

# hadbm set Command Fails

## Description

`hadbm set` brings the database instance to a state that is hard to recover from.

Changing a database configuration variable with the `hadbm set` command fails. For example, setting `DataBufferPoolSize` to a larger size fails due to lack of shared memory on node-0. The

hadbm set command leaves the database with node-0 in a stopped state and node-1 in a running state. Resetting the pool size back to the original value with the help of hadbm set fails with the message:

```
22073: The operation requires restart of node 1. Its mirror node is
currently not available. Use hadbm status --nodes to see the status of
the nodes.
```

The hadbm startnode 0 command is also of no use in this situation.

## Solution

Stop the database, then restore the old values using hadbm set, then restart the database.

# Failure in configure-ha-cluster: Creating an HADB Instance Fails

## Description

Creation of an HADB cluster fails with the message:

```
cresqldict: HADB-E-00208: The transaction was aborted.
```

This indicates that the booting transaction populating the SQL dictionary tables was aborted.

## Solution

Run the configure-ha-cluster again. If the hadbm create command fails with the above message, rerun the command.



# Security Problems

---

This chapter covers problems that you may encounter as a result of security settings:

- “`java.security.AccessControlException: Access Denied Error`” on page 69
- “`javax.ejb.AccessLocalException: Client Not Authorized Error`” on page 70
- “`Authentication is Not Working With the Solaris Realm`” on page 70
- “`Mutual Authentication Not Working With the Application Client`” on page 71

## `java.security.AccessControlException: Access Denied Error`

### Description

The following error occurs from an application client or in the `server.log`:

```
java.security.AccessControlException: access denied  
(java.util.PropertyPermission name write...)
```

There is a permissions issue in the policy files. Either the `client.policy` for the application client or the `server.policy` for server side components does not have permission to set the property.

### Solution

Add the permission in `client.policy` (for the application client), or in `server.policy` (for EJB/web modules) for the application that needs to set the property. By default, applications only have “read” permission for properties.

For example, to grant read/write permission for all the files in the codebase directory, add or append the following to `client.policy` or `server.policy`:

```
grant codeBase "file:/../build/sparc_SunOS/sec/-" {  
    permission java.util.PropertyPermission "*", "read,write";  
};
```

## javax.ejb.AccessLocalException: Client Not Authorized Error

### Description

Role-mapping information is available in Sun-specific XML (for example, `sun-ejb-jar.xml`), and authentication is okay, but the following error message is displayed:

```
[...INFO|sun-appserver-pe8.0|javax.enterprise.system.container.ejb|...|  
javax.ejb.AccessLocalException: Client not authorized for this invocation.  
at com.sun.ejb.containers.BaseContainer.preInvoke(BaseContainer.java:...  
at com.sun.ejb.containers.EJBObjectInvocationHandler.invoke(...)
```

### Solution

Check whether the EJB module (`.jar`) or web module (`.war`) is packaged in an application (`.ear`) and does not have role-mapping information in application level, Sun-specific, `sun-application.xml`. For any application (`.ear`), security role-mapping information must be specified in `sun-application.xml`. It is acceptable to have both module-level XML and application-level XML.

## Authentication is Not Working With the Solaris Realm

Check whether the installation and server startup was performed as a local user, instead of as the root user. Always start the Enterprise Server as the root user, because the Solaris realm works only with the root user. It was not designed to work with any other local user. Note also that role mapping can happen on the local user.

# Mutual Authentication Not Working With the Application Client

## Description

This failure can occur when the keystore and truststore properties are not set properly.

## Solution

Set the following properties on the JVM:

```
javax.net.ssl.keyStore=  
<keystore-file-path\>;javax.net.ssl.trustStore=<truststore-file-path\>
```

To use the application client, set the environment variable VMARGS to the following value:

```
-Djavax.net.ssl.keyStore=  
${admin.domain.dir}/${admin.domain}/config/keystore.jks  
-Djavax.net.ssl.trustStore=  
${admin.domain.dir}/${admin.domain}/config/cacerts.jks
```





## Frequently Asked Questions

---

This section covers some common questions asked about the Enterprise Server:

- “What Happens When No Server Side Realm is Configured?” on page 73
- “Can I Use a PKCS12 Certificate for My Client Certificate?” on page 73
- “Can I See the TLS/SSL Handshake Information for an SSL Client?” on page 74
- “Can I Change the Keystore Password?” on page 74
- “How Do I Maintain a Session in JAX-RPC?” on page 74
- “How Do I Access the Naming Service From a Standalone Java Client?” on page 75
- “Are RMI-IIOP Stubs Needed to Access Remote EJBs?” on page 76
- “How Do I Change the Log Level for an Application Logger?” on page 77

### **What Happens When No Server Side Realm is Configured?**

When the application is configured (within XML files), but no server side realm is configured, the application is authenticated in the default realm. No error is thrown that indicates “No such realm.”

### **Can I Use a PKCS12 Certificate for My Client Certificate?**

Is there a way to use my PKCS12 certificate for an SSL the application client or standalone client during mutual authentication?

No. You cannot use a PKCS12 certificate directly, but you can write your own client using the JSSE, which supports `storetype=PKCS12` (read only, no write to keystore).

## Can I See the TLS/SSL Handshake Information for an SSL Client?

Yes. Set the Java debugging property on the JVM. To see the handshake information from the application client, append the following:

```
-Djavax.net.debug=ssl,handshake to the VMARGS variable.
```

## Can I Change the Keystore Password?

Yes. Use the following J2SE properties to change the keystore password:

```
-Djavax.net.ssl.keyStorePassword=password  
-Djavax.net.ssl.trustStorePassword=password
```

Note that the keystore password must match the individual key passwords to perform operations on the keys, so you will need to change the keystore password with the property mentioned above and then change the password to each key to match that password.

## How Do I Maintain a Session in JAX-RPC?

Clients cannot maintain sessions with JAX-RPC endpoints. There is a client and server aspect to sessions, and it is not obvious how to set this up.

The situation is that a client makes a call to the service, and the server responds and sets a cookie on the connection. From then on, the client sends back that same cookie with each call and the server can check it.

A JAX-RPC stub normally ignores the cookie that comes back. When the `SESSION_MAINTAIN_PROPERTY` is set to `true`, it sends back whatever cookie the server set on it.

On the server side, you need to add one field to your class and a method that sets it. The endpoint must implement `javax.xml.rpc.server.ServiceLifecycle`, and two methods must be added: `destroy()` (which can be empty) and `init(Object context)`.

Add a `ServletEndpointContext` object to the endpoint; for example `myServletEndpointContext`. The `init(Object context)` method can be set as follows:

```
myServletEndpointContext = (ServletEndpointContext) context;
```

From then on, the business methods can access to the `HttpSession` with `myServletEndpointContext.getHttpSession()`. The first call to `getHttpSession` creates the session, if one does not already exist.

With this model, any method the client calls can get the session, set session attributes, get values from it, and so on. From then on, the client will send back the same cookie information.

## How Do I Access the Naming Service From a Standalone Java Client?

### ▼ To access the naming service from an application client

- 1 **Include `appserv-rt.jar` in the CLASSPATH when starting the client Java VM.**

The JNDI bootstrapping machinery looks for a file called `jndi.properties`, which is located in `appserv-rt.jar`. This file contains all the bootstrapping properties for the Enterprise Server's naming service. It is better to have these properties read from `appserv-rt.jar` than to hard-code them in either the client startup script or in the application code.

- 2 **When accessing remote EJBs from a standalone client, it is not necessary to retrieve the client JAR from the deployment or to put it in the client JVM's CLASSPATH, because static RMI-IIOP stubs are not needed when using the Enterprise Server naming service. This removes a step that was required in previous releases. (See [“Are RMI-IIOP Stubs Needed to Access Remote EJBs?”](#) on [page 76](#) for more details).**

- 3 **Code the client to use the default constructor `InitialContext` that does not require an argument. For example:**

```
InitialContext ic = new InitialContext();
```

It is a common misconception that the client should be coded to explicitly reference the `CosNaming` service. `CosNaming` is only used for some kinds of Enterprise Server objects, so doing this will not provide access to many of the kinds of resources you might need in the client such as JMS queues, connection factories, and so on. Furthermore, explicit use of `CosNaming` bypasses the Enterprise Server's naming service code. This often means that the client cannot take advantage of desirable value-added behavior built in to the Enterprise Server's naming service.

- 4 **Use the global JNDI name of the target resource when doing the lookup. `java:comp/env` cannot be used from standalone Java clients, because by definition such clients run outside of a J2EE container. The only client component in which `java:comp/env` can be used is in a J2EE Application Client.**

- 5 If the client is running on a different host machine than the server instance, set the following system property when starting the Java VM:**

```
-Dorg.omg.CORBA.ORBInitialHost=hostname_of_target_server
```

This value defaults to localhost so it is only needed if the client and server instance are not colocated. For example:

```
java -Dorg.omg.CORBA.ORBInitialHost=server1 ... com.foo.MyMainClass
```

- 6 By default, the client attempts to contact port 3700 to access the naming service in the server. Since 3700 is the default naming service port used by the Enterprise Server, there is no additional port configuration needed in the client. In some cases, due to port conflicts, the server instance uses a different naming service port. The naming service port used by the server instance is listed in the `<iiop-listener id="orb-listener-1" port="3700">` element in `domain.xml`.**

To change the naming service port used by the client, set the following system property when starting the client Java VM:

```
-Dorg.omg.CORBA.ORBInitialPort=naming_port_of_target_server
```

## Are RMI-IIOP Stubs Needed to Access Remote EJBs?

No. Unlike previous releases of the Enterprise Server, the current version does not require static RMI-IIOP stubs at runtime.

Removing this requirement provides the following benefits:

- Faster deployment and redeployment time for applications containing remote EJBs or clients of remote EJBs
- Fewer runtime errors caused by stub CLASSPATH configuration problems

In addition, the Enterprise Server achieves these benefits without significant impact on runtime performance, while maintaining full RMI-IIOP interoperability.

The only scenario where RMI-IIOP stubs are still required is for standalone clients that explicitly instantiate an `InitialContext` for the `CosNaming` naming service. This is *not* the recommended approach for using the naming service in the Enterprise Server. However, to maintain compatibility for these kinds of standalone clients, there is a deployment-time option that forces the generation of RMI-IIOP stubs in a way that matches previous releases. To use it, set

```
--generatermistubs=true
```

when deploying with `asadmin` or the Administration Console. The RMI-IIOP stubs are placed in the `client.jar` file, just as they were in previous releases.

## How Do I Change the Log Level for an Application Logger?

Each application uses its own application logger to log messages. To configure the log level for a particular application, use one of two options:

- In the Admin GUI Log Level configuration page add a property with a property name representing the logger name, and the value representing one of seven log levels (FINEST, FINER, FINE, CONFIG, INFO, WARNING, SEVERE) or OFF.

For example, to change the log level of application logger named `com.X.Y` to `FINEST`, the property name would be `com.X.Y` and the property value would be `FINEST`. The change is reflected in the `domain.xml` file, and takes effect immediately. No Server restart is required.

- Directly add the property to the `<module-log-levels>` element in `domain.xml`, as shown below.

```
<module-log-levels admin="INFO" classloader="INFO" cmp="INFO"
cmp-container="INFO" configuration="INFO" connector="INFO"
corba="INFO" deployment="INFO" ejb-container="INFO" javamail="INFO"
jaxr="INFO" jaxrpc="INFO" jdo="INFO" jms="INFO" jta="INFO"
jts="INFO" mdb-container="INFO" naming="INFO" node-agent="INFO"
resource-adapter="INFO" root="INFO" saaj="INFO" security="INFO"
server="INFO" synchronization="INFO" util="INFO" verifier="INFO"
web-container="INFO">
  <property name="com.X.Y" value="FINEST" />
</module-log-levels>
```



# Index

---

## A

- ACC log service, 19
- admin password
  - invalid, 34-35
  - not known, 35-36
- admin username
  - invalid, 34-35
  - not known, 35-36
- Administration Console, cannot access, 32-33
- application, cannot access, 34-36

## C

- can't recall admin user name or password, 35-36
- cannot access Administration Console, 32-33
- cannot access local server, 30-31
- cannot access remote server, 31-32
- cannot access server application, 34-36
- Cannot find log files, 29-30

## D

- diagnosis of a problem, 14
- documentation, 20
- domain
  - finding, 32
  - restarting, 25-29

## E

- error messages, 16

## F

- forum, 21

## H

- HADB, 43
  - setting the environment, 59
- HADB will not start, 64
- hadbm command not working, 59
- high-load situation, 54
- hostname, finding, 32
- HTTP server access logs, 29

## I

- invalid admin username or password, 34-35
- investigation of a problem, 16

## J

- JAVA\_HOME setting, 59

## **K**

Knowledge Base, 20

## **L**

local server, access fails, 30-31

log files, 17

log levels, 18

logs

cannot find, 29-30

directory, 29

## **M**

message types, 29

## **N**

node agent, restarting, 25-29

## **O**

online forum, 21

## **P**

pkginfo, 14

port number, how the default value can change, 31

power failure and HADB, 64

## **R**

remote server, access fails, 31-32

response-timeout-in-seconds, 48

restarting server, 25-29

rpm, 14

## **S**

server

how default port number can change, 31

Server, log file, description of, 29

server

log file, location of, 29

will not start on windows, 36

server.log, 18

showrev, 14

start-domain command fails, 24-25

SunSolve, 20

## **T**

thread dump, obtaining, 19

## **U**

ungraceful shutdown, 64

## **W**

Windows

log levels, 18

server will not start, 36