



Sun Java™ System

Application Server 8.1
Enterprise Edition
Troubleshooting Guide

2005Q1

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

Part No: 819-0086

Copyright © 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 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 of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.

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.

Use is subject to license terms. This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo, Java, and the Java Coffee Cup logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

This product is covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical 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 © 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. 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 plus des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou les brevets supplémentaires ou les applications de brevet en attente aux Etats - Unis et dans les autres pays.

CE PRODUIT CONTIENT DES INFORMATIONS CONFIDENTIELLES ET DES SECRETS COMMERCIAUX DE SUN MICROSYSTEMS, INC. SON UTILISATION, SA DIVULGATION ET SA REPRODUCTION SONT INTERDITES SANS L'AUTORISATION EXPRESSE, ECRITE ET PREALABLE DE SUN MICROSYSTEMS, INC.

L'utilisation est soumise aux termes de la Licence. Cette distribution peut comprendre des composants développés par des tierces parties.

Sun, Sun Microsystems, le logo Sun, Java, et le logo Java Coffee Cup sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, 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

Preface	vii
Who Should Use This Book	vii
How This Book Is Organized	vii
Conventions Used in This Book	viii
Typographic Conventions	viii
Symbols	ix
Default Paths and File Names	ix
Shell Prompts	x
Related Documentation	x
Books in This Documentation Set	x
Other Server Documentation	xii
Accessing Sun Resources Online	xii
Contacting Sun Technical Support	xii
Related Third-Party Web Site References	xiii
Sun Welcomes Your Comments	xiii
Chapter 1 Overview	1
Planning Ahead	1
Identifying the Problem	2
Seeking a Solution	4
Verify System Configuration	4
Evaluate Messages	4
Examine Log Files	6
See if the Problem has been Solved Before	8
Search the Product Documentation	8
Search the Knowledge Base	8
Search or Participate in the Online Forum	9
Contact Support	9
Chapter 2 Common Problems	11
asadmin start-domain Command Fails	12

Error: "CLI143 There is more than one domain..."	12
Error: "Could Not Start Domain"	12
Automatically Restarting a Domain or Node Agent on Reboot	13
Restarting Automatically on UNIX Platforms	13
Restarting Automatically on the Microsoft Windows Platform	14
Cannot Find Log Files	17
Accessing Local Server Fails (http://localhost:8080)	19
Did the Server Start?	19
Was the Server Started at the Expected Port?	20
Accessing Remote Server Fails	20
Is the Server Available Locally?	21
Is the Proxy Setting Causing a Problem?	21
Cannot Access the Administration Console	22
Is the Application Server Running?	22
Is the Administration Console Running on the Expected Port?	23
Is the Security Manager Disabled?	23
Cannot Access a Server Application	24
Is the Application Server Running?	24
Was Application Deployment Successful?	24
Server Will Not Start on Windows (Port Conflict)	24
Is Another Application Running on the Server's Port?	25
Has an Ungraceful Shutdown Occurred on a Previously Running Server?	25
Port Conflicts Debugging Multiple Instances on the Same Server	25
Two Server Instances Bind to Same Port on Windows	26
Error: "System cannot find the path specified"	27
Application Generates Error persistence.support.JD0DataStoreException	27
Using asadmin set Command May Produce Unexpected Results	28
Cannot Undeploy Or Redeploy Application With Open Streams to Jar Files (Windows Only)	28
Cannot Reinstall Application Server After Manual Deletion of Directories	29
Cannot Produce a JVM Thread Dump After Server Crash	30
Chapter 3 HADB Problems	31
HADB Database Creation Fails	32
Problems Related to Shared Memory	32
Too Few Semaphores	34
Database Nodes Cannot Be Reached and the Database Does Not Function	34
The Management Agents Could Not Establish a Domain	34
hadbm create or hadbm addnodes Command Hangs	35
ma (Management Agent Process) Crashes	36
Server Responds Slowly After Idle Period	36
Requests Are Not Succeeding	37
Is the Load Balancer Timeout Correct?	37
Are the System Clocks Synchronized?	37

Is the Application Server Communicating With HADB?	38
Session Persistence Problems	38
The create-session-store Command Failed	39
Configuring Instance-Level Session Persistence Did Not Work	39
Session Data Seems To Be Corrupted	39
HADB Performance Problems	40
Is There a Shortage of CPU or Memory Resources, or Too Much Swapping?	41
Is There Disk Contention?	41
Is There a Shortage of HADB Data Device Space?	43
Is There a Shortage of Other HADB Resources?	43
High Load Problems	44
Is the Tuple Log Out Of Space?	45
Is the node-internal Log Full?	45
Are There Enough Locks?	46
Can You Fix the Problem by Doing Some Performance Tuning?	46
Connection Problem Caused by Lack of Semaphore Resources	47
Improving CPU Utilization	48
HADB Administration Problems	48
hadbm Command Fails: The agents could not be reached	49
hadbm Command Fails: command not found	49
hadbm Command Fails: JAVA_HOME not defined	50
hadbm createdomain fails, but two split domains are created	51
create Fails: "path does not exist on a host"	51
Database Does Not Start	52
clear Command Failed	54
create-session-store Failed	54
hadbm Command Hangs	55
Cannot Restart the HADB	55
Shared Memory Segment Key Already Exists (Windows only)	56
Failure in configure-ha-cluster	57
Unable to Run configure-ha-cluster	57
hadbm set Command Fails	58
Failure in configure-ha-cluster: Creating an HADB Instance Fails	58
Chapter 4 Security Problems	61
java.security.AccessControlException: Access Denied Error	61
javax.ejb.AccessLocalException: Client Not Authorized Error	62
Authentication is Not Working With the Solaris Realm	62
Mutual Authentication Not Working With the Application Client	63
Chapter 5 Frequently Asked Questions	65
What Happens When No Server Side Realm is Configured?	65

Can I Use a Pkcs12 Certificate for My Client Certificate?	65
Can I See the TLS/SSL Handshake Information for an SSL Client?	66
Can I Change the Keystore Password?	66
How Do I Maintain a Session in JAX-RPC?	66
How Do I Access the Naming Service From a Standalone Java Client?	67
Are RMI-IIOP Stubs Needed to Access Remote EJBs?	68
How Do I Change the Log Level for an Application Logger?	69
Index	71

Preface

This guide describes common problems encountered when using the Sun Java™ System Application Server Server 8.1 Standard and Enterprise Editions, and how to solve them.

Who Should Use This Book

The *Troubleshooting Guide* is intended for developers and administrators who use the Sun Java™ System Application Server to assemble and deploy distributed and Web-based applications.

- Java APIs as defined in the Java™ Servlet, JavaServer Pages™ (JSP™), Enterprise JavaBeans™ (EJB™), and Java™ Database Connectivity (JDBC™) specifications
- The SQL structured database query languages
- Relational database concepts

How This Book Is Organized

This book describes workarounds for some problems you may encounter when configuring, using, or deploying applications with the Sun Java™ System Application Server 8.1 EE or PE. The following table summarizes the content of this book.

Table 1 How This Book Is Organized

Chapter	Description
Chapter 1, Overview	Provides a general overview of the troubleshooting process.
Chapter 2, Common Problems	Covers the most common problems you may encounter when using the product.
Chapter 3, HADB Problems	Covers problems related to the High Availability Database (HADB).
Chapter 4, Security Problems	Covers problems related to security settings.
Chapter 5, Frequently Asked Questions	Covers commonly asked questions about Application Server software.

Conventions Used in This Book

The tables in this section describe the conventions used in this book.

Typographic Conventions

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

Table 2 Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123 (Monospace)	API and language elements, HTML tags, web site URLs, command names, file names, directory path names, onscreen computer output, sample code.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123 (Monospace bold)	What you type, when contrasted with onscreen computer output.	% su Password:
<i>AaBbCc123</i> (Italic)	Book titles, new terms, words to be emphasized. A placeholder in a command or path name to be replaced with a real name or value.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. Do <i>not</i> save the file. The file is located in the <i>install-dir/bin</i> directory.

Symbols

The following table describes the symbol conventions used in this book.

Table 3 Symbol Conventions

Symbol	Description	Example	Meaning
[]	Contains optional command options.	ls [-l]	The -l option is not required.
{ }	Contains a set of choices for a required command option.	-d {y n}	The -d option requires that you use either the y argument or the n argument.
-	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.
>	Indicates menu item selection in a graphical user interface.	File > New > Templates	From the File menu, choose New. From the New submenu, choose Templates.

Default Paths and File Names

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

Table 4 Default Paths and File Names

Term	Description
<i>IdentityServer-base</i>	Represents the base installation directory for Identity Server. The Identity Server 2004Q2 default base installation and product directory depends on your specific platform: Solaris™ systems: /opt/SUNWam Linux systems: /opt/sun/identity
<i>DirectoryServer-base</i>	Represents the base installation directory for Sun Java System Directory Server. Refer to the product documentation for the specific path name.
<i>ApplicationServer-base</i>	Represents the base installation directory for Sun Java System Application Server. Refer to the product documentation for the specific path name.

Table 4 Default Paths and File Names (*Continued*)

Term	Description
<i>WebServer-base</i>	Represents the base installation directory for Sun Java System Web Server. Refer to the product documentation for the specific path name.

Shell Prompts

The following table describes the shell prompts used in this book.

Table 5 Shell Prompts

Shell	Prompt
C shell on UNIX or Linux	<i>machine-name%</i>
C shell superuser on UNIX or Linux	<i>machine-name#</i>
Bourne shell and Korn shell on UNIX or Linux	\$
Bourne shell and Korn shell superuser on UNIX or Linux	#
Windows command line	C:\

Related Documentation

The <http://docs.sun.com>SM website enables you to access Sun technical documentation online. You can browse the archive or search for a specific book title or subject.

Books in This Documentation Set

The Sun Java System Application Server Standard and Enterprise Edition manuals are available as online files in Portable Document Format (PDF) and Hypertext Markup Language (HTML).

The following table lists tasks and concepts described in the Sun Java™ System Application Server manuals. The manuals marked (*updated for 8 2004Q5*) have been updated for the Sun Java System Application Server Standard and Enterprise Edition 8 2004Q5 release. The manuals not marked in this way have not been updated since the version 8 Enterprise Edition release.

Table 6 Books in This Documentation Set

Book Title	Description
<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, JDK, and JDBC/RDBMS.
<i>Quick Start Guide</i>	How to get started with the Sun Java System Application Server product.
<i>Installation Guide</i>	Installing the Sun Java System Application Server software and its components.
<i>Deployment Planning Guide</i>	Evaluating your system needs and enterprise to ensure that you deploy Sun Java System Application Server in a manner that best suits your site. General issues and concerns that you must be aware of when deploying an application server are also discussed.
<i>Developer's Guide</i>	Creating and implementing Java™ 2 Platform, Enterprise Edition (J2EE™ platform) applications intended to run on the 2004Q4 Beta Sun Java System Application Server that follow the open Java standards model for J2EE components and APIs. Includes general information about developer tools, security, assembly, deployment, debugging, and creating lifecycle modules.
<i>J2EE 1.4 Tutorial</i>	Using J2EE 1.4 platform technologies and APIs to develop J2EE applications and deploying the applications on the Sun Java System Application Server.
<i>Administration Guide</i>	Configuring, managing, and deploying the Sun Java System Application Server subsystems and components from the Administration Console.
<i>High Availability Administration Guide</i>	Post-installation configuration and administration instructions for the high-availability database.
<i>Administration Reference</i>	Editing the Sun Java System Application Server configuration file, <code>domain.xml</code> .
<i>Upgrade and Migration Guide</i>	Migrating your applications to the new Sun Java System Application Server programming model, specifically from Application Server 6.x and 7. This guide also describes differences between adjacent product releases and configuration options that can result in incompatibility with the product specifications.
<i>Performance Tuning Guide</i>	Tuning the Sun Java System Application Server to improve performance.
<i>Error Message Reference</i>	Solving Sun Java System Application Server error messages.
<i>Reference Manual</i>	Utility commands available with the Sun Java System Application Server; written in manpage style. Includes the <code>asadmin</code> command line interface.

Other Server Documentation

For other server documentation, go to the following:

- Directory Server documentation
http://docs.sun.com/coll/DirectoryServer_04q2
- Web Server documentation
http://docs.sun.com/coll/S1_websvr61_en
- Application Server documentation
http://docs.sun.com/coll/s1_asseu3_en
- Web Proxy Server documentation
<http://docs.sun.com/prod/s1.webproxys#hic>

Accessing Sun Resources Online

For product downloads, professional services, patches and support, and additional developer information, go to the following:

- Download Center
<http://www.sun.com/software/download/>
- Professional Services
<http://www.sun.com/service/sunps/sunone/index.html>
- Sun Enterprise Services, Solaris Patches, and Support
<http://sunsolve.sun.com/>
- Developer Information
<http://developers.sun.com/prodtech/index.html>

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in the product documentation, go to <http://www.sun.com/service/contacting>.

Related Third-Party Web Site References

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 Send Comments. 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. For example, the title of this book is *Sun Java™ System Application Server 2005Q1 Troubleshooting Guide*, and the part number is 819-0789.

Sun Welcomes Your Comments

Overview

This chapter provides a description of the tools, methods, and information sources available for troubleshooting the Sun Java™ System Application Server Server 8.1 Enterprise Edition. Guidelines for evaluating and investigating a problem are included.

- [Planning Ahead](#)
- [Identifying the Problem](#)
- [Seeking a Solution](#)

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 Application 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 *Application Server 8.1 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:

1. Start the Application 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 Application 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

J2EE application servers are typically deployed in complex and highly sophisticated operating environments. The Sun Java™ System Application 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 Application Server are LDAP, Web Server, Sun ONE Message Queue, 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.

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 Application Server? On the same machine or not?
7. How is the Application Server connected to the directory server?
8. Are Sun Java™ Systems 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 Application 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](#)
- [Evaluate Messages](#)
- [Examine Log Files](#)
- [See if the Problem has been Solved Before](#)
- [Search the Product Documentation](#)
- [Search the Knowledge Base](#)
- [Search or Participate in the Online Forum](#)
- [Contact Support](#)

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 Java™ System Application Server 8.1 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.ipplanet.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 Application Server error messages, refer to the *Sun Java™ System Application Server Error Messages Reference* at the following location:

http://docs.sun.com/db/coll/ApplicationServer8_04q4

Examine Log Files

A number of the Application 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 – Application Server-specified log level ID or name
- Process identifier (PID) – PID of the Application 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 Application Server problem area are discussed in the associated chapters of this manual.

Log Levels

The Application Server has many log levels that can be set in the Administration interface (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 Java™ System Application Server Administrator's Guide*.

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

Log Options

The Administration interface 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.)

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 Sun ONE Application 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 Java™ System Application Server Developer's Guide to Clients* for more information.

Obtaining a Thread Dump

This section explains how to obtain a thread dump for Application Server 8.1. By default, the server dumps a core file and restarts with the `-Xrs java-option` flag in the `server.xml` file.

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.

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 Application 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 Application Server product release is available at:

<http://docs.sun.com/db/coll/ApplicationServer81>

Descriptions of the Application Server manuals are listed in “[Conventions Used in This Book.](#)”

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:

1. Go to SunSolve.
<http://sunsolve.sun.com/pub-cgi/show.pl?target=home>
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 Application Server online forum is available at:

<http://swforum.sun.com/jive/index.jsp?cat=7>

Contact Support

When necessary, gather together the information you have acquired and contact technical support, as described in “[Contacting Sun Technical Support](#).”

Seeking a Solution

Common Problems

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

- `asadmin start-domain` Command Fails
- Automatically Restarting a Domain or Node Agent on Reboot
- Cannot Find Log Files
- Accessing Local Server Fails (<http://localhost:8080>)
- Accessing Remote Server Fails
- Cannot Access the Administration Console
- Cannot Access a Server Application
- Server Will Not Start on Windows (Port Conflict)
- Port Conflicts Debugging Multiple Instances on the Same Server
- Two Server Instances Bind to Same Port on Windows
- Error: "System cannot find the path specified"
- Application Generates `java.lang.UnsupportedOperationException`
- Using `asadmin set` Command May Produce Unexpected Results
- Cannot Undeploy Or Redeploy Application With Open Streams to Jar Files (Windows Only)
- Cannot Reinstall Application Server After Manual Deletion of Directories
- Cannot Produce a JVM Thread Dump After Server Crash

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..."](#)
- [Error: "Could Not Start Domain"](#)

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 Application Server 8. 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 Application 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 Application Server 8 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](#)
- [Restarting Automatically on the Microsoft Windows Platform](#)

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 Application Server 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 Application Server 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 Java™ System Application 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_scm_lite.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= "SJSAS_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.
 - a. In the Services Control Panel, double-click the service you created.

- b. In the Properties window, click the Log On tab.
- c. 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:

- a. 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
```

- b. 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
```

- c. 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 Application 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 Application 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`

- Windows installation/uninstallation:

`%TEMP%`

Log file names are different for each product distribution, but are platform independent:

- Platform Edition, both SDK and Application Server distributions, main installation/uninstallation log files:

`Install_Application_Server_8PE_<timestamp>.log`
`Uninstall_Application_Server_8PE_<timestamp>.log`

- Platform Edition, Application Server only distribution, low-level installation/uninstallation log files:

`Sun_Java_System_Application_Server_Platform_Edition_install.B<timestamp>`
`Sun_Java_System_Application_Server_Platform_Edition_uninstall.B<timestamp>`

- Platform Edition, SDK distribution, low-level installation/uninstallation log files:

`Java_2_Platform_Enterprise_Edition_1.4_SDK_install.B<timestamp>`
`Java_2_Platform_Enterprise_Edition_1.4_SDK_uninstall.B<timestamp>`

- Enterprise Edition, main installation/uninstallation log files:
 Install_Application_Server_8EE_<timestamp>.log
 Uninstall_Application_Server_8EE_<timestamp>.log
- Enterprise Edition, low-level installation/uninstallation log files:
 Sun_Java_System_Application_Server_Enterprise_Edition_install.B<timestamp>
 Sun_Java_System_Application_Server_Enterprise_Edition_uninstall.B<timestamp>

Accessing Local Server Fails (http://localhost:8080)

Things to check for this error include the following:

- [Did the Server Start?](#)
- [Was the Server Started at the Expected Port?](#)

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	<p>The server's default port number is 8080, however, there are a number of ways in which the expected value can change:</p> <ul style="list-style-type: none">• 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?](#)
- [Is the Proxy Setting Causing a Problem?](#)

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

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 19](#).

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 hostname and domain for the localhost machine: <ul style="list-style-type: none">• 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 <code>hostname</code> 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?](#)
- [Is the Administration Console Running on the Expected Port?](#)
- [Is the Security Manager Disabled?](#)

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 19](#) 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 20](#) 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

According to the *J2EE 1.4 Specification*, the Security Manager is not optional; it must be enabled in the Application Server. Since there is no configuration interface in the Application 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 Application Server, some things to check include the following:

- [Is the Application Server Running?](#)
- [Was Application Deployment Successful?](#)

Is the Application Server Running?

Description

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

Solution

Review the information in [“Did the Server Start?” on page 19](#) 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
```

Server Will Not Start on Windows (Port Conflict)

If a message similar to the following is displayed when starting the Application 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 Application Server did not shut down cleanly.

Other things to check include the following:

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

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 Application 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 Application 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.

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 Application Server EE (not PE) software, and is due to a known Windows security flaw rather than a problem with the Application Server itself.

The problem occurs when two or more instances of the Application 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"

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 Application 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`

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.

Example

```
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 Application Server After Manual Deletion of Directories

Description

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

Solution 1

Remove Application 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 Application Server in a different directory.

Cannot Produce a JVM Thread Dump After Server Crash

Description

If the Application 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

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

HADB Problems

This section covers problems you may encounter when using the Application Server 8.1 EE 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 Application Server.

Topics in this chapter include:

- [HADB Database Creation Fails](#)
- [Server Responds Slowly After Idle Period](#)
- [Requests Are Not Succeeding](#)
- [Session Persistence Problems](#)
- [HADB Performance Problems](#)
- [High Load Problems](#)
- [Connection Problem Caused by Lack of Semaphore Resources](#)
- [Improving CPU Utilization](#)
- [HADB Administration Problems](#)
- [Shared Memory Segment Key Already Exists \(Windows only\)](#)
- [Failure in configure-ha-cluster](#)
- [Unable to Run configure-ha-cluster](#)
- [hadbm set Command Fails](#)
- [Failure in configure-ha-cluster: Creating an HADB Instance Fails](#)

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](#)
- [Too Few Semaphores](#)
- [Database Nodes Cannot Be Reached and the Database Does Not Function](#)
- [The Management Agents Could Not Establish a Domain](#)
- [hadbm create or hadbm addnodes Command Hangs](#)
- [ma \(Management Agent Process\) Crashes](#)

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 Java™ System Application Server 8.1 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 Java™ System Application 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.mainternal.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 appserver 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?](#)
- [Are the System Clocks Synchronized?](#)
- [Is the Application Server Communicating With HADB?](#)

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 Application 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 Application Server Communicating With HADB?

Description

HADB may be created and running, but if the persistence store has not yet been created, the Application 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](#)
- [Configuring Instance-Level Session Persistence Did Not Work](#)
- [Session Data Seems To Be Corrupted](#)

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?](#)
- [Is There Disk Contention?](#)
- [Is There a Shortage of HADB Data Device Space?](#)
- [Is There a Shortage of Other HADB Resources?](#)

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* msecs).”

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 RH 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.”](#)

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?](#)
- [Is the node-internal Log Full?](#)
- [Are There Enough Locks?](#)
- [Can You Fix the Problem by Doing Some Performance Tuning?](#)

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.”](#)

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?](#)

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.”](#)

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?](#)

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	hadbm attribute DataBufferPoolSize
Size of Tuple Log Buffer	hadbm attribute LogBufferSize
Size of Node Internal Log Buffer	hadbm attribute InternalLogBufferSize
Number of Database Locks	hadbm attribute NumberOfLocks

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:
```

You must configure more semaphore unso structures on the host computer. See the “Preparing for HADB Setup” chapter in the *Sun Java™ System Application Server 8.1 Installation Guide* for information on how to configure `semmnu` 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.)

Solutions

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 Java™ System Application 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](#)
- [hadbm Command Fails: command not found](#)
- [hadbm Command Fails: JAVA_HOME not defined](#)
- [hadbm createdomain fails, but two split domains are created](#)
- [create Fails: "path does not exist on a host"](#)
- [Database Does Not Start](#)
- [clear Command Failed](#)

- [create-session-store Failed](#)
- [hadbm Command Hangs](#)
- [Cannot Restart the HADB](#)

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 hostname1
ping 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 Java™ System Application Server 8.1 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 Java™ System Application Server 8.1 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?](#)
- [Do the History Files Contain Errors?](#)

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”](#) 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.
- [Contacting Sun Technical Support](#).

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`
- `SQLException: No suitable driver`

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` system 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 Application 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 confutation, 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...
```

Solution 2

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:

```
HADBGMGT007: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.

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

Security Problems

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

- [java.security.AccessControlException: Access Denied Error](#)
- [javax.ejb.AccessLocalException: Client Not Authorized Error](#)
- [Authentication is Not Working With the Solaris Realm](#)
- [Mutual Authentication Not Working With the Application Client](#)

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 Application 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
```

Mutual Authentication Not Working With the Application Client

Frequently Asked Questions

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

- [What Happens When No Server Side Realm is Configured?](#)
- [Can I Use a Pkcs12 Certificate for My Client Certificate?](#)
- [Can I See the TLS/SSL Handshake Information for an SSL Client?](#)
- [Can I Change the Keystore Password?](#)
- [How Do I Maintain a Session in JAX-RPC?](#)
- [How Do I Access the Naming Service From a Standalone Java Client?](#)
- [Are RMI-IIOP Stubs Needed to Access Remote EJBs?](#)
- [How Do I Change the Log Level for an Application Logger?](#)

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.

For more information on the `HttpSession` object, see <http://java.sun.com/j2ee/1.4/docs/api/javax/servlet/http/HttpSession.html>

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 Application 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 Application Server naming service. This removes a step that was required in previous releases. (See [Are RMI-IIOP Stubs Needed to Access Remote EJBs?](#) 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 Application 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 Application Server's naming service code. This often means that the client cannot take advantage of desirable value-added behavior built in to the Application 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 Application 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 Application 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 Application 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 Application 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

```
--generateterminstubs=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 [7](#)
- Administration Console
 - cannot access [22](#)
- application
 - cannot access [24](#)

C

- cannot access Administration Console [22](#)
- cannot access local server [19](#)
- cannot access remote server [20](#)
- cannot access server application [24](#)
- Cannot find log files [17](#)

D

- diagnosis of a problem [2](#)
- documentation [8](#)
- domain
 - finding [22](#)
 - restarting [13](#)

E

- error messages [5](#)

F

- forum [9](#)

H

- HADB [31](#)
 - setting the environment [49](#)
- HADB will not start [55](#)
- hadbm command not working [49](#)
- high-load situation [44](#)
- hostname
 - finding [22](#)
- HTTP server access logs [17](#)

I

- investigation of a problem [4](#)

J

JAVA_HOME setting [50](#)

K

Knowledge Base [8](#)

L

Linux, default base directory for [ix](#)

local server

access fails [19](#)

log files [6](#)

log levels [7](#)

logs

cannot find [17](#)

directory [18](#)

M

message types [17](#)

N

node agent, restarting [13](#)

O

online forum [9](#)

P

pkginfo [2](#)

port number

how the default value can change [20](#)

power failure and HADB [55](#)

R

remote server

access fails [20](#)

response-timeout-in-seconds [37](#)

restarting server [13](#)

rpm [2](#)

S

Server

log file, description of [17](#)

server

how default port number can change [20](#)

log file, location of [18](#)

will not start on windows [24](#)

server.log [7](#)

showrev [2](#)

Solaris

patches [xii](#)

support [xii](#)

start-domain command fails [12](#)

SunSolve [8](#)

support

Solaris [xii](#)

T

thread dump, obtaining [7](#)

U

ungraceful shutdown [55](#)

W

Windows

server will not start [24](#)

