

# Troubleshooting Guide

*Sun™ ONE Application Server*

**Version 7, Enterprise Edition**

817-2161-10  
September 2003

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

Copyright © 2003 Sun Microsystems, Inc. All rights reserved.

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

UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

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.

---

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

CE LOGICIEL 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, ÉCRITE ET PRÉALABLE DE SUN MICROSYSTEMS, INC. Droits du gouvernement américain, utilisateurs gouvernementaux - logiciel commercial. Les utilisateurs gouvernementaux sont soumis au contrat de licence standard de Sun Microsystems, Inc., ainsi qu'aux dispositions en vigueur de la FAR (Federal Acquisition Regulations) et des suppléments à celles-ci. 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, Sun™ ONE, le logo Java Coffee Cup et le logo Sun™ ONE sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

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.

# Contents

|  |           |
|--|-----------|
| <b>About This Guide</b> .....                            | <b>11</b> |
| Who Should Use This Guide .....                          | 11        |
| How This Guide is Organized .....                        | 12        |
| Using the Documentation .....                            | 13        |
| Documentation Conventions .....                          | 15        |
| General Conventions .....                                | 15        |
| Conventions Referring to Directories .....               | 16        |
| Product Support .....                                    | 17        |
| <br>   |           |
| <b>Chapter 1 Troubleshooting Approach</b> .....          | <b>19</b> |
| Identifying the Problem .....                            | 19        |
| Seeking a Solution .....                                 | 21        |
| Interpreting Messages .....                              | 21        |
| Types of Messages .....                                  | 21        |
| Error Messages .....                                     | 22        |
| Examining Log Files .....                                | 23        |
| Log Levels .....   | 23        |
| Log Options .....  | 24        |
| Client Side Logging .....                                | 24        |
| Obtaining a Thread Dump .....                            | 24        |
| Determining If This Problem Has Been Solved Before ..... | 25        |
| Searching the Product Documentation .....                | 25        |
| Searching the Knowledge Base .....                       | 26        |
| Searching or Participating in the Online Forum .....     | 26        |
| .....  | 26        |

|  |               |
|--|---------------|
| <b>Chapter 2 Installation and Uninstallation Problems</b>                              | <b>27</b>     |
| Install/Uninstall Logs   | 28            |
| Can't install remotely using the graphical interface.                                  | 29            |
| Can't reinstall the server.  | 29            |
| Silent installation is not working correctly.  | 31            |
| Is the silent installation configuration file correct?                                 | 31            |
| Uninstallation failure needs cleanup.  | 31            |
| Can't install the load balancer plug-in.   | 31            |
| Is your web server installed?  | 32            |
| Is there a previously installed load balancer or reverse proxy plug-in on your system? | 32            |
| Has the load balancer plug-in already been installed?                                  | 32            |
| Are the configuration files correct?   | 32            |
| Load balancer won't start.   | 33            |
| Shared memory creation failed.   | 33            |
| Have you configured shared memory?   | 34            |
| Is there an error in your /etc/system file?  | 34            |
| Did you reboot the machine after configuring shared memory?                            | 34            |
| clsetup is not working.  | 35            |
| Have you configured shared memory?   | 35            |
| Has remote communication been set up correctly?  | 35            |
| Under SSH, are the HADB and the Application Server co-located on the same machine?     | 36            |
| Are the application server and HADB installed in the same directories on each machine? | 36            |
| Are all the Admin Servers on the application server instances in the cluster running?  | 37            |
| Are the input files on all instances in the cluster identical?                         | 37            |
| Unable to create HADB database using clsetup.  | 37            |
| Problems when running clsetup as non-root.   | 38            |
| Insufficient space.  | 38            |
| Is the number of semaphores too low?   | 38            |
| Can't test the ssh setting as root.  | 39            |
| Can't get ssh to skip the login prompt.  | 39            |
| <br><b>Chapter 3 Startup and Login Problems</b>  | <br><b>41</b> |
| Server does not appear to be running.  | 41            |
| Is the Application Server running?   | 42            |
| During installation, did the initial server startup run successfully?                  | 43            |
| Is the server available locally?   | 43            |
| Was the server started at the expected port?   | 44            |
| Is your proxy setting causing a problem?   | 45            |
| Has an ungraceful shutdown occurred on a previously-running server?                    | 46            |
| Can't access the Admin Server.   | 46            |
| Has the Admin Server been started?   | 47            |
| Are you the user who installed the Application Server?                                 | 47            |
| Is the Admin Server running at the expected port?                                      | 47            |

|   |           |
|---|-----------|
| Can't access a server application. ....   | 48        |
| Is the Application Server running? ....   | 48        |
| Is the application enabled? ....  | 48        |
| Was application deployment successful? ....   | 49        |
| Forgot the user name or password. ....  | 50        |
| Forgot the admin Server port number. ....   | 50        |
| Server won't start (CGI error). ....  | 50        |
| Are limits on file descriptors adequate? ....   | 51        |
| Are kernel parameters adequate? ....  | 51        |
| Restart operation fails ....  | 51        |
| SSL/TLS are enabled ....  | 51        |
| JMS failed to start. ....   | 52        |
| Are you attempting to start the instance as a non-root user? ....                       | 52        |
| Do Solaris bundled and unbundled domains and instances have the same names? ....        | 53        |
| Do the imq logs have out of memory errors? ....   | 53        |
| <br><b>Chapter 4 Runtime Problems</b> ....  | <b>55</b> |
| Runtime Logs ....   | 56        |
| Load Balancer / Web Server won't start ....   | 56        |
| Instance goes unused after restarting ....  | 57        |
| Can't access a web page. ....   | 57        |
| Can't access an application. ....   | 57        |
| Is the application deployed? ....   | 58        |
| Is your loadbalancer.xml file correct? ....   | 58        |
| Is the web server running? ....   | 58        |
| Has the correct port been specified for the web server? ....                            | 59        |
| Server responds slowly after being idle ....  | 59        |
| Does the log contain "Lost connection" messages? ....                                   | 59        |
| My application suddenly went away. ....   | 59        |
| Is the application you are using being quiesced by the load balancer? ....              | 60        |
| Requests are not succeeding. ....   | 60        |
| Is the load balancer timeout correct? ....  | 60        |
| Are the system clocks synchronized? ....  | 61        |
| Have you enabled the instances of the cluster? ....                                     | 61        |
| Is the AppServer communicating with HADB? ....  | 61        |
| Server log: app persistence-type = memory ....  | 62        |
| Dynamic reconfiguration failed. ....  | 62        |
| Session Persistence Problems ....   | 63        |
| The create-session-store command failed. ....   | 63        |
| Are the HADB and the application server instance on different sides of a firewall? .... | 63        |
| Configuring instance-level session persistence didn't work. ....                        | 63        |
| Session data seems to be corrupted. ....  | 64        |
| HTTP session failover is not working. ....  | 64        |

|   |           |
|---|-----------|
| Are the system clocks synchronized? .....   | 65        |
| Do all objects bound to a distributed session implement the java.io.Serializable interface? ..... | 65        |
| Is session information stored in HttpSession? .....   | 65        |
| Is your web application distributable? .....  | 66        |
| Is the persistence type set to ha? .....  | 66        |
| An object is cloned instead of shared .....   | 66        |
| Has a session store been created? .....   | 67        |
| Are all the machines in the cluster homogenous? .....   | 67        |
| Has high availability been enabled? .....   | 67        |
| Out of Memory and Stack Overflow errors. ....   | 67        |
| HADB Transaction Failures .....   | 68        |
| Is there a shortage of HADB data devices space? .....   | 68        |
| Is there a shortage of other HADB resources? .....  | 69        |
| Have history files grown too large? .....   | 70        |
| Other Transaction Problems .....  | 70        |
| Can't set TRANSACTION_SERIALIZABLE level for a bean. ....   | 70        |
| Sporadic failures during high loads. ....   | 70        |
| No space left in Tuple Log .....  | 71        |
| Node-internal log gets full .....   | 71        |
| Not enough locks .....  | 72        |
| Frequent "High Load" Warnings .....   | 73        |
| Client cannot connect to HADB. ....   | 73        |
| Connection Queue Problems .....   | 74        |
| Full connection queue closes socket .....   | 74        |
| Connection Pool Problems .....  | 74        |
| Single sign-on requires larger connection pool. ....  | 75        |
| Server: Unable to obtain connection from pool .....   | 76        |
| JDBC connection is not available. ....  | 76        |
| Is the max-pool-size setting adequate? .....  | 77        |
| Exception occurs while calling DataSource.getConnection(). ....                                   | 77        |
| Exception occurs while executing against MSSQL. ....  | 78        |
| IOException: Connection in invalid state .....  | 78        |
| <b>Chapter 5 Administration Problems .....</b>  | <b>81</b> |
| Server Logs .....   | 81        |
| Application Server logs .....   | 82        |
| HADB History Files .....  | 82        |
| HADB logs .....   | 83        |
| Command-Line Interface Problems .....   | 83        |
| Can't access the command-line utility. ....   | 84        |
| Can't access the Application Server man pages. ....   | 84        |
| Graphical Interface Problems .....  | 84        |
| Can't access the Administration interface. ....   | 85        |

|   |     |
|---|-----|
| Can't undo accidental "changes."  | 85  |
| Monitoring Problems   | 85  |
| Load Balancer Plug-in isn't being monitored   | 85  |
| Authentication/Authorization Problems   | 86  |
| Can't import the certificate for my server.   | 86  |
| Has the trust database been created?  | 86  |
| Was the certificate generated with the right tool?                                  | 87  |
| The server does not recognize my certificate.                                       | 87  |
| LDAP authentication/equalization is not working.                                    | 87  |
| HADB Administration Problems  | 88  |
| hadbm command fails: host unreachable.  | 89  |
| Is the host up and running?   | 89  |
| Is RSH or SSH set up and running?   | 89  |
| Are the SSH binaries in the proper location?  | 90  |
| hadbm command fails: command not found  | 90  |
| hadbm command fails: JAVA_HOME not defined  | 90  |
| create fails: "path does not exist on a host"                                       | 91  |
| database doesn't start.   | 91  |
| Was there a shared memory get segment failure?                                      | 91  |
| Do the history files contain errors?  | 92  |
| Do you need a simple solution?  | 94  |
| clear command failed  | 94  |
| create-session-store failed   | 94  |
| Invalid user name or password   | 94  |
| SQLException: No suitable driver  | 95  |
| "No space left on device" appears in server.log                                     | 95  |
| node status is "starting" after issuing stopnode.                                   | 95  |
| Node failure occurred.  | 96  |
| Double node failure occurred.   | 97  |
| Can't restart the HADB after an ungraceful shutdown.                                | 97  |
| hadbm command doesn't return control to user.                                       | 98  |
| Cluster Administration Problems   | 98  |
| Refragmentation of the HADB fails.  | 99  |
| Is there enough space on the data devices?  | 99  |
| The cladmin command is not working.   | 100 |
| Are the Admin Servers of all the instances in the cluster started?                  | 100 |
| Do all the instances in the cluster have same administrator user name and password? | 100 |
| Are the input files correct?  | 101 |
| Are the input files on all instances in the cluster identical?                      | 101 |
| Application is not available on the cluster.  | 101 |
| Did the application deploy successfully to the cluster?                             | 101 |
| Common Administration and Recovery Actions  | 101 |
| Examine the history files   | 102 |

|   |            |
|---|------------|
| Maintain service while taking HADB offline .....                                    | 102        |
| Clear the database and recreate session store .....                                 | 103        |
| <b>Chapter 6 Deployment Problems .....</b>  | <b>105</b> |
| EJB Problems .....  | 105        |
| CMP Problems .....  | 105        |
| Are you using a supported JDBC driver? .....  | 105        |
| Ant Problems .....  | 105        |
| Can't run Ant remotely. ....  | 106        |
| <b>Chapter 7 Integration Problems .....</b>   | <b>107</b> |
| Web Server Problems .....   | 107        |
| Sun ONE Web Server Problems .....   | 107        |
| Apache Web Server Problems .....  | 108        |
| Reverse Proxy Plug-In Problems .....  | 108        |
| Messaging Problems .....  | 108        |
| JMS broker won't start. ....  | 108        |
| Can't reestablish Sun ONE Message Queue broker in a cluster environment. ....       | 108        |
| IDE Problems .....  | 109        |
| Application Server is not in the Sun ONE Studio Server Registry. ....               | 109        |
| Is the Application Server plug-in for Sun ONE Studio installed? .....               | 109        |
| Is the plug-in configured with the Application Server installation directory? ..... | 109        |
| Is the Application Server registered? .....   | 109        |
| <b>Chapter 8 Migration Problems .....</b>   | <b>111</b> |
| Migration Toolbox doesn't run correctly. ....                                       | 111        |
| Is the script customized for your environment? .....                                | 112        |
| Are JAR files in your extension directory? .....                                    | 112        |
| Are you using the correct JDK? .....  | 112        |
| Non-fatal error occurs during extraction. ....                                      | 112        |
| Are you doing a NetDynamics migration? .....  | 113        |
| Is manual migration warranted? .....  | 113        |
| Fatal error occurs during extraction. ....  | 113        |
| Translation error occurs. ....  | 114        |
| Trouble migrating from WebLogic. ....   | 114        |
| <b>Chapter 9 Upgrade Problems .....</b>   | <b>115</b> |
| Upgrade Log .....   | 115        |
| Package-based installation consistently fails. ....                                 | 116        |
| User ID errors occur. ....  | 116        |
| Errors occur during the backup. ....  | 116        |
| Errors occur during reconfiguration. ....   | 116        |

|  |            |
|--|------------|
| Upgrade interrupted; server is in inconsistent state. ....                                       | 117        |
| Doing a partial online upgrade of a cluster causes problems. ....                                | 117        |
| <b>Appendix A Summary of High Availability Commands</b> .....                                    | <b>119</b> |
| <b>Appendix B Frequently Asked Questions (FAQs)</b> .....  | <b>121</b> |
| Installation and Initial Access .....  | 121        |
| How do I obtain version information? .....   | 121        |
| How do I access the sample applications? .....   | 122        |
| Licensing .....  | 122        |
| How do I renew an expired evaluation license? .....  | 122        |
| Can I upgrade from an evaluation license to a full Sun ONE Application Server 7 license? ...     | 122        |
| In a media package, how do I get the license key? .....  | 123        |
| How can I get an evaluation license for Enterprise Edition? .....                                | 123        |
| Platform Support .....   | 123        |
| What's the difference between Platform Edition, Standard Edition, and Enterprise Edition? ..     | 124        |
| What Sun ONE Message Queue versions are supported? .....   | 124        |
| What J2SE versions are supported? .....  | 125        |
| What database drivers are supported? .....   | 125        |
| Is JADO supported for Sun ONE Application Server 7? .....  | 125        |
| Does Sun ONE Application Server 7 have a JCA connector to talk to SAP R3? .....                  | 125        |
| Are custom authorization providers supported? .....  | 126        |
| How is load balancing supported? .....   | 126        |
| Functional Support .....   | 126        |
| Can remote EJB clients use something other than RMI over IIOP, or is that the only support       |            |
| protocol? .....  | 127        |
| Is sticky load balancing supported? .....  | 127        |
| Are multiple JVMs for hosting components supported? .....  | 128        |
| High-Availability .....  | 128        |
| Where are the load balancer files? .....   | 128        |
| How do I enable the health checker? .....  | 128        |
| Server Administration and Operation .....  | 129        |
| How do I start a single application server instance rather than all the instances in a domain? . | 129        |
| What is the 'Apply Changes Required' message about? .....  | 129        |
| Why are multiple IIOP listeners useful and how many listeners are allowed? .....                 | 130        |
| Does an application server instance have to be running in order to deploy to it? .....           | 130        |
| Does the Admin Server need to be running in order for me to run my application? .....            | 130        |
| Can I configure Sun ONE Application Server 7 to run as non-root? .....                           | 131        |
| Application Debugging .....  | 131        |
| Can I modify roles in the web or EJB deployment descriptors without restarting the application   |            |
| server? .....  | 131        |
| Upgrade/Migration .....  | 131        |
| Can I upgrade from other versions of Sun ONE Application Server 7 to Enterprise Edition? ..      | 132        |

**Index .....133**

# About This Guide

This *Troubleshooting Guide* provides instructions for resolving common problems with the Sun™ Open Net Environment (Sun ONE) Application Server 7, Enterprise Edition product. Many issues that also affect the Standard Edition product are also included.

The following topics are addressed in the preface:

- [Who Should Use This Guide](#)
- [How This Guide is Organized](#)
- [Using the Documentation](#)
- [Documentation Conventions](#)
- [Product Support](#)

## Who Should Use This Guide

This manual is intended to be used by anyone who is charged with solving failure-to-operate problems with the Sun ONE Application Server 7 software. This can include technical support engineers, system administrators, network administrators, application server administrators, and developers who are responsible for restoring operational functionality.

---

|             |   |
|-------------|---|
| <b>NOTE</b> | This document does not address performance tuning, optimization, or best practices for using the Application Server. It focuses on the problems that are most likely to occur during operation, and provides guidelines for restoring operation as quickly as possible. |
|-------------|---|

---

This guide assumes the reader is familiar with the following:

- UNIX<sup>®</sup> operating system
- Client/server programming model
- Java programming language
- How to install enterprise-level software products
- How Sun ONE Application Server 7 works
- High-availability and clustering concepts
- Internet and World Wide Web

## How This Guide is Organized

This guide contains the following chapters and appendixes:

- [Chapter 1, “Troubleshooting Approach”](#)
- [Chapter 2, “Installation and Uninstallation Problems”](#)
- [Chapter 3, “Startup and Login Problems”](#)
- [Chapter 4, “Runtime Problems”](#)
- [Chapter 5, “Administration Problems”](#)
- [Chapter 6, “Deployment Problems”](#)
- [Chapter 7, “Integration Problems”](#)
- [Chapter 8, “Migration Problems”](#)
- [Chapter 9, “Upgrade Problems”](#)
- [Appendix A, “Summary of High Availability Commands”](#)
- [Appendix B, “Frequently Asked Questions \(FAQs\)”](#)

# Using the Documentation

The Sun ONE Application Server 7, 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 ONE Application Server manuals.

**Table 1** Sun ONE Application Server Documentation Roadmap

| For information about  | See the following  |
|--|--|
| Late-breaking information about the software and the documentation.  | <i>Release Notes</i>   |
| Comprehensive, table-based summary of supported hardware, operating system, JDK, and JDBC/RDBMS.   | <i>Platform Summary</i>  |
| Sun ONE Application Server 7 overview, including the features available with each product edition.   | <i>Product Overview</i>  |
| Diagrams and descriptions of server architecture and the benefits of the Sun ONE Application Server architectural approach.  | <i>Server Architecture</i>   |
| New enterprise, developer, and operational features of Sun ONE Application Server 7.   | <i>What's New</i>  |
| How to get started with the Sun ONE Application Server 7 product. Includes a sample application tutorial.  | <i>Getting Started Guide</i>   |
| Installing the Sun ONE Application Server software and its components, such as sample applications, the Administration interface, and the high-availability components. Instructions for implementing a basic high-availability configuration are included.  | <i>Installation Guide</i>  |
| Evaluating your system needs and enterprise to ensure that you deploy Sun ONE 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>System Deployment Guide</i>   |
| Best practices for HTTP session availability that application architects and developers can use.   | <i>Application Design<br/>Guidelines for Storing Session<br/>State</i> |
| Creating and implementing Java™ 2 Platform, Enterprise Edition (J2EE™ platform) applications intended to run on the Sun ONE Application Server 7 that follow the open Java standards model for J2EE components such as servlets, Enterprise JavaBeans™ (EJBs™), and JavaServer Pages™ (JSPs™). Includes general information about application design, developer tools, security, assembly, deployment, debugging, and creating lifecycle modules. A comprehensive Sun ONE Application Server glossary is included. | <i>Developer's Guide</i>   |

**Table 1** Sun ONE Application Server Documentation Roadmap (*Continued*)

| For information about  | See the following  |
|--|--|
| Creating and implementing J2EE web applications that follow the Java™ Servlet and JavaServer Pages (JSP) specifications on the Sun ONE Application Server 7. Discusses web application programming concepts and tasks, and provides sample code, implementation tips, and reference material. Topics include results caching, JSP precompilation, session management, security, deployment, SHTML, and CGI.            | <i>Developer's Guide to Web Applications</i>                         |
| Creating and implementing J2EE applications that follow the open Java standards model for enterprise beans on the Sun ONE Application Server 7. Discusses Enterprise JavaBeans (EJB) programming concepts and tasks, and provides sample code, implementation tips, and reference material. Topics include container-managed persistence, read-only beans, and the XML and DTD files associated with enterprise beans. | <i>Developer's Guide to Enterprise JavaBeans Technology</i>          |
| Creating Application Client Container (ACC) clients that access J2EE applications on the Sun ONE Application Server 7.   | <i>Developer's Guide to Clients</i>                                  |
| Creating web services in the Sun ONE Application Server environment.   | <i>Developer's Guide to Web Services</i>                             |
| Java™ Database Connectivity (JDBC™), transaction, Java Naming and Directory Interface™ (JNDI), Java™ Message Service (JMS), and JavaMail™ APIs.  | <i>Developer's Guide to J2EE Services and APIs</i>                   |
| Creating custom NSAPI plug-ins.  | <i>Developer's Guide to NSAPI</i>                                    |
| Information and instructions on the configuration, management, and deployment of the Sun ONE Application Server subsystems and components, from both the Administration interface and the command-line interface. Topics include cluster management, the high-availability database, load balancing, and session persistence. A comprehensive Sun ONE Application Server glossary is included.                         | <i>Administrator's Guide</i>   |
| Editing Sun ONE Application Server configuration files, such as the <code>server.xml</code> file.  | <i>Administrator's Configuration File Reference</i>                  |
| Configuring and administering security for the Sun ONE Application Server operational environment. Includes information on general security, certificates, and SSL/TLS encryption. HTTP server-based security is also addressed.   | <i>Administrator's Guide to Security</i>                             |
| Configuring and administering service provider implementation for J2EE™ Connector Architecture (CA) connectors for the Sun ONE Application Server 7. Topics include the Administration Tool, Pooling Monitor, deploying a JCA connector, and sample connectors and sample applications.  | <i>J2EE CA Service Provider Implementation Administrator's Guide</i> |
| Migrating your applications to the new Sun ONE Application Server 7 programming model, specifically from iPlanet Application Server 6.x and from Netscape Application Server 4.0. Includes a sample migration.   | <i>Migrating and Redeploying Server Applications Guide</i>           |

**Table 1** Sun ONE Application Server Documentation Roadmap (*Continued*)

| For information about  | See the following  |
|--|--|
| How and why to tune your Sun ONE Application Server to improve performance.  | <i>Performance Tuning Guide</i>  |
| Information on solving Sun ONE Application Server problems.  | <i>Troubleshooting Guide</i>   |
| Messages that you may encounter while running Sun ONE Application Server 7. Includes a description of the likely cause and guidelines on how to address the condition that caused the message to be generated. | <i>Error Message Reference</i>   |
| Utility commands available with the Sun ONE Application Server; written in manpage style.  | <i>Utility Reference Manual</i>  |
| Using the Sun™ Open Net Environment (Sun ONE) Message Queue software.  | The Sun ONE Message Queue documentation at:<br><br><a href="http://docs.sun.com/db?p=prod/sl.slmsgqu">http://docs.sun.com/db?p=prod/sl.slmsgqu</a> |

## Documentation Conventions

This section describes the types of conventions used throughout this guide:

- [General Conventions](#)
- [Conventions Referring to Directories](#)

### General Conventions

The following general conventions are used in this guide:

- **File and directory paths** are given in UNIX® format (with forward slashes separating directory names).
- **URLs** are given in the format:

`http://server.domain/path/file.html`

In these URLs, *server* is the server name where applications are run; *domain* is your Internet domain name; *path* is the server's directory structure; and *file* is an individual filename. Italic items in URLs are placeholders.

- **Font conventions** include:

- The `monospace` font is used for sample code and code listings, API and language elements (such as function names and class names), file names, pathnames, directory names, and HTML tags.
- *Italic* type is used for code variables.
- *Italic* type is also used for book titles, emphasis, variables and placeholders, and words used in the literal sense.
- **Bold** type is used as either a paragraph lead-in or to indicate words used in the literal sense.
- **Installation root directories** for most platforms are indicated by *install\_dir* in this document. Exceptions are noted in [“Conventions Referring to Directories” on page 16](#).

By default, the location of *install\_dir* on **most** platforms is:

- Solaris™ 8 non-package-based Evaluation installations:

*user's home directory/sun/appserver7*

- Solaris unbundled, non-evaluation installations:

*/opt/SUNWappserver7*

For the platforms listed above, *default\_config\_dir* and *install\_config\_dir* are identical to *install\_dir*. See [“Conventions Referring to Directories” on page 16](#) for exceptions and additional information.

- **Instance root directories** are indicated by *instance\_dir* in this document, which is an abbreviation for the following:  
*default\_config\_dir/domains/domain/instance*
- **UNIX-specific descriptions** throughout this manual apply to the Linux operating system as well, except where Linux is specifically mentioned.

---

**NOTE** Forte for Java 4.0 has been renamed to Sun ONE Studio 4 throughout this manual.

---

## Conventions Referring to Directories

By default, when using the Solaris™ 8 and 9 package-based installation and the Solaris 9 bundled installation, the application server files are spread across several root directories. These directories are described in this section.

- **For Solaris 9 bundled installations**, this guide uses the following document conventions to correspond to the various default installation directories provided:
  - *install\_dir* refers to `/usr/appserver/`, which contains the static portion of the installation image. All utilities, executables, and libraries that make up the application server reside in this location.
  - *default\_config\_dir* refers to `/var/appserver/domains`, which is the default location for any domains that are created.
  - *install\_config\_dir* refers to `/etc/appserver/config`, which contains installation-wide configuration information such as licenses and the master list of administrative domains configured for this installation.
- **For Solaris 8 and 9 package-based, non-evaluation, unbundled installations**, this guide uses the following document conventions to correspond to the various default installation directories provided:
  - *install\_dir* refers to `/opt/SUNWappserver7`, which contains the static portion of the installation image. All utilities, executables, and libraries that make up the application server reside in this location.
  - *default\_config\_dir* refers to `/var/opt/SUNWappserver7/domains` which is the default location for any domains that are created.
  - *install\_config\_dir* refers to `/etc/opt/SUNWappserver7/config`, which contains installation-wide configuration information such as licenses and the master list of administrative domains configured for this installation.

## Product Support

If you have general feedback on the product or documentation, please send this to [appserver-feedback@sun.com](mailto:appserver-feedback@sun.com).

If you have problems with your system, contact customer support using one of the following mechanisms:

- The online support web site at:  
<http://www.sun.com/supporttraining/>
- The telephone dispatch number associated with your maintenance contract

Please have the following information available prior to contacting support. This helps to ensure that our support staff can best assist you in resolving problems:

- Description of the problem, including the situation where the problem occurs and its impact on your operation
- Machine type, operating system version, and product version, including any patches and other software that might be affecting the problem
- Detailed steps on the methods you have used to reproduce the problem
- Any error logs or core dumps

# Troubleshooting Approach

This chapter provides a description of the tools, methods, and information sources available for troubleshooting the Sun™ Open Net Environment (ONE) Application Server 7 product. Guidelines for evaluating and investigating a problem are included.

The following sections are contained in this chapter:

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

## Identifying the Problem

J2EE application servers are typically deployed in complex and highly sophisticated operating environments. The Sun ONE 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 help you classify the problem, so you can more easily search for a solution.

---

**NOTE** 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 Sun ONE Application Server? On the same machine or not?
7. How is the Sun ONE Application Server connected to the directory server?
8. Are Sun ONE Application Servers in a cluster or not?
9. Was any upgrade done? If so, what were source and target versions?
10. Was a migration done? If so, what were source and target versions?
11. Have any new applications been deployed?
12. Is SSL enabled or not?
13. What database version is 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. Are the interoperating component versions in compliance with the compatibility matrix specified in the release notes?

After gathering this information, you will want to do the following:

- Collect web server error and access log data (web server instance-specific). Refer to [“Server Logs” on page 81](#) for information on specific logs.
- Collect any Sun ONE Application Server stack traces.

---

**NOTE** 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 the user may have already taken in attempting to resolve the problem.

## Seeking a Solution

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

The following topics are addressed in this section:

- [Interpreting Messages](#)
- [Examining Log Files](#)
- [Determining If This Problem Has Been Solved Before](#)
- [Searching the Product Documentation](#)
- [Searching the Knowledge Base](#)
- [Searching or Participating in the Online Forum](#)

## Interpreting 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:

- **Information**—These messages mark normal completion of particular 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 status to be reported as Failed. Error messages generally provide detailed information about the nature and the cause of the problem that occurred.

## Error Messages

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

- In some cases, the message is very clear about what is wrong and what needs to be done to fix it. For example, you start up a domain using the `asadmin start-domain` command, then inadvertently issue the same command a few moments later, after the domain has been started. You get the following message:

```
userD:\Sun\studio5_se\appserver7\bin>asadmin start-domain
Domain already started : domain1
Domain domain1 Started.
```

In this case, the message gives clear guidance and the problem should be relatively easy to fix or can be disregarded.

- Sometimes an error message gives only general information on 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 as obvious, or there could be multiple things wrong. You will have to consider various possibilities and perhaps a number of solutions. If your proposed fix is time consuming or costly, you should take steps to ensure that the fix is likely to be correct before actually doing anything.

- Some error messages are either not helpful, or give you no 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:
HTTP/4144: error sending
D:/Sun/studio5_se/appserver7/domains/domain1/server1/applications/j2ee-modules/SupplierServiceClient1_1/SupplierServiceClient1_SOAP.html (Overlapped I/O operation
is in progress.) status=1:5
```

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

For descriptions of all the Application Server error messages, refer to the *Sun ONE Application Server Error Messages Reference* at the following location:

[http://docs.sun.com/db/coll/s1\\_asseul\\_en](http://docs.sun.com/db/coll/s1_asseul_en)

## Examining 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 Sun ONE 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.

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 ONE 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 Admin Server log file (`server.log`) and to tightly integrate with Solaris. The log levels ALERT and FATAL are specific to 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 ONE 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 7. 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 your server instance does not include the `-Xrs java-option` flag. Remove the `-Xrs java-option` flag if it exists.
2. If the option was changed, restart your server instance.
3. Run the following command on the application server instance:

```
kill -3 pid
```

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

## Determining If This 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 if you need to seek more details, explanations, or examples.

Additional resources for researching an existing problem are available within Sun Microsystems:

- Sun Microsystems call center database
- Sun Microsystems bugtraq database

## Searching the Product Documentation

It is idea to be familiar with all of the product documentation at your disposal. Skimming them quickly makes it easier to find things later and, more importantly, lets you know what's there so that when a problem occurs, it triggers a search in your "mental index" of the docs you've seen.

When you have a problem, make it a point to look through every potentially relevant document you can lay your hands on. Several great things happen in this process:

- You may well find the answer you need.
- You'll find other things you'll wind up needing later.
- You'll acquire the terminology you need to describe the problem accurately, and know when a recommendation to "read document X" won't be help.

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

The documentation for all the Application Server product releases is here:

<http://docs.sun.com/db/prod/sl.asse#hic>

All the information needed may be in one of the existing product manuals. Being familiar with the documentation in the Application Server doc collection will help you identify the relevant material more quickly. A description of the Application Server manuals can be found in ["Using the Documentation" on page 13](#).

## Searching the Knowledge Base

The Knowledge Base is a collection of articles on product issues that may provide information helpful to you in 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 Free SunONE Articles.
4. Click Next.
5. Enter your search criteria.
6. Click Go.

## Searching or Participating in the Online Forum

You can browse directly in any of the online forums, or log in and register to start posting your own messages. The Application Server online forum is here:

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

# Installation and Uninstallation Problems

The high-availability components of Sun™ Open Net Environment (ONE) Application Server 7, Enterprise Edition product include the HADB, the HADB Management Client, and the load balancer plug-in. During installation, these components can be installed with the rest of the Application Server components, or separately. The load balancer plug-in is *usually* installed separately from the Application Server components.

This chapter addresses problems that you may encounter while performing installation or uninstallation of the Sun ONE Application Server 7, Enterprise Edition product or its components or plug-ins.

The following sections are contained in this chapter:

- [Install/Uninstall Logs](#)
- [Can't install remotely using the graphical interface.](#)
- [Can't reinstall the server.](#)
- [Silent installation is not working correctly.](#)
- [Silent installation is not working correctly.](#)
- [Uninstallation failure needs cleanup.](#)
- [Can't install the load balancer plug-in.](#)
- [Load balancer won't start.](#)
- [Shared memory creation failed.](#)
- [clsetup is not working.](#)
- [Unable to create HADB database using clsetup.](#)

- Problems when running clsetup as non-root.
- Insufficient space.
- Can't test the ssh setting as root.
- Can't get ssh to skip the login prompt.

## Install/Uninstall Logs

The following Application Server logs can be useful for troubleshooting problems you may have with installation or uninstallation:

```
/var/sadm/install/logs/Sun_ONE_Application_Server_install.log
/var/sadm/install/logs/Sun_ONE_Application_Server_uninstall.log
```

Use the following log for troubleshooting problems with the `clsetup` command:

```
/var/tmp/clsetup.log
```

In addition to these log files, low-level installation and uninstallation log files are created at these locations:

```
/var/sadm/install/logs/Sun_ONE_Application_Server_install.<timestamp>
/var/sadm/install/logs/Sun_ONE_Application_Server_uninstall.<timestamp>
```

The following logs are associated with the high-availability components:

- Web server errors, including load balancer error messages, are written into the `web server error.log`.
- Application server messages are logged in the respective instance server's `server.log` file (the default location is `/var/opt/SUNWappserver7/domains/domain1/server1/logs`). This log includes admin server messages and deployment errors.
- Admin-server messages, including application deployment messages, are logged in the admin-server's `server.log` file (the default location is `/var/opt/SUNWappserver7/domains/domain1/admin-server/logs`).
- Database creation errors are written to `/var/sadm/install/logs/clsetup.log`.
- Initial cluster setup errors are written to `/var/tmp/clsetup.log`.
- Cluster administration errors are written to `/var/tmp/cladmin.log`.

Some guidelines on using logs:

- Set the value of the `require-monitor-data` property to true in the `loadbalancer.xml` file in order to see monitoring details in the log.
- The `UnhealthyInstances` messages that appear in the log should be particularly helpful in troubleshooting load balancer problems.
- Using `hadbm` to set a large tuple log size will increase performance of the logging facility.
- The `cladmin.log` file may be useful in troubleshooting cluster administration.
- The `clsetup.log` file may be helpful in finding out what went wrong during installation when you establish a new cluster.

Device directory location: `/var/opt/SUNWhadb`

Configuration files location: `/etc/opt/SUNWhadb/dbdef`

## Can't install remotely using the graphical interface.

On UNIX, if you are installing the Application Server software remotely using the graphical interface, you must enable the display configuration on the machine where you are installing the product.

### *Solution*

Set the `DISPLAY` environment variable to contain the name of the server and domain, using this format:

```
host_name.domain_name.com:0.0
```

Then run the following command on the remote client:

```
xhost +
```

## Can't reinstall the server.

If installation and uninstallation are performed according to the documented instructions and they complete normally, you will be able to reinstall the server with no problems. However, if you have used another method to remove the Application Server files, or if there as been a failure during installation or

uninstallation, the system might be in an inconsistent state, leaving behind files or processes specific to the Application Server in the `/var/sadm/install/productregistry` file. These leftover files and processes will provoke an error message similar to the following on a subsequent installation:

```
UnsatisfiedLinkError: Can't load library: libinstallCore.so
```

You will need to clean up these files or processes before attempting a new installation.

*Solution: Clean up leftover files and processes*

1. Log in as root.
2. Navigate to your installation directory and check the content of the `/var/sadm/install/productregistry` file for installed packages, that is, files having the SUNW string. For example:

```
cat /var/sadm/install/productregistry | grep SUNW
```

3. Run `pkgrm` for the SUNW packages that were found in the product registry. For example:

```
pkgrm SUNWasaco
```

4. Remove the following files, if they are present:

```
/tmp/setupSDKNative
```

```
/tmp/SolarisNativeToolkit_3.0_1
```

5. After the packages have been removed, use the `prodreg` registry editor to remove the Application Server-specific entries.
6. At the command line, kill all `appservd` processes that may be running by typing the following:

```
ps -ef | grep appservd
```

```
pkill appservd
```

7. Remove all remaining files under the Sun ONE Application Server installation directories. Refer to [“Conventions Referring to Directories” on page 16](#) for further information and bundled and unbundled structures.
8. Remove the Sun ONE Application Server 7 directories.
  - Installation directory—by default, `/opt/SUNWappserver7`
  - Configuration directory—by default, `/etc/opt/SUNWappserver7`
  - Default domains directory—by default, `/var/opt/SUNWappserver7`

## Silent installation is not working correctly.

Consider the following:

- [Is the silent installation configuration file correct?](#)

### Is the silent installation configuration file correct?

To run a silent installation, you must have created a silent installation configuration file by running a standard installation using the `savestate` option as described in the *Sun ONE Application Server Installation Guide*.

```
./setup -savestate
```

In tailoring the file for your silent installation, if you have introduced any errors in the configuration file, for example mistyping a variable name, the silent installation may not run.

#### *Solution*

Verify that the silent installation configuration file is correct and that you have not introduced any errors that may invalidate the file.

## Uninstallation failure needs cleanup.

If an uninstallation fails, you may need to clean up some leftover files or processes before attempting a new installation.

#### *Solution*

Follow the instructions in [“Solution: Clean up leftover files and processes” on page 30](#).

## Can't install the load balancer plug-in.

Consider the following possibilities:

- [Is your web server installed?](#)
- [Is there a previously installed load balancer or reverse proxy plug-in on your system?](#)
- [Has the load balancer plug-in already been installed?](#)

- [Are the configuration files correct?](#)

## Is your web server installed?

Before you can install the load balancer plug-in, you must have the web server already installed (Sun ONE Web Server 6.0, SP6 or Apache Web Server 1.3.27). The web server is not required for the other Enterprise Edition components, just for the load balancer plug-in.

### *Solution*

Install the web server before installing the load balancer plug-in.

## Is there a previously installed load balancer or reverse proxy plug-in on your system?

The Sun ONE Application Server 7, Enterprise Edition requires that any existing load balancer or reverse proxy plug-in that exist on your system be removed before installing the load balancer plug-in.

### *Solution*

Remove the existing plug-in using the uninstallation program. On a clean system, the following message should display if you try to access the plug-in:

```
ERROR: information for "SUNWaspx" was not found.
```

## Has the load balancer plug-in already been installed?

If the load balancer plug-in component is disabled or grayed out on the Component Selection page, the correct version is already installed.

## Are the configuration files correct?

The installation program checks to see if the appropriate configuration files for the load balancer plug-in are found in the location you specify.

For the Sun ONE Web Server plug-in, the following files are searched:

```
<install_dir>/config/magnus.conf
<install_dir>/config/obj.conf
```

For the Apache Web Server plug-in, this file is searched:

```
<install_dir>/conf/httpd.conf
```

### *Solution*

Specify the correction location.

## Load balancer won't start.

A message similar to the following might appear in the load balancer log file when you try to start the load balancer:

```
lb.configurator: CNFG1008 : Multiple instances with the same name
: are not allowed for the cluster cluster1
```

The most likely problem is that the load balancer configuration file, `loadbalancer.xml`, is not configured correctly.

### *Solution*

Verify your `loadbalancer.xml` file and make sure that the instance name is unique.

## Shared memory creation failed.

This error occurs while running `hadbm create` or `clsetup` (which calls `hadbm create`). When the HADB server processes are booted for the first time on each machine in the HADB configuration, they create the shared memory segments which constitute the database.

The typical message in this case is:

```
Failed to create shared memory
```

This message indicates that the `hadbm create` command could not allocate the shared memory to the database segments.

If you see this error in the history file, consider the following:

- [Have you configured shared memory?](#)
- [Is there an error in your `/etc/system` file?](#)

- [Did you reboot the machine after configuring shared memory?](#)

## Have you configured shared memory?

Shared memory must be configured for the HADB host machines before you can work with the HADB.

### *Solution*

Configure shared memory by following the instructions in the Configuring Shared Memory and Semaphores section in the Preparing for HADB Setup chapter of the *Sun ONE Application Server Installation Guide*.

## Is there an error in your `/etc/system` file?

You may have made a mistake or a typing error when you configured shared memory for the HADB.

### *Solution*

Verify that you have followed the instructions in the Configuring Shared Memory and Semaphores section in the Preparing for HADB Setup chapter of the *Sun ONE Application Server Installation Guide*. Correct any typing error.

## Did you reboot the machine after configuring shared memory?

The shared memory changes in the `/etc/system` file will not take affect until you have rebooted the machine.

### *Solution*

Reboot the machine.

# clsetup is not working.

The `clsetup` command is used to automate the process of setting up a cluster in a typical single-machine configuration. After the Sun ONE Application Server 7, Enterprise Edition software and high-availability components are installed, this script uses three input files to set up a basic cluster. The most likely problems are errors in the input files (if they have been edited) and `clsetup` requirements not being met.

Consider the following possibilities:

- Have you configured shared memory?
- Has remote communication been set up correctly?
- Under SSH, are the HADB and the Application Server co-located on the same machine?
- Are the application server and HADB installed in the same directories on each machine?
- Are all the Admin Servers on the application server instances in the cluster running?
- Are the input files on all instances in the cluster identical?

## Have you configured shared memory?

Shared memory must be set up before you can use the `clsetup` command. Instructions for setting up shared memory are contained in the *Sun ONE Application Server Installation Guide*.

## Has remote communication been set up correctly?

RSH or SSH must be set up before the `clsetup` command can be run.

To verify that remote communication has been established, `rsh` to each host in the cluster. The identity should be returned from the remote host. For example:

```
rsh computer99.zmtn.company.com uname -a
```

Instructions for setting up host communications are contained in the Preparing for HADB Setup chapter of the Sun ONE Application Server *Installation Guide*.

clsetup is not working.

---

**TIP** After the SSH environment is set up, the very first time the SSH is invoked, you need to add the node machine name to the `known_hosts` file. Type the following: `ssh machine_name`

SSH will prompt you to add the *machine\_name* to the `known_hosts` file by asking a yes/no question. Answer yes.

---

### *Solution*

If the verification does not work, remote communication for the cluster has not been set up correctly. Instructions for doing this are contained in the Setting Up Host Communication section of the *Sun ONE Application Server Installation Guide*.

## Under SSH, are the HADB and the Application Server co-located on the same machine?

If you are co-locating the HADB and the Application Server on the same machine using SSH, a `known_hosts` file must exist under the `.ssh` directory. That file is necessary for the nodes to communicate with each other, so the `hadb` cluster functions properly

### *Solution*

If the `known_hosts` file is not under the `.ssh` directory, run either the `ssh localhost` command or the `ssh hostname` command before using the `clsetup` command.

## Are the application server and HADB installed in the same directories on each machine?

The `clsetup` program can not work when the files are installed in different directories on different machines.

### *Solution:*

Reinstall the Sun ONE Application Server and HADB in the same directories on each machine."

## Are all the Admin Servers on the application server instances in the cluster running?

Before running the `clsetup` command, all the Admin Servers in the cluster must be running.

## Are the input files on all instances in the cluster identical?

The `clsetup` command is not designed to set up each instance with different values. For example, this command cannot create a JDBC connection with different settings for each instance.

### *Solution*

Verify that the input files are identical on all instances in the cluster.

## Unable to create HADB database using clsetup.

This could happen when you run `clsetup` to configure the cluster. You might see errors similar to the following in the `/var/tmp/clsetup.log` file:

```
CREATING HADB DATABASE...

/opt/SUNWhadb/4.2.2-17/bin/hadbm create
--installpath=/opt/SUNWhadb/4.2.2-17
--configpath=/etc/opt/SUNWhadb/dbdef --historypath=/var/tmp
--devicepath=/opt/SUNWhadb/4 --datadevices=1 --portbase=15200
--spares=0 --inetd=false --inetdsetupdir=/tmp --devicesize=512
--dbpassword=password --hosts=eas-v880-1,eas-v880-1 hadb

hadbm:Error 22024: Specified hosts are not reachable: [ eas-v880-1 ]

HADB Database creation failed.
```

### *Solution*

Make sure your communication protocol (RSH/SSH) is configured properly before running the `clsetup` command. If you plan to use RSH for your communication, make sure you uncomment the following line in the `clresource.conf` file before running the `clsetup` command:

```
set managementProtocol=rsh
```

If you are using SSH, make sure you closely follow all the SSH configuration steps contained in the *Sun ONE Application Server Installation Guide*.

## Problems when running clsetup as non-root.

If you want to run the `clsetup` command as a user other than root, you'll need to set up administration for non-root.

### *Solution*

Follow the instructions in the Setting Up Administration for Non-Root section in the *Sun ONE Application Server Installation Guide*.

## Insufficient space.

Consider the following possibilities:

- [Is the number of semaphores too low?](#)

### Is the number of semaphores too low?

The typical message in this case is:

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

The history file then contains the following entry:

```
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 ONE Application Server Installation Guide*.

## Can't test the ssh setting as root.

In trying to test the SSH setting using the following command:

```
# ssh hostname date
```

the console prompts for the root password:

```
# root@hostname's password:
```

When running the HADB admin clients as root, the `sshd` configuration (`/etc/sshd_config`) on all machines in the cluster must have `PermitRootLogin` set to yes. Sun SSH does not permit root login by default; it is set to no.

### *Solution*

1. Change `PermitRootLogin` on all machines to yes.
2. Restart `sshd`.

## Can't get ssh to skip the login prompt.

An error similar to the following occurs, suggesting that the `sshd` server is not running on the destination machine:

```
Secure connection to vortex-dev1 refused; reverting to insecure method.
Using rsh. WARNING: Connection will not be encrypted.
Password:
```

You can set up your local environment to use the HADB commands from anywhere by setting the `PATH` variable after you have implemented SSH. You should not have to log in.

### *Solution*

1. Verify that the SSH server is running by issuing the following command on the server machine:

```
ps -e |grep sshd
```

2. If the SSH server is not running, start it as follow:

```
/etc/init.d/sshd start
```

3. Check the `~<ssh-user>/.ssh/authorized_keys` file on each destination machine to ensure that all the public keys from all the machines are listed in that file.

Can't get ssh to skip the login prompt.

4. For both the users home directory (~<ssh-user>) and the .ssh subdirectory, ensure that write permission is not granted for other or for group

For further information on setting up host communications for the HADB, refer to the Preparing for HADB Setup chapter of the *Sun ONE Application Server Installation Guide*.

---

|             |  |
|-------------|--|
| <b>NOTE</b> | If you are using a different version of SSH than the version described in the Sun ONE Application Server documentation, consult the documentation for that SSH version for instructions on setting up public key authentication. |
|-------------|--|

---

# Startup and Login Problems

This chapter addresses common problems that can occur when the Sun™ Open Net Environment (ONE) Application Server or Admin Server are starting, or when a user is logging in.

The following sections are contained in this chapter:

- [Server does not appear to be running.](#)
- [Can't access the Admin Server.](#)
- [Can't access a server application.](#)
- [Forgot the user name or password.](#)
- [Forgot the admin Server port number.](#)
- [Server won't start \(CGI error\).](#)
- [Restart operation fails](#)
- [Restart operation fails](#)

## Server does not appear to be running.

When you visit the start page of the Application Server, the initial screen does not appear. Consider the following:

- [Is the Application Server running?](#)
- [During installation, did the initial server startup run successfully?](#)
- [Is the server available locally?](#)
- [Was the server started at the expected port?](#)

Server does not appear to be running.

- [Is your proxy setting causing a problem?](#)
- [Has an ungraceful shutdown occurred on a previously-running server?](#)

## Is the Application Server running?

Use one of the following commands to determine if the Admin Server has been started:

- On Unix:

```
netstat -a | grep <admin-server port>
```

Which produces a result like this:

```
*.9889      *.*      0    0 24576    0 LISTEN
```

- On Solaris:

```
ptree `pgrep appserver`
```

which produces a list of all the processes and their directories. Check to see if `admin-server` appears in the output listing.

### *Solution*

If the Application Server is not running, start the initially-configured administrative domain by running the following command:

```
asadmin start-domain --domain domain1
```

As the command completes, you should observe the following results:

```
asadmin start-domain --domain domain1
Instance domain1:admin-server started
Instance domain1:server1 started
Domain domain1 Started.
```

If other problems occur, you can use the following command to stop both the Admin Server as well as the Application Server instance of the initially-configured domain, `domain1`:

```
asadmin stop-domain --domain domain1 --local
```

As the command completes, you should observe the following results:

```
asadmin stop-domain --domain domain1 --local
Instance domain1:server1 stopped
Domain domain1 Stopped.
```

Now restart the domain as explained above.

Syntax on the `asadmin` command is contained in the Application Server man pages and the *Sun ONE Application Server Administrator's Guide*.

## During installation, did the initial server startup run successfully?

If the console window is still open, it should display a message like this:

```
Domain domain1 Started
```

where `domain1` is the name of the default domain. This indicates that the default domain was started successfully.

If you have already closed the console window, you can check for messages in the Application Server log file here:

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

If startup was successful, you should see a message similar to the following at the end of the log file:

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

## Is the server available locally?

To verify that the server is running locally:

1. Log on to the machine where the server is running.
2. Access the local web page. For example

```
http://localhost:80/ (the default port)
```

Situation 1: If the start page does not appear on the local machine, it is most likely that the application server isn't running or didn't start normally.

### *Solution 1*

In addition to checking the server logs for any errors during startup, check the following:

- [“Is the Application Server running?” on page 42](#)
- [“Was the server started at the expected port?” on page 44.](#)

**Situation 2:** If the start page appears locally but *not* on remote machines, there is a networking problem from the remote clients. For example, DNS might be set incorrectly (so the request is being sent to the wrong machine), the network configuration on the remote machine could be incorrect, a network router could be down, and so on.

### *Solution 2*

This is probably not an Application Server issue. Check your network.

## Was the server started at the expected port?

The server could be running at a different port number than the one you expect, 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 here:

```
domain_config_dir/domains/domain1/server/config/server.xml
```

2. Find the `http-listener` element.
3. Inspect the value of the port attribute.

*Explanation of how the expected port number can change during installation*—The server's default port number is 80, however, you can specify a different port number during installation. During installation, if the specified port number is already taken by another application when you start the server, the port number rolls forward to the next available number. For example, if a server was already running on the default port 80, the Application Server would be running on port number 81. If two servers were running, the port number would be 82, and so on.

If `http-listener` is running at a port that is in use, you may see a message similar to the following:

```
[21/Jan/2003:01:41:15] WARNING (10364): ADM0011: Could not  
reregister HttpListener with DomainRegistry.
```

Sun ONE Application Server and HADB port assignments must not conflict with other port assignments on the same machine. Default and recommended port assignments are as follows:

- Sun ONE Message Queue: 7676
- IIOP: 3700

- HTTP server root: 80
- HTTP server non-root: 1024
- Admin server root: 4848
- HADB nodes: Each node uses five consecutive ports. If the default portbase (15200) is used, node 0 uses 15200 through 15204, node 1 uses 15220 through 15224, and so on.

### *Solution 1*

Kill any other process that is running under the same port, or change the port number of the http-listener as follows:

1. Open the Administration interface (*hostname:admin\_port*).
2. Browse to the HTTP Server.
3. Browse to the HTTP Listeners.

The default listener is `http-listener-1`.

4. Click that default listener and find the port number (default is 80).
5. Change it to any unused port.
6. Save the settings.

You should no longer receive this warning.

### *Solution 2*

Change to another port and be sure to enter the correct port number when invoking the server.

## Is your proxy setting causing a problem?

You should be able to access the server directly from your local system (localhost) as follows (for the default port 80):

```
http://localhost:80/
```

You may not be able to access your local system if your browser connects to the web through a proxy. (A proxy is a program that looks like a direct web connection, but which is actually a separate program that makes that connection for you.)

A typical error message situation is:

```
The requested item could not be loaded by the proxy.
```

```
Netscape Proxy's network connection was refused by the server:  
localhost:4848 The server may not be accepting connections or may  
be busy. Try connecting again later
```

### *Solution*

To solve this problem, do one of the following:

- Direct your browser to bypass the proxy server when accessing your local system. Check your 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:80/
```

To find your hostname and domain:

- Type *hostname* at the command prompt.

## Has an ungraceful shutdown occurred on a previously-running server?

If a crash has occurred, the server could be in an inconsistent state.

### *Solution*

Use the `asadmin stop-domain` command to stop the Application Server, then restart the server using `asadmin start-domain` command.

Refer to [“Is the Application Server running?” on page 42](#) for guidelines.

## Can't access the Admin Server.

The Admin Server provides the administration facilities for the Application Server (one Admin Server per domain). The Application Server log file, at `domains/domain1/admin-server/logs/server.log`, may be helpful in determining the reason the Admin Server is not running.

If you cannot access the Admin Server, consider the following:

- [Has the Admin Server been started?](#)
- [Are you the user who installed the Application Server?](#)

- [Is the Admin Server running at the expected port?](#)

## Has the Admin Server been started?

See [“Is the Application Server running?” on page 42.](#)

## Are you the user who installed the Application Server?

When the `start-domain` or `stop-domain` command fails with the error:

```
Could not start the domain.
You don't have permission to access
<install_dir>/domains/domain1/admin-server/config
```

The error indicates that you are not logged on as the user who installed the system.

### *Solution*

You have to start the domain's admin server using the same login name as the user who installed the app server. You can then start *other* server instances using AdminGUI, once the admin server has been started, but the admin server can only be started by the person who installed the server.

## Is the Admin Server running at the expected port?

The default port number for the Admin Server is 4848. However, the server could be running at a different port number than the one you expect, either because it was intentionally installed there, or because another server was already running on the installation port when the server was started.

### *Solution*

Refer to [“Was the server started at the expected port?” on page 44](#) for guidelines on checking the port your Admin Server is actually running on.

Be sure to enter the correct port number when invoking the Admin Server.

# Can't access a server application.

If you are unable to access a particular application, find the application's context root in the deployed application's `application.xml` file in `domains/domain1/server1/applications/j2ee-apps`.

Then consider the following:

- [Is the Application Server running?](#)
- [“Is the application enabled?” on page 48](#)
- [Was application deployment successful?](#)

## Is the Application Server running?

The server must be running before an application can be accessed.

### *Solution*

Use the `asadmin` command to determine if the application server is running:

```
asadmin list-components --user admin --password password server1
```

This command should produce a listing like the following:

```
hello1 <application>
dukesbook <application>
There are no standalone WAR modules
There are no standalone EJB modules
There are no connector modules
```

For more information, see [“Server does not appear to be running.” on page 41](#).

## Is the application enabled?

Use the following command to see if the application is enabled:

```
asadmin show-component-status --user admin --password password
dukesbook
```

where `dukesbook` is the application (component) name.

## Was application deployment successful?

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

### *Solution*

To verify that deployment was successful, do the following:

1. Check *install\_dir*/domains/domain1/server/server.log for Admin Server.  
You may see entries similar to the following:

```
[20/Jul/2003:11:41:41] INFO ( 1600): DPL5109: EJBC - START of EJBC for
[stateless-converter]

[20/Jul/2003:11:41:41] INFO ( 1600): CORE3282: stdout: Remote message: Processing
beans ....

[20/Jul/2003:11:41:42] INFO ( 1600): DPL5108: EJBC - Generated code for remote
home and EJBObject implementations for [stateless-converter]

[20/Jul/2003:11:41:42] INFO ( 1600): CORE3282: stdout: Remote message: Compiling
wrapper code ....

[20/Jul/2003:11:41:46] INFO ( 1600): CORE3282: stdout: Remote message: Compiling
RMI-IIOP code ....

[20/Jul/2003:11:41:55] INFO ( 1600): DPL5110: EJBC - END of EJBC for
[stateless-converter]

[20/Jul/2003:11:41:56] INFO ( 1600): Total Deployment Time: 17605 msec, Total EJB
Compiler Module Time: 14100 msec, Portion spent EJB Compiling: 80%

Breakdown of EJBC Module Time: Total Time for EJBC: 14100 msec, CMP Generation: 0
msec (0%), Java Compilation: 10 msec (0%), RMI Compilation: 13239 msec (93%),

[20/Jul/2003:11:41:56] INFO ( 1600): ADM1041:Sent the event to
instance:[ApplicationDeployEvent -- deploy stateless-converter]

[20/Jul/2003:11:42:03] INFO ( 1600): ADM1042:Status of event to instance:[success]
```

2. Check the file system hierarchy under your server (such as server1) and look for your new application directory under j2ee-apps. If it was a module you deployed, look under the j2ee-modules directory to see your new module directory.
3. Check the instance's server.xml file in the /config directory for the instance. Look for an entry similar to the following for your application or module:

```
<j2ee-application enabled="true"
location="D:\Sun\studio5_se\appserver7\domains\domain1\server1\applications\j2ee-
apps\stateless-converter_1" name="stateless-converter"
virtual-servers="server1"/>
```

## Forgot the user name or password.

If you do not remember the administrative user name that was supplied during installation, try these solutions in this order:

1. Enter the user name `admin`. This is the default user name specified in the server configuration dialog during installation.
2. If this doesn't work, look in the following file:

`domain_config_dir/domain1/admin-server/config/admpw`

This file contains the administrator's user name followed by the encrypted form of the administrative user's password. Seeing the user name may jog your memory.

3. If this doesn't work, delete the administrative domain and recreate it with a new password.
4. As a last resort, uninstall and reinstall the Application Server.

## Forgot the admin Server port number.

If you do not remember the port number of the Admin Server, you can inspect the Admin Server configuration file to determine the HTTP server port number:

1. Navigate to `domain_config_dir/domain1/admin-server/config/` and open the `server.xml` file in a text editor.
2. Look for the following element:

```
http-listener id="http-listener-1" address="0.0.0.0"
port="4848"...
```

In this example, port 4848 is the HTTP port number in use.

## Server won't start (CGI error).

If the Application Server won't start, you may receive the following error:

```
[05/Aug/2002:01:12:12] SEVERE (21770): cgi_init reports:
HTTP4047: could not initialize CGI subsystem

(Cgistub path /export/home/sun/appserver7/appserv/lib/Cgistub),
err fork() failure [Not enough space]
```

This indicates resource shortage. Consider the following:

- [Are limits on file descriptors adequate?](#)
- [Are kernel parameters adequate?](#)

## Are limits on file descriptors adequate?

On UNIX, you can use the `ulimit` command to determine the number of available file descriptors or to set limits on the system's available file descriptors. The `ulimit` command displays the limits for the current shell and its descendants.

For `sh` shell, the `ulimit -a` command lists all the current resource limits. The `ulimit -n` command lists the maximum file descriptors plus 1.

## Are kernel parameters adequate?

### *Solution*

On UNIX, increase the system resources by modifying the `/etc/system` file to include the following entries:

```
set rlim_fd_max=4086
set rlim_fd_cur=1024
```

Reboot the system for the new kernel parameters to take effect.

After you have set the shell resources, the Application Server should start.

## Restart operation fails

When an attempted restart fails, consider the following:

- [SSL/TLS are enabled](#)

### SSL/TLS are enabled

Restart does not work if SSL/TLS are enabled.

### *Solution*

Stop and then start the instance.

# JMS failed to start.

The JMS failed to start.

---

**TIP** If your application does not require JMS functionality, disabling it can improve performance. To disable JMS, change the following settings in the `server.xml` file:

```
<jms-service port="7676" admin-user-name="admin"
admin-password="admin" init-timeout-in-seconds="30"
enabled="false">
```

---

- [Are you attempting to start the instance as a non-root user?](#)
- [Do Solaris bundled and unbundled domains and instances have the same names?](#)
- [Do the imq logs have out of memory errors?](#)

## Are you attempting to start the instance as a non-root user?

When attempting to start the application server instance as a non-root user, the command fails and the following message is displayed:

```
Could not start the instance
```

In the log file for the instance (`server.log`), the following error message occurs:

```
JMS5035: Timed out after 30000 milliseconds while trying to
verify if the JMS service startup succeeded.
```

When started as root, the application server instance starts normally.

### *Solution*

Verify the correct user owns the JMS broker instance by running the following command:

```
ls -l /opt/imq/var/instances/
```

For example, the broker files for server1 in domain1 will be in the domain1\_server1 directory. If this directory is owned by root, the ownership of the broker files must be changed to the appropriate user. For example, the following command changes the ownership of these files to the UNIX user greg in the staff group:

```
chown -R greg:staff /opt/imq/var/instance/domain1_server1
```

Unless this change is made, it is not possible for the Application Server to access these files, so the JMS broker (and ultimately the Application Server) cannot start.

## Do Solaris bundled and unbundled domains and instances have the same names?

If your machine has the Solaris 9 bundled version of the Application Server software installed, and you then install the unbundled version of the Application Server, the Message Queue broker for these application server installations will be shared.

---

|             |   |
|-------------|---|
| <b>NOTE</b> | In general, only one or the other type of bundle should be used. It not necessary to install an unbundled Application Server if a bundled version is already available. |
|-------------|---|

---

If you do not uniquely name your new domains and instances, you may receive the following errors when starting up the second instance with the same domain or instance name:

```
SEVERE: JMS5024: JMS service startup failed
```

```
SEVERE: CORE5071: An error occurred during initialization
```

### *Solution*

Give the (unbundled) domains and instances names that are different from the instances and domains in the bundled installation.

To avoid these errors, refer to the JMS Support chapter in the *Application Server Administrator's Guide* for guidance.

## Do the imq logs have out of memory errors?

If the imq logs show out of memory errors, system tuning is necessary.

JMS failed to start.

### *Solution*

1. Upgrade system memory.
2. Decrease the app server's heap size.
3. Add more swap space.

---

|             |   |
|-------------|---|
| <b>NOTE</b> | Adding more swap space will increase the number of applications that can run, which may adversely affect system performance, as more swapping will occur. |
|-------------|---|

---

For more information on optimizing your system, consult the Performance and Tuning Guide.

# Runtime Problems

This chapter addresses problems that you may encounter while running the Sun™ Open Net Environment (ONE) Application Server 7 product.

The following sections are contained in this chapter:

- [Runtime Logs](#)
- [Load Balancer / Web Server won't start](#)
- [Instance goes unused after restarting](#)
- [Can't access a web page.](#)
- [Can't access an application.](#)
- [Server responds slowly after being idle](#)
- [My application suddenly went away.](#)
- [Requests are not succeeding.](#)
- [Server log: app persistence-type = memory](#)
- [Dynamic reconfiguration failed.](#)
- [Session Persistence Problems](#)
- [HTTP session failover is not working.](#)
- [Out of Memory and Stack Overflow errors.](#)
- [HADB Transaction Failures](#)
- [Other Transaction Problems](#)
- [Sporadic failures during high loads.](#)
- [Frequent "High Load" Warnings](#)

- [Client cannot connect to HADB.](#)
- [Connection Queue Problems](#)
- [Connection Pool Problems](#)

## Runtime Logs

Refer to the logs and information in “[Server Logs](#)” on [page 81](#) for information on using the logs to troubleshooting runtime problems.

## Load Balancer / Web Server won't start

This problem occurs when two instances have the same listener value—for example, if instance `foo` has listener value `bar:80` and instance `spam` has listener value `bar:80`.

The error messages that result look like this:

```
04/Sep/2003:13:01:08] warning ( 2938): reports: lb.runtime: RNTM2029:
DaemonMonitor :http://hostname:81 : could be because of connection saturation

[04/Sep/2003:13:01:08] failure ( 2938): ServerInstance.cpp@265: reports:
lb.runtime:RNTM3002 : Failed to add listener multiple times: <instance name>

[04/Sep/2003:13:01:08] failure ( 2938): FailoverGroup.cpp@102:
reports:lb.failovermanager: FGRP1002: Instance <instance name> could not be added to
theFailoverGroup: cluster1

[04/Sep/2003:13:01:08] failure ( 2938): LBConfigurator.cpp@209:
reports:lb.cofigurator: CNFG1007 :ServerInstance <instance name> could not be added
onFailoverGroup cluster1

[04/Sep/2003:13:01:08] failure ( 2938): lbplugin.cpp@168: reports:
lb.runtime:RNTM3004 : Failed to initialise load balancing subsystem
```

### *Solution*

Make sure that the listener values for each instance are unique.

## Instance goes unused after restarting

An instance was down, and is now back up, but the access log shows that is not getting any requests.

This situation occurs when the Load Balancer has not been configured to regularly check the health of instances. When the instance was down, the Load Balancer marked it as unhealthy, and recorded that fact in the load balancer log. But now that the instance is up and running again, the Load Balancer doesn't know that its health has been restored.

### *Solution*

Configure the Load Balancer to check the health of instance regularly by adding the health checker URL to loadbalancer.xml with a line like this:

```
<health-checker url="/pathToHealthChecker"
  interval-in-seconds="10" timeout-in-seconds="30" />
```

## Can't access a web page.

A typical error that displays on the browser is the following:

```
404 Not Found
The requested URL destination_URL was not found on this server.
```

This means that the web page you are attempting to access is not available at the location you have specified. The most likely causes are a change in the location, or an error in the URL you have specified.

### *Solution*

1. Try again.
2. If you still receive this error, verify that you have entered the location correctly.
3. If you have entered the URL correctly, verify that the location has not changed or been deleted. You will have to contact the page owner to verify this.

## Can't access an application.

There are a number of possible reasons that you cannot access an application. A typical error message in this case is

Consider the following:

Can't access an application.

- [Is the application deployed?](#)
- [Is your loadbalancer.xml file correct?](#)
- [Is the web server running?](#)
- [Has the correct port been specified for the web server?](#)

## Is the application deployed?

The most likely cause is that the application is not deployed. When an application is deployed to a cluster, an entry for it appears in the web server plug-in's `loadbalancer.xml` file. If an application has been successfully deployed to a cluster, the following snippet shows how the `loadbalancer.xml` file should look:

The `cladmin` command is used to deploy an application to all instances in your cluster. Refer to the *Sun ONE Application Server Administrator's Guide* for deployment instructions. Otherwise, refer to the *Sun ONE Application Server Developer's Guide* for non-cluster deployment guidelines.

### *Solution*

If needed, redeploy the application.

## Is your loadbalancer.xml file correct?

Check the web server log files to verify that the load balancer started. If it hasn't, there may be errors about the `loadbalancer.xml` file written in the to the error log.

Consider the following:

- [Is the web server running?](#)
- [Has the correct port been specified for the web server?](#)

## Is the web server running?

Verify that the web server has started.

## Has the correct port been specified for the web server?

Determine the correct web server port number and verify that the correct port has been specified. Refer to [“Is the Admin Server running at the expected port?” on page 47](#) for guidelines on determining the port number.

## Server responds slowly after being idle

If the server takes a while to service a request after a long period of idleness, consider the following:

[Does the log contain “Lost connection” messages?](#)

Does the log contain “Lost connection” messages?

If the server log shows error messages of the form,

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

then server has 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.

### **Important Note:**

Make sure the HADB connection time out is greater than the JDBC connection pool time out. If the JDBC connection time out is more than the HADB connection time out, the connection will be closed from the HADB side, but it will be there in appserver connection pool. So when the application tries to use the connection, the application server will have to re-create the connection, which incurs significant over head

## My application suddenly went away.

Consider the following:

- [Is the application you are using being quiesced by the load balancer?](#)

## Is the application you are using being quiesced by the load balancer?

When an application is being quiesced, you may experience loss of service when the application is disabled, until the application is re-enabled.

## Requests are not succeeding.

The following problems are addressed in this section:

- [Is the load balancer timeout correct?](#)
- [Have you enabled the instances of the cluster?](#)
- [Are the system clocks synchronized?](#)
- [Is the AppServer communicating with HADB?](#)

## Is the load balancer timeout correct?

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.

On the other hand, 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?

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!

### *Solution*

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

## Have you enabled the instances of the cluster?

Even if you start an application server instance and define it to be a part of the cluster, the instance will not receive requests from the load balancer until you enable the instance. Enabling makes the instance an active part of the cluster. The correct sequence of events for activating and deactivating an instance is:

1. Start the Application Server.
2. Create an Application Server instance.
3. Enable the instance
4. Disable the instance.
5. Stop the instance.
6. Start the instance (only if it has been stopped).
7. Enable the instance.

## Is the AppServer communicating with HADB?

HADB may be created & running, but if the persistence store has not yet been created, the Application Server won't 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 --dbsystempassword super123
```

## Server log: app persistence-type = memory

The server.log shows that the J2EE application is using memory persistence instead of High Availability, with a message like this:

```
Enabling no persistence for web module [Application.war]'s sessions:  
persistence-type = [memory]
```

This situation occurs when the application server has not been configured to use HA.

### *Solution*

Enable the availability service with a command like this:

```
asadmin set --user admin --password netscape --host localhost  
--port 4848 serverName.availability-service.availabilityEnabled=true
```

## Dynamic reconfiguration failed.

The load balancer plug-in detects changes to its configuration by examining the time stamp of the `loadbalancer.xml` file. If a change has been made to the `loadbalancer.xml` file, the load balancer automatically reconfigures itself. The load balancer ensures that the modified configuration data is compliant with the DTD before overwriting the existing configuration.

If changes to the `loadbalancer.xml` file are not in the correct format, as specified by the `sun-loadbalancer_1_0.dtd` file, the reconfiguration fails and a failure notice is printed in the web server's error log files. The load balancer continues to use the old configuration in memory.

---

|             |   |
|-------------|---|
| <b>NOTE</b> | If the load balancer encounters a hard disk read error while attempting to reconfigure itself, it uses the configuration that is currently in memory, and a warning message is logged to the web server's error log file. |
|-------------|---|

---

*Solution*

Edit the `loadbalancer.xml` file as needed until it follows the correct format as specified in the DTD file.

## Session Persistence Problems

The following problems are addressed in this section:

- [The create-session-store command failed.](#)
- [Configuring instance-level session persistence didn't work.](#)

### The create-session-store command failed.

Consider the following:

- [Are the HADB and the application server instance on different sides of a firewall?](#)

Are the HADB and the application server instance on different sides of a firewall?

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 didn't 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.

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 you determine that the data has been corrupted, there are three possible solution.

### *Solution*

To bring the session store back to a consistent state, do the following:

1. Use the `asadmin clear-session-store` command to clear the session store.
2. If clearing the session store doesn't work, reinitialize the data space on all the nodes and clear the data in the HADB using the `hadbm clear` command.
3. If clearing the HADB doesn't work, delete and then recreate the database.

## HTTP session failover is not working.

The Sun ONE Application Server 7, Enterprise Edition includes the high-availability database (HADB) for storing session data. The HADB is not a general-purpose database but instead is an `HttpSession` store.

If HTTP session failover is not working correctly, consider the following:

- [Are the system clocks synchronized?](#)
- [Do all objects bound to a distributed session implement the `java.io.Serializable` interface?](#)
- [Is session information stored in `HttpSession`?](#)
- [Is your web application distributable?](#)
- [An object is cloned instead of shared](#)
- [Has a session store been created?](#)
- [Are all the machines in the cluster homogenous?](#)
- [Has high availability been enabled?](#)

## Are the system clocks synchronized?

For HTTP session failover to work, the clocks of all the computers on which the application server instances in a cluster reside must be synchronized. (For more detail, see [“Are the system clocks synchronized?” on page 61.](#))

### *Solution*

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

## Do all objects bound to a distributed session implement the `java.io.Serializable` interface?

If an object does not implement the `java.io.Serializable` interface, it will not be persisted. No errors or warnings are produced, because the lack of persistence may well be the desired behavior. The remaining session objects are successfully persisted and will fail over.

---

**NOTE** If an object which is not serializable implements `java.io.Serializable`, an exception will be thrown in the log. In this case, the entire user session is not persisted. A failover will produce an empty session.

---

### *Solution*

Make sure that every class in the session that is supposed to persist implements `java.io.Serializable`, as in this example:

```
public class MyClass implements java.io.Serializable
{
    ..
}
```

## Is session information stored in `HttpSession`?

Sun ONE Application Server 7, Enterprise Edition does not support session failover if session data is stored in a stateful session bean.

## Is your web application distributable?

For a web application to be highly available, it should be distributable. An application is non-distributable if the `webapp` element of the `web.xml` file does not contain a distributable subelement.

For additional information, refer to the *Sun ONE Application Server Developer's Guide to Web Applications*.

### *Solution*

Verify that the `webapp` element of the `web.xml` file contains a distributable subelement.

## Is the persistence type set to ha?

The persistence type must be set to `ha` for session failover to work. When you run the `clsetup` command, the persistence type is set to `ha` by default. If you do not use the `clsetup` command to set up your initial cluster, the persistence type is specified as `memory`, the default. The `memory` type offers no session persistence upon failover, while the failover capabilities offered by the `file` persistence type are intended for use only in development systems where session failover capabilities are not strictly required.

Instructions for setting the persistence type are contained in the Session Persistence chapter of the *Sun ONE Application Server Administrator's Guide*.

### *Solution*

Verify that the persistence type is set to `ha`. If it isn't, modify either the entire instance or your particular application to use the `ha` persistence type. For details, see the Session Persistence chapter of the Administration Guide.

## An object is cloned instead of shared

When using `modified attribute` persistence scope, and a session fails over or is activated after being passivated, an object that is shared between two attributes comes back as two separate copies of the object instead of as a single shared object.

The situation occurs because you can not have one object referred to by two separate attributes when using `modified attribute`. The application server serializes and persists each attribute separately, so the shared object gets serialized twice, once for each attribute. When the objects get deserialized, they are now two separate objects.

## Has a session store been created?

HTTP session failover will not work until a session store has been created using the `asadmin create-session-store` command.

Instructions for creating a session store are contained in the Session Persistence chapter of the *Sun ONE Application Server Administrator's Guide*

## Are all the machines in the cluster homogenous?

All Application Server instances in a cluster must have the same applications deployed to them. For these applications to take part in failover, they must have a consistent session persistence configuration and point to the same session store.

Any new instance that you add to a cluster must have the same version and same patch level as all existing instances in a cluster.

## Has high availability been enabled?

HTTP session failover will not work until high availability has been enabled using the `availability-enabled` attribute.

### *Solution*

Set the `availability-enabled` attribute using the `asadmin` command.

# Out of Memory and Stack Overflow errors.

When using the Sun ONE Application Server to deploy web applications, out of memory errors and stack overflow errors are occurring., even though the volume of data is not large.

Memory: 2048M  
CPU : 1 900Mhz Ultra SparcIII

### *Solution 1*

Java VM runs out of memory because the web applications are quite demanding in terms of creating Java objects.

This can usually be solved by setting the VM `Xms/Xmx` and `Xss` parameters. If you have default values similar to these:

-Xmx256m - for initial app. server install  
 -Xss512k - default for 1.4.0 server VM

try something like:

-Xmx1024m -Xss1024k

### *Solution 2*

Java VM runs out of memory because of bad web application design.

For instance, the application creates lots of references to Java objects and they are being cross referenced through the lifecycle of this application preventing the garbage collector from doing its job. With each additional user using the application, it consumes more and more heap space until it runs out.

There is nothing you can do here except for changing the application. You are most likely dealing with an application memory leak which can be traced using a profiling tool such as OptimizeIt.

### *Solution 3*

The other thing you can try is to configure your Application Server with JDK 1.4.2 instead the default 1.4.0\_02 which has better memory management and more aggressive garbage collection.

## HADB Transaction Failures

Transaction failures are generally caused by a shortage of system resources. The causes of the failures will generally be found in the application server log.

In your efforts to isolate the problem, consider the following:

- [Is there a shortage of HADB data devices space?](#)
- [Is there a shortage of other HADB resources?](#)
- [Have history files grown too large?](#)

### Is there a shortage of HADB data devices space?

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.

To find the messages, follow the instructions in [“Examine the history files” on page 102](#). 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 an explanation.

### *Solution 1*

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

```
hadbm set TotalDataDevicePerNode=size
```

(See the Administrator's Guide)

This solution requires that there is space available on the physical disk which is used for the HADB data device.

HADB will automatically re-start each node of the database.

Once all nodes have re-started, the database will start accepting insert and update requests.

### *Solution 2*

Stop and clear the HADB, and create a new instance with more nodes and/or larger data devices and/or several data devices per node, see the Administrator's guide.

Unfortunately, using this solution will erase all persistent data.

## 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 in the period where the system is short on resources.

These problems frequently occur in “High Load” situations. For details, see [“Frequent “High Load” Warnings” on page 73](#).

## Have history files grown too large?

As history files grow, they consume more and more disk space. Using the `hadbm clearhistory` command recovers that space, with the option of saving those files to a specified directory. For more information, consult the Administrator's guide.

## Other Transaction Problems

- [Can't set TRANSACTION\\_SERIALIZABLE level for a bean.](#)

## Can't set TRANSACTION\_SERIALIZABLE level for a bean.

Since the Release Notes say that `check-modified-at-commit` flag is not implemented for this release, is there an equivalent of the Weblogic `<transaction-isolation>` element in Sun ONE Application Server?

### *Solution*

First identify which resource is being used in that bean, then add the following attribute to the `jdbc-connection-pool` in the `server.xml` file:

```
transaction-isolation-level="serializable"
```

This is an optional element and does not exist in the `server.xml` file by default. You must explicitly add it.

You can do this either using the Administration interface or using the command-line interface to modify the `server.xml` file, then running the `asadmin reconfig` command.

## Sporadic failures during high loads.

High load scenarios are recognizable by the following symptoms:

- User requests do not succeed.
- The database gives multiple timeout and 'transaction aborted' messages.
- The history files have 'HIGH LOAD' warnings

High load scenarios are commonly caused by one of the following problems:

- [No space left in Tuple Log](#)
- [Node-internal log gets full](#)
- [Not enough locks](#)

**Note:**

Frequently, these problems can be solved by making more CPU horsepower available, as described in [“Solution 2: Improve CPU Utilization” on page 72](#). (For even more information, consult the Tuning Guide.)

See also: [“Frequent “High Load” Warnings” on page 73](#).

## No space left in Tuple Log

The 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 ('log throw due to..' messages in the history files), which happens as a result of:
  - Network contention, so the log records don't reach the mirror node
  - CPU contention at the mirror node which keeps it from processing the received log records quickly enough.

### *Solution 1*

Check CPU usage, as described in [“Solution 2: Improve CPU Utilization” on page 72](#).

### *Solution 2*

If CPU utilization is not a problem, and there is sufficient memory, increase the `LogBufferSize` using the `hadbm set LogbufferSize=` command.

### *Solution 3*

Look for evidence of network contention, and resolve bottlenecks.

## Node-internal log gets full

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

### *Solution 1*

Check CPU usage, as described in [“Solution 2: Improve CPU Utilization” on page 72](#).

### *Solution 2*

If CPU utilization is not a problem, and there is sufficient memory, increase the `InternalLogbufferSize` using the `hadbm set InternalLogbufferSize=` command.

## Not enough locks

Some extra symptoms that identify this condition are:

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

### *Solution 1: Increase the number of locks*

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

### *Solution 2: Improve CPU Utilization*

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.)

If there are more 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, as described in [“Sporadic failures during high loads.” on page 70](#). Otherwise, add new machines and add new nodes on them.

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

# Frequent “High Load” Warnings

In some cases HADB will write informational or warning “HIGH LOAD” messages to the history file. If such messages are frequently encountered in the history file, the database administrator should take certain steps to improve the availability of resources. In most situations, reducing load or improving host performance will increase the availability of resources. This may be accomplished by:

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

See the Administrator’s Guide for information on how to perform these tasks.

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

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`

## Client cannot connect to HADB.

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

```
HADB-E-11626: Error in IPC operations, iostat = 28: No space left on device
```

where:

- 11626 is an HADB error code, 'Error in IPC operations', which means that some Inter Process Communication operation failed.

- 'iostat = 28' means (on solaris) that the operating system set errno to ENOSPC, which again translates to 'No space left on device'.

The most likely explanation is that a `semget ( )` call failed (see the unix man pages). If HADB started successfully, and you get this message at runtime, it means that the host computer has too few semaphore “undo structures”. See the “Preparing for HADB Setup” chapter in the Installation guide for how to configure `semmnu` in `/etc/system`.

### *Solution 1*

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

### *Solution 2*

Stop the HADB database, reconfigure and reboot the affected host, re-start the HADB database. HADB will be unavailable during the process.

## Connection Queue Problems

The following problems may occur with connection queues:

- [Full connection queue closes socket](#)

### Full connection queue closes socket

This problem typically occurs in “high load” scenarios. The `server.log` file shows the error: `SEVERE (17357): Connection queue full, closing socket.`

#### *Solution: Increase configuration values*

Increase the value for `ConnQueueSize` and `MaxKeepAliveConnections` in file `magnus.conf` under the `config` directory of your Sun ONE WebServer, for example:

- `MaxKeepAliveConnections 1024`
- `ConnQueueSize 1024`

## Connection Pool Problems

The following problems may occur in relation to connection pools:

- Single sign-on requires larger connection pool.
- Server: Unable to obtain connection from pool
- JDBC connection is not available.
- Exception occurs while calling `DataSource.getConnection()`.
- Exception occurs while executing against MSSQL.
- `IOException`: Connection in invalid state

## Single sign-on requires larger connection pool.

When single sign-on (or session persistence) requires connections and the wait time is exceeded, the following error occurs:

```
Unable to get connection - Wait-Time has expired
```

The Sun ONE Application Server uses the same connection pool for both HADB session persistence and single sign-on. Single sign-on is enabled by default. If an application requires single sign-on functionality, the connection pool setting must be doubled.

---

**TIP** If your application does not require single sign-on functionality, disabling it can improve performance. To disable single sign-on, change the following settings in the `server.xml` file:

```
<virtual-server id="server1" ... >
  ...
  <property name="sso-enabled" value="false"/>
  ...
</virtual-server>
```

---

### *Solution*

Double the size of the connection pool.

For example, if an application indicates that 16 is the optimal number of connections to a single HADB node, the number of connections should be doubled to 32 if single sign-on functionality is required. In this case, the JDBC connection pool settings look like this:

```
<jdbc-connection-pool steady-pool-size="32" max-pool-size="32"
max-wait-time-in-millis="60000" pool-resize-quantity="2"
idle-timeout-in-seconds="10000"
is-isolation-level-guaranteed="true"
is-connection-validation-required="true"
connection-validation-method="auto-commit"
fail-all-connections="false" datasource
classname="com.sun.hadb.jdbc.ds.HadbDataSource" name="CluJDBC"
transaction-isolation-level="repeatable-read">
```

You should also double the `loadbalancer.xml` file setting for `response-timeout-in-seconds` from 60 seconds to 120 seconds.

```
<property name="response-timeout-in-seconds" value="120"/>
```

This value must be equal to or greater than the following:

```
Max Response Time for any activity + (<jdbc-connection-pool
max-wait-time-in-millis="90000")
```

## Server: Unable to obtain connection from pool

The application server is having trouble connecting with HADB, as evidenced by a message like the following in the `server.log` file:

```
ConnectionUtilgetConnectionsFromPool failed using connection URL:
null Unable to obtain connection from pool
```

### *Solution*

Make sure HADB is running. Make sure that session-store, JDBC connection pool, and JNDI name (`jdbc/hastore`) are created. Configure the session persistence for High Availability with a command like the following:

```
asadmin configure-session-persistence --user admin
--password netscape --host localhost --port 4848
--type ha --frequency web-method --scope session
--store jdbc/hastore server1
```

## JDBC connection is not available.

Consider the following:

- [Is the max-pool-size setting adequate?](#)

## Is the max-pool-size setting adequate?

The `server.xml` file defines the following default values:

- `<ejb-container... : steady-pool-size="32"`
- `<jdbc-connection-pool...: max-pool-size="32"`

During server start/restart, the `ejb-container` steady pool for deployed enterprise beans will be created. Since the default `steady-pool-size` is 32, 32 enterprise beans will be created for each bean unless a different value is specified in the `sun-ejb-jar.xml` file.

The `setEntityContext` method will be called for each of the beans created. If more than one bean is grabbing JDBC connections in the `setEntityContext` method from the same JDBC connection pool, the following happens:

- During the steady pool creation, all the JDBC connections from the JDBC connection pool will be used (since default `max-pool-size = 32`)
- No connections will be left for any other beans that are created.

If the newly-created beans attempt to grab a JDBC connection from the same pool in their `setEntityContext` method, an exception is thrown with the following message:

```
No available resource . Wait-time expired.
```

### *Solution*

Increase the `max-pool-size` of the default `jdbc-connection-pool` to a higher value, such as 256.

## Exception occurs while calling `DataSource.getConnection()`.

This exception occurs when an invalid `DataSource` class property is registered within the JDBC connection pool. Misspelling is a common cause. For instance, while creating a `jdbc-connection-pool` for Oracle, one might specify the following:

```
< property name="OracleURL" value="jdbc:oracle:..."/>
```

This will result in the following exception since `OracleURL` is not a property of any Oracle datasource:

```
NoSuchMethodException: setOracleURL
```

### *Solution*

Verify that all `jdbc-connection-pool` properties in use are valid.

Verify that the `datasource` classname specified is for the required vendor `datasource` class.

## Exception occurs while executing against MSSQL.

The following exception occurs while executing a statement against the MSSQL server using a non-XA driver:

```
java.sql.SQLException: [DataDirect] [SQLServer JDBC Driver] Can't
start a cloned connection while in manual transaction mode
```

This happens when the `selectMethod` property is not set to `cursor`.

### *Solution*

Ensure the `selectMethod` property is set correctly during JDBC connection pool registration:

```
<property name="selectMethod" value="cursor"/ >
```

Using the command-line interface, issue the following command:

```
asadmin set -u < admin user> -p < admin password > -H < host >
-p < port > < instance >.jdbc-connection-pool.< pool name >
.property.selectMethod="cursor"
```

The options can also be set using the graphical Administration interface.

## IOException: Connection in invalid state

The error log shows the following message:

```
WEB2001: Web application service failed
java.io.IOException: Error from HA Store:
    Connection is in invalid state for LOB operations.
    <stack trace>
```

This error occurs when you have the HADB JDBC Connection pool `transaction-isolation-level` entry set to `read-committed` and `read-uncommitted`.

*Solution*

Change the `transaction-isolation-level` value of your HADB JDBC Connection pool to `repeatable-read` and re-start the application server.



# Administration Problems

This chapter discusses problems that you may encounter while administering the Sun™ Open Net Environment (ONE) Application Server 7 product. Full reference material and instructions for performing administration tasks can be found in the *Sun ONE Application Server Administrator's Guide* and *Administrator's Guide to Security*.

The following sections are contained in this chapter:

- [Server Logs](#)
- [Command-Line Interface Problems](#)
- [Graphical Interface Problems](#)
- [Monitoring Problems](#)
- [Authentication/Authorization Problems](#)
- [Authentication/Authorization Problems](#)
- [HADB Administration Problems](#)
- [Cluster Administration Problems](#)
- [Common Administration and Recovery Actions](#)

## Server Logs

This section covers:

- [Application Server logs](#)
- [HADB History Files](#)
- [HADB logs](#)

## Application Server logs

The Application Server collects and stores event information in two log files which are located in the logs directory:

- `access`—Application Server instance HTTP events
- `server.log`—Application Server events

Log entries can also be directed to another log file as specified by the administrator. In addition, each virtual server within an Application Server instance has its own identity and can have its own log file.

The following components and subsystems can utilize selective logging of server messages:

- CORBA-based clients (ORB)
- Web container
- Enterprise JavaBeans (EJB) container
- Message-driven bean (MDB) container
- Java Transaction Service (JTS)
- Java Message Service (JMS)
- Virtual Servers

Extensive information on how these logs work and the information gathered in them is available in the Using Logging chapter of the *Sun ONE Application Server Administrator's Guide*. Log levels are also described in the online help of the Administration interface.

## HADB History Files

The HADB history files are named using the format `dbname.out.node_number`. There should be a file for each node on a host. If there are four nodes on a host locally, then four files are created, each with the node number corresponding to the node.

For example, for an HADB instance named `failover`, with two nodes on the same system, the history file names would be `failover.out.0` and `failover.out.1`.

The history files for the current server are located in the `/var/tmp` directory.

For additional information on the history files, refer to the Configuring the High Availability Database chapter in the Sun ONE Application Server *Administrator's Guide*

## HADB logs

The logs associated with high-availability administration include the following:

- Web server errors are written to the Application Server log file, `server.log`, default location `/admin-server/logs/server.log` (equates to the old web server `errors.log` file).
- Database creation errors are written to `server.log`.
- Cluster administration errors are written to:

`/var/tmp/cladmin.log`

Some guidelines on using logs:

- Set the value of the `require-monitor-data` property to `true` in the `loadbalancer.xml` file in order to see monitoring details in the log.
- The `UnhealthyInstances` messages that appear in the log should be particularly helpful in troubleshooting.
- Setting a large tuple log size will increase performance of the logging facility.
- The `cladmin.log` file may be useful in troubleshooting cluster administration.

Device directory location: `/var/opt/SUNWhadb`

Configuration files location: `/etc/opt/SUNWhadb/dbdef`

## Command-Line Interface Problems

This section discusses problems that you may encounter while using the command-line interface of the Application Server.

- [Can't access the command-line utility.](#)
- [Can't access the Application Server man pages.](#)

## Can't access the command-line utility.

After installing the Application Server software, you will need to configure your environment to include the `bin` directory of the Application Server if you are going to do any of the following:

- Run the `asadmin` command
- Run the command-line utility
- Access the `asant` utility for working with the sample applications.

### *Solution*

Add the `install_dir/bin` directory to your `PATH` environment variable. If you are not familiar with the process of setting environment variables, refer to the post-installation instructions in the *Sun ONE Application Server Installation Guide*.

---

**NOTE** If your Admin Server is running under SSL, the `--secure` flag must be used.

---

## Can't access the Application Server man pages.

For the Solaris unbundled version of the product, you will not be able to access the man pages until you add the `install_dir/man` to the `MANPATH` environment variable.

### *Solution*

Add `install_dir/man` to your `MANPATH` environment variable.

## Graphical Interface Problems

This section discusses problems that you may encounter while using the Administration interface of the Application Server.

---

**NOTE** If your Admin Server is running under SSL, `https://...` must be used for browser access.

---

This section addresses the following issues:

- [Can't access the Administration interface.](#)
- [Can't undo accidental "changes."](#)

## Can't access the Administration interface.

If the connection was refused when attempting to invoke the graphical Administration interface, it is likely that the Admin Server is not running.

### *Solution*

Refer to ["Can't access the Admin Server."](#) on page 46 for information on troubleshooting this problem.

## Can't undo accidental "changes."

If an instance has been flagged for Apply Changes Required, and you decide NOT to make changes (perhaps the changes were a mistake and you want to forget the whole thing), there is no obvious method to *unset* the Apply Changes Required condition. Clicking Apply Changes seems to be forced at this point.

Shutting down your browser, restarting the Application Server instance, and so on. does NOT clear the Apply Changes flag. You are still prompted to apply the changes (since the backup configuration file is different from the current and applied configuration file).

### *Solution*

Copy over the latest updated `server.xml` file from the backup before updates are applied. This should effectively turn off the Apply Changes Required flag.

# Monitoring Problems

This section covers:

- [Load Balancer Plug-in isn't being monitored](#)

## Load Balancer Plug-in isn't being monitored

Logging for the load balancer plug-in is not automatically turned on. To turn on load balancer plug-in log messages:

1. Set the web server logging level to DEBUG.
2. Set the value of the `require-monitor-data` property to true. For example:

```
<property name=?require-monitor-data" value="true" />
```

---

**TIP** When logging is enabled on the load balancer plug-in, the load balancer writes HTTP session IDs in the web server log files. Therefore, if the web server hosting the load balancer plug-in is located in the DMZ, we recommend that you do not use the DEBUG or similar log level in production environments. If you must use the DEBUG logging level, then you should turn off load balancer logging by setting the `require-monitor-data` property to false in `loadbalancer.xml` file.

---

For more information, refer to the Configuring Load Balancer chapter of the Sun ONE Application Server *Administrator's Guide*.

## Authentication/Authorization Problems

This section addressed the following problems:

- [Can't import the certificate for my server.](#)
- [The server does not recognize my certificate.](#)
- [LDAP authentication/equalization is not working.](#)

### Can't import the certificate for my server.

Consider the following:

- [Has the trust database been created?](#)
- [Was the certificate generated with the right tool?](#)

#### Has the trust database been created?

If you haven't created the trust database in Sun ONE Application Server, you need to do that.

*Solution*

In the Security page of the Administration interface, click the Manage Database tab and create the trust database by entering its password.

### Was the certificate generated with the right tool?

The app server supports NLS database only. So, `certutil` and `openssl` are compatible tools. You can't use certificates generated by `keytool` directly on Appserver.

*Solution*

Generate the certificate with `certutil` or `openssl`.

## The server does not recognize my certificate.

There are three certificates involved in client certificate authentication.

1. First is the server certificate with which you will enable security in the server instance. This must be installed in the server as a Certificate for "This server."
2. Second is the client certificate which you will install in the browser to authenticate yourself to the server when client-cert authentication is enabled.
3. Third is the server certificate chain which links the prior two certificates. This must be installed in the server instance as the certificate for "Server certificate chain." If this certificate is not installed on the server instance, the instance doesn't know which client certificate to authenticate.

*Solution*

Verify that all the certificates have been implemented correctly. Be sure that you implement the chain in #3 and that the ROOT Certificate Authority (CA) is trusted.

## LDAP authentication/equalization is not working.

In order for the Application Server to use an LDAP-based directory server for authentication and authorization, the security realm must be configured and the LDAP realm must be activated.

*Solution*

1. In the left pane of the Administration interface, expand the `server1/Security/Realms/ldap` tree.

2. In the right panel, verify that the Classname field contains the following information:

```
com.ipplanet.ias.security.auth.realm.ldap.LDAPRealm
```

This class is the interface between the Application Server and the LDAP-based directory server.

3. Click Properties to display the pane for configuring specifics for the Directory Server implementation. Enter data similar to the following:
  - Name: directory Value: ldap://localhost:389
  - base-dn Value: dc=sun,dc=com
  - jaas-context Value ldapRealm
4. In the left pane of the Administration interface, expand the server1/Security hierarchy and change the Default Realm to ldap.
5. Apply Changes and Restart your instance as prompted.

## HADB Administration Problems

In the Sun ONE Application Server 7, Enterprise Edition, the `hadbm` and its many subcommands and options is provided for administering the high-availability database (HADB). A summary of the `hadbm` commands is contained in [“Summary of High Availability Commands” on page 119](#).

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 ONE 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: host unreachable](#).
- [hadbm command fails: command not found](#)
- [hadbm command fails: JAVA\\_HOME not defined](#)
- [create fails: “path does not exist on a host”](#)
- [database doesn't start](#).
- [clear command failed](#)

- `create-session-store` failed
- node status is “starting” after issuing `stopnode`.
- Node failure occurred.
- Double node failure occurred.
- Can’t restart the HADB after an ungraceful shutdown.
- `hadbm` command doesn’t return control to user.

## hadbm command fails: host unreachable.

The command fails with the error, “Host unreachable: <hostname>”.

The host could be unreachable either because it is down, or because the communication pathway has not been established. To isolate the problem, consider the following:

- Is the host up and running?
- Is RSH or SSH set up and running?
- Are the SSH binaries in the proper location?

### Is the host up and running?

If the remote host isn’t running or can’t accept connections, attempts to access it will fail.

#### *Solution*

Try pinging the host to see if it is up and running, ready to accept communications:

```
ping <hostname>
```

### Is RSH or SSH set up and running?

The communication pathway must be established before the `hadbm` command can succeed.

#### *Solution*

The `hadbm` commands will not work if host communication has not been set up. That is, the HADB nodes must have been configured for Remote Shell (RSH) or Secured Shell (SSH). Refer to “Preparing for HADB setup” in the Sun ONE Application Server Installation Guide for guidelines on verifying RSH and SSH.

If the verification does not work, remote communication for the cluster has not been set up correctly. Instructions for doing this are contained in the Setting Up Host Communication section of the *Sun ONE Application Server Installation Guide*.

### Are the SSH binaries in the proper location?

When using SSH, the relevant binaries must be in the proper location.

#### *Solution*

If you use `ssh`, make sure that the binaries are in `/usr/bin`.

## hadbm command fails: command not found

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*

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

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

#### *Solution 2*

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 ONE Application Server Installation Guide*.

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

```
which asadmin
which hadbm
```

These commands should echo the paths to the utilities.

## hadbm command fails: JAVA\_HOME not defined

The message “Error: `JAVA_HOME` is not defined correctly” indicates that the `JAVA_HOME` environment variable has not been set properly.

If multiple Java versions are installed on the system, you must ensure that the `JAVA_HOME` environment variable points to the correct Java version (1.4.1\_03 for Enterprise Edition).

Instructions for setting the `PATH` variable are contained in the Preparing for HADB Setup chapter of the Sun ONE Application Server *Installation Guide*.

## create fails: “path does not exist on a host”

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

```
./hadbm create ...
...
hadbm:Error 22022: Specified path does not exist on a host. Please
specify a valid path: [ machineName ... ]
```

This error message indicates that the HADB server component is not installed on the machine on which you are trying to create the HA database.

### *Solution*

Install the HADB server on the machine you are creating the HADB on and run the command again.

## database doesn't 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?](#)
- [Do you need a simple solution?](#)

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

If the history files show the error message:

```
..'systemerr'..HADB-S-01760: Shared memory get segment failed..
```

*Solution 1: Use sync;sync and reboot instead of init 6*

The `hadbm create` command can fail with this error occurs after making changes to `/etc/system` and doing a system reset with the `init 6` command.

Instead of re-spinning the machine with `init 6`, do `sync;sync` as root user and then `reboot`.

*Solution 2: Increase the amount of shared memory*

There may not be as much shared memory as the HADB needs. The amount of shared memory required by HADB depends on parameters like `DataBufferPoolSize`, `LogbufferSize`, and other parameters. Look into the file `/etc/system` and set `shmsys:shminfo_shmmax` to the maximum value possible (the preferred value is `0xffffffff`).

Verify that other shared memory settings are configured correctly. After making your changes, issue the `hadbm stop` command and reboot the machine.

For more information on the mechanics of configuring shared memory, consult the chapter, “Preparing for HADB Setup” in the *Sun ONE Application Server Installation Guide*. For guidelines on choosing the best settings, consult the *Performance Tuning Guide*.

*Solution 3: Verify /etc/system settings*

Verify the settings in the system file. Even a single mistyped character will create problems.

*Solution 4: Resolve conflicts*

Use `ipcs` to see if there are any shared memory segments or semaphores occupied unnecessarily by you or the other users. Use `ipcrm` to free them and then try starting the database.

*Solution 5: Increase the number of semaphores*

If the problem persists, then the operating system may not have enough shared memory or semaphores, etc. Increase them according to the number of nodes you have in the machine. (For details, see the Deployment Guide). Note that after making these changes, you must restart the machine to make them available.

**Do the history files contain errors?**

If the problem still persists, look into the HADB history files.

**Note:**

The `hadbm` utility cleans up all the files it created when `hadbm create` fails. In that case, the messages about the cause of the error are lost. But if the client machine has a `historypath` directory (default `/var/tmp`), then the history files are preserved there when the command fails.

If the `historypath` directory does not exist, you need to examine the syslogs on the hosts for error and warning messages from HADB. Messages are prefixed with “HADB” (the default `syslogprefix` value, which can be changed using the `create` command’s `--set` option).

Some of the more likely error messages to look for are:

- `Shared memory get segment failed`  
The system has not been set up with enough shared memory. (Discussed in the previous section.)
- `Could not verify my node address`  
Another process is using the port which one of the HADB servers processes. This can be resolved by stopping the other process, or by setting the `PortBase` attribute to another value using the command `hadbm set portbase=<value>`.
- `hadbm <command>' fails with internal error:`  
"The database could not be started"

Check the following:

- RSH and SSH are set up correctly, and you can communicate with all the machines in the HADB configuration.
- 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.

Once you’ve verified that none of the above errors have occurred, try the following remedies, in order:

- Delete the database and retry.
- Delete the database, reboot, and retry.
- Delete database, reinstall the HADB software, and retry.
- Contact Technical Support, as described in [“Product Support” on page 17](#).

For more information, refer to the Error Message Reference.

## Do you need a simple solution?

As a last resort, try the following possible solutions.

### *Solution 1: Delete the database*

Issue the `hadbm delete` command, and see if that allows the `hadbm create` to proceed normally.

### *Solution 2: Reboot the machine.*

Sometimes a system reboot is the necessary last resort. Issue `hadbm delete`, `reboot`, and then rerun the `hadb create` command.

## clear command failed

When this command fails, the history files are likely to explain why. See [“Do the history files contain errors?” on page 92](#) for instructions on viewing the history files and a list of some common 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 `--dbsystempassword` you supplied to `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 can't remember the `dbsystem` password, you'll need to clear the database using `hadbm clear` and provide a new `dbsystem` system password.

## SQLException: No suitable driver

The `create-session-store` produced 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 the `asenv.conf` in application server config directory.

To solve it, change the `AS_HADB` to point to the location of your HADB installation

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

```
AS_HADB=/export/home0/hercules/0815/SUNWhadb/4.2.2-17
```

### *Solution 2*

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

## “No space left on device” appears in server.log

When the error message “No space left on device” appears at regular intervals in the `server.log`, it can indicate that the HADB has run out of shared memory. To solve the problem, see [“Solution 2: Increase the amount of shared memory” on page 92](#).

On the other, if the message does not come and go intermittently, then you need to add a device or increase the size of existing devices.

### *Solution: Determine available space for user data*

To determine the space available for user data, take 99% of the total device size, then subtract 4 times the `LogBufferSize`. The difference between the total device size and the free size is the user data size. If the data may be refragmented in the future, the user data size should not exceed 50% of the space available for user data. If refragmentation is not relevant, close to 100% may be used.

## node status is “starting” after issuing stopnode.

After issuing the `stopnode` command, `hadbm status --nodes` shows the node as starting.

This situation occurs if you have set up `inetd`, because the node is automatically restarted when you stop a node. The result is that the node resumes the running state but is in the offline role.

Therefore, if you have set up `inetd`, and you want to make changes to the host which requires a reboot, then you need to perform some additional tasks to stop the node:

### *Solution*

1. Comment out the `inetd` entry for that node from the `inetd` configuration files (or the node is automatically restarted as soon as you stop it).
2. Re-add the entry to the `inetd` files after you have restarted the node.
3. Restart `inetd` by sending a `SIGHUP` to the process. For example:

```
ps -e |grep inetd (to find PID)
kill -HUP <PID_inetd>
```

For additional information, refer to the HADB Setup chapter of the *Sun ONE Application Server Administrator's Guide*.

## Node failure occurred.

If `hadbm status` shows that a node has stopped, you should first follow the instructions in [“Examine the history files” on page 102](#) to see what caused the node failure. A common situation is a power failure on the machine where the node resides. (If a power failure hasn't occurred, examine the logs to find the cause.)

---

**TIP** You may want to restart a node if you notice strange behavior in a node (for example excessive CPU consumption) and want to check whether a restart cures the problem. Use the `hadbm restartnode` command to restart an HADB node.

---

### *Solution*

To get the node back up and running, do the following (example for node 2):

- For a spare node:

```
host1$hadbm startnode 2
```

- For an active node:

```
host1$hadbm startnode -l=repair 2
```

## Double node failure occurred.

A properly configured HADB installation will handle single node failures. Double node failures (two mirror nodes are down) are not handled—all persistent data are lost. In the case of power failure, data will be consistent and available, as long as there isn't a double node failure.

When a double node failure occurs, the `hadbm status` command shows the database as “non-operational”.

### *Solution*

Reset the database by following the instructions in [“Clear the database and recreate session store” on page 103](#).

## Can't restart the HADB after an ungraceful shutdown.

**Situation:** On a machine with a running HADB, the machine was shut down without first stopping the database. After restarting the machine, the database does not start. 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.

In the case of a power failure, it is most likely that data is in an inconsistent state, that is, a record may have been in the process of being committed when the power failed.

---

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

When you restart the HADB in that manner, data remains available. On the other hand, if you stop and start HADB in separate operations using `hadbm stop` and `hadbm start`, data is unavailable while HADB is stopped.

---

### *Solution*

To rectify the situation, verify that the node states show Starting/Recovering, then reset the database by following the instructions in [“Clear the database and recreate session store” on page 103](#).

## hadbm command doesn't return control to user.

Many hadbm commands, in particular hadbm set, restart all the nodes of the database in order. If some problem has occurred, then the command may not return.

### *Solution 1*

From another window/shell, look at the history files for all the nodes to see if an error has occurred or if the command is still in progress. Run `hadbm status --nodes` to see if all the nodes are up and running. If they are not and there appears to be a permanent failure, you will need to cancel the command, and then try running `hadbm restart`.

### *Solution 2*

If Solution 1 fails, and your command was an attempt to set a configuration value for hadbm, try resetting it back to its old value and see if the database restarts correctly.

If the restart continues to fail, follow the instructions in [“Clear the database and recreate session store” on page 103](#) to reset the database.

### *Solution 3*

If clearing the database is unsuccessful, you'll have to delete the database using `hadbm delete`, recreate it using `hadbm create`, and then recreate the session store using `asadmin create-session-store`.

## Cluster Administration Problems

In the Sun ONE Application Server 7, Enterprise Edition, you can use the `cladmin` command to run the following `asadmin` commands simultaneously on all application server instances in a cluster: `start-instance`, `stop-instance`, `deploy`, `undeploy`, `create-jdbc-resource`, `create-jdbc-connection-pool`, `configure-session-persistence`, `delete-jdbc-resource`, `delete-jdbc-connection-pool`. This simplifies the task of cluster administration.

The `cladmin` command is located in the `install_dir/bin` directory. The default location of the `cladmin` input files, `clinstance.conf` and `clpassword.conf`, is `/etc/opt/SUNWappserver7`.

Refer to the chapter on Using the `cladmin` Command in the Sun ONE Application Server *Administrator's Guide* for a full explanation of this command.

This section addresses the following problems:

- [Refragmentation of the HADB fails.](#)
- [The cladmin command is not working.](#)
- [Application is not available on the cluster.](#)

## Refragmentation of the HADB fails.

The attempt to refragment the HADB failed.

Consider the following possibility:

- [Is there enough space on the data devices?](#)

Is there enough space on the data devices?

Messages like these indicate that refragmentation failed for lack of space on the data devices:

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

The problem occurs when data devices are filled beyond 50% or 60% of the available space, which does not leave enough extra space to carry out the refragmentation. (To see how much space has been used on the machine, use the `df` command. To calculate the amount of space that can be used for user data, see [“Solution: Determine available space for user data” on page 95.](#))

---

**TIP** Monitor your data device usage using the `hadbm deviceinfo` command.

---

### *Solution 1*

If your history files are becoming too large, clear them using the `hadbm clearhistory` command. This is the simplest solution, so try it first. History files are located in `/var/tmp`.

### *Solution 2*

Find out what disk the data devices are on with the `hadbm get DevicePath` command and check the for space on that disk. If there is room, increase the size of the data devices using the following command:

```
hadbm set TotalDataDevicePerNode=size
```

### *Solution 3*

If your data devices are using more than 50-60% of capacity and you cannot increase the size of your device as suggested above, do one of the following:

- If your machine has extra disks or has the possibility of adding additional disks, use `hadbm stop` and `hadbm delete` to stop and delete the database. Then create a new database with an increased number of devices.
- Add the nodes without refragmenting using

```
hadbm addnodes --no-refragment
```

Then recreate the session store so that it is applied to the new nodes by following the instructions in [“Clear the database and recreate session store” on page 103](#).

### *Solution 4*

If your devices are running at 80% or 90%, and all else fails, follow the instructions in [“Clear the database and recreate session store” on page 103](#).

## The cladmin command is not working.

Consider the following possibilities:

- [Are the Admin Servers of all the instances in the cluster started?](#)
- [Do all the instances in the cluster have same administrator user name and password?](#)
- [Are the input files correct?](#)
- [Are the input files on all instances in the cluster identical?](#)

Are the Admin Servers of all the instances in the cluster started?

Before running the `clsetup` command, all the Admin Servers in the cluster must be running.

Do all the instances in the cluster have same administrator user name and password?

During installation, the installation program creates a `clinstance.conf` file with entries for two instances. If you add more instances to the cluster, you must add information about these instances in the `clinstance.conf` file.

### Are the input files correct?

The order in which entries appear in the `clinstance.conf` file is important and must not be changed from the default order. If you add information about more application server instances, entries for these instances must in the correct order.

#### *Solution*

Verify that any changes you have made to the input files follow the format specified in the *Sun ONE Application Server Administrator's Guide*.

### Are the input files on all instances in the cluster identical?

The values in the input files must be identical on all instances in the cluster. The `cladmin` command is not designed to set up each instance with different values.

#### *Solution*

Verify that the `cladmin` input files are identical on all instances in the cluster.

## Application is not available on the cluster.

Consider the following possibilities:

- [Did the application deploy successfully to the cluster?](#)

### Did the application deploy successfully to the cluster?

It's possible that the deploy operation failed. To find out, run this command against each instance in the cluster.

```
asadmin list-components --type web
```

#### *Solution*

If the application isn't listed, try redeploying it and look for errors during deployment.

# Common Administration and Recovery Actions

This section describes common administrative and recovery procedures that are used in a variety of situations.

This section covers:

- [Examine the history files](#)

- [Maintain service while taking HADB offline](#)
- [Clear the database and recreate session store](#)

## Examine the history files

The history files are generally found at their default location, `/var/tmp`. If they are not at that location, use `hadbm get HistoryPath` to find the path to the history files.

The history file names are of the form `<dbname>.out.<nodenumber>`. The default database name is `hadb`, so for the default database name, the history file for node 0 would be `hadb.out.0`.

## Maintain service while taking HADB offline

Any command that makes HADB unavailable (such as `hadbm clear`) causes the application servers to start reporting errors in the error log. Client requests will then take a long time to get handled as the application continues retrying its requests to HADB, which can't answer because it is unavailable.

You can avoid this situation by disabling session persistence prior to clearing or stopping the database. This procedure takes time, but it lets the system maintain full service of your application(s) while HADB is down.

Perform the following steps:

1. Disable session persistence by using `cladmin` to set `availability-enabled` to `false` for the cluster. (See the “Session Persistence” section of the Admin Guide for the details of this procedure.)
2. Restart all your instances using the following approach:
  - Disable half of the instances in your cluster (or as many as you can at a time to maintain the necessary level of service for your application) by marking them as disabled in the load balancer configuration file. (See the load balancing section of the Admin Guide for details)
  - After the quiescence period has been reached, restart the disabled instances, and then re-enable them in the load balancer.
  - Repeat those steps for the next batch of instances until you have restarted all the instances

3. Once HADB is back up and running again, set availability-enabled to true and follow the restart process again.

## Clear the database and recreate session store

Clearing the database, restarting it, and recreating the session store is always the quickest way to fix your database. All session data in HADB will be lost, but all session data will still be available because it exists in the application server cache. (The only exception is sessions that have been passivated. They will not be in the application server cache and thus will be lost when you clear HADB.)

---

**TIP** If you need to keep servicing user requests, following the instructions in the previous section, “[Maintain service while taking HADB offline](#)”.

---



---

**TIP** **Avoid losing important data.**

For transient session data, losing the data is generally not an issue. The problem concerns losing passivated session state, that is, when a session no longer has any data in the Application Server cache and the state has been passivated to the HADB. Sessions are passivated when the maximum number of sessions exceeds the number specified in `sun-web.xml` file for each application.

To avoid losing data, configure the application to reject any new session requests when the maximum number of sessions is met by setting the maximum number of sessions to a very high value and rejecting any sessions beyond that. This prevents passivation, thus avoiding the risk of losing session data if you need to clear the HADB.

---

1. Use this command to clear the database, reinitialize all data devices, and recreate all system tables:

```
hadbm clear --spares x --dbpassword=tttt smokedb
```

where *x* is the number of spares you originally had, and *tttt* is the database password.

This command clears the database—all your old data is lost.

**2. Get the JDBC URL:**

```
asadmin hadbm get jdbcURL smokedb
```

**3. Recreate the appserver schemas and set up session persistence:**

```
asadmin create-session-store  
--storeurl <jdbc url returned from step 4>  
--storeuser appservusr --storepassword <password>
```

When issuing this command, make sure the user and password used in create-session-store match the user and password specified in the JDBC connection pool for HADB.

For more information, see the *Sun ONE Application Server Administrator's Guide*.

# Deployment Problems

This chapter addresses problems that you may encounter while deploying applications in the Sun™ Open Net Environment (ONE) Application Server 7 environment.

The following sections are contained in this chapter:

- [EJB Problems](#)
- [Ant Problems](#)

## EJB Problems

### CMP Problems

Consider the following:

- [Are you using a supported JDBC driver?](#)

Are you using a supported JDBC driver?

For a list of the supported JDBC drivers for the Application Server, refer to the *Sun ONE Application Server Platform Summary*.

## Ant Problems

Consider the following:

- [Can't run Ant remotely.](#)

## Can't run Ant remotely.

To run Ant from a remote machine, the Ant utility must be installed on that machine. You can deploy to the remote machine using the Administration interface, and you can use the `asadmin` command or the Sun ONE Studio for deployment on the remote machine.

Additionally, the default settings that are set up during Application Server installation assume the local instance to be the target for all administrator operations from `asant` (deployment, resource registration, and so on). This can be changed by editing the `install_dir/samples/common.properties` file. This file is referred to by all the samples that use `common.xml`, that is, most of the samples. A variety of settings can be customized: admin host, port, user, password, target instance, and so on.

### *Solution*

Edit the file to set the values for the remote machine or instance and run the following command on the sample.

```
asant deploy
```

# Integration Problems

This chapter discusses problems that may occur between the Sun™ Open Net Environment (ONE) Application Server 7 product and an associated product or component.

The following sections are contained in this chapter:

- [Web Server Problems](#)
- [Messaging Problems](#)
- [IDE Problems](#)

## Web Server Problems

This section addresses the following types of web server problems:

- [Sun ONE Web Server Problems](#)
- [Apache Web Server Problems](#)
- [Reverse Proxy Plug-In Problems](#)

## Sun ONE Web Server Problems

For Sun ONE Web Server, the load balancer DTD and XML files should be located in the *webserver\_install\_dir/webserver\_instance/config* directory.

## Apache Web Server Problems

For Apache Web Server, the load balancer DTD and XML files should be located in the `webserver_install_dir/webserver_instance/conf` directory.

## Reverse Proxy Plug-In Problems

Check to be sure that the file path for `libpassthrough.so/libpassthrough.dll` specified in the `magnus.conf` file is correct.

# Messaging Problems

This section addresses the following types of messaging problems:

- [JMS broker won't start.](#)
- [Can't reestablish Sun ONE Message Queue broker in a cluster environment.](#)

## JMS broker won't start.

Refer to [“Restart operation fails” on page 51](#) for information on this problem.

## Can't reestablish Sun ONE Message Queue broker in a cluster environment.

Situation: Running three Sun ONE Message Queue brokers in a cluster environment. If a broker is killed and brought up again, the following error message occurs:

```
[B1039]: Broker "<instance name>@host:port" ready.  
ERROR [B3010]: Interest 0[consumer:0] exists in the store already  
WARNING [B2030]: Internal Error: Could not add remote interest 0[consumer:0]
```

The publisher finds this broker running and publishes the message, but the subscriber cannot receive the message.

# IDE Problems

This section addresses the following problems:

- [Application Server is not in the Sun ONE Studio Server Registry.](#)

## Application Server is not in the Sun ONE Studio Server Registry.

I start the Application server 7 with the command line `asadmin start-appserv` and everything seems to work fine. I can log on to the Administration interface using `http://localhost:4848/admin` and see that the server1 instance has started and exists. I go back to the runtime tab in Sun ONE Studio and the Server Registry has no entry for the Sun ONE Application Server. Do I need to refresh the Server Registry or the Installed servers underneath?

Is the Application Server plug-in for Sun ONE Studio installed?

If not, you can get it from update center if you have Enterprise Edition of Studio. Make sure that you have registered for the product.

Is the plug-in configured with the Application Server installation directory?

1. In the IDE, bring up the Explorer window and Runtime tab at the bottom of this window.
2. Expand the Server Registry, Installed Servers, and Sun ONE Application Server7 nodes.
3. Right-click the Sun ONE Application Server 7 node and select Properties.  
The properties editor is displayed.
4. Ensure that the Sun ONE Application Server Server Home property is configured with the location of either the Application Server or Application Server's administrative client component.

Is the Application Server registered?

1. Under Installed Servers, expand the Sun ONE Application Server 7 node.  
If you see a localhost node under the Sun ONE Application Server 7 node, an Application Server instance is already registered.

2. If you either do not see a registered Application Server, or the registered server is not the server of interest, right-click the Sun ONE Application Server 7 node and choose Add Admin Server.

The Add Admin Server dialog box appears.

# Migration Problems

This chapter discusses problems you may encounter while migrating applications to the Sun™ Open Net Environment (ONE) Application Server 7 environment.

The following sections are contained in this chapter:

- [Migration Toolbox doesn't run correctly.](#)
- [Non-fatal error occurs during extraction.](#)
- [Fatal error occurs during extraction.](#)
- [Translation error occurs.](#)
- [Trouble migrating from WebLogic.](#)

Full instructions for performing migration tasks can be found in the *Sun ONE Application Server Migrating and Redeploying Server Applications* manual.

## Migration Toolbox doesn't run correctly.

If you have difficulty running the Sun ONE Application Server Migration Toolbox, consider the following possibilities:

- [Is the script customized for your environment?](#)
- [Are JAR files in your extension directory?](#)
- [Are you using the correct JDK?](#)

## Is the script customized for your environment?

It is important that the `%MIGTBX_HOME%/bin/setenv.bat` script be customized for your environment. Because of limitations of the JDK, you cannot install the Migration Toolbox in a path containing directory names with spaces. For example, do not unpack the archive in your `C:\Program Files` directory. Instead, unpack the archive either in `c:\AppServer7` or `c:\.`

## Are JAR files in your extension directory?

To avoid class version issues, all JAR files should be removed from your JDK extension directory (`%JAVA_HOME%/jre/lib/ext`) while running the Migration Toolbox. All the classes necessary for running the Migration Toolbox are included with the distribution.

---

|             |   |
|-------------|---|
| <b>NOTE</b> | Simply renaming the JAR files in the extension directory is not sufficient; you must move them to a different location. |
|-------------|---|

---

## Are you using the correct JDK?

Because many development machines have several installed copies or versions of the JDK, be sure you know which copy of the JDK you are using. Set the `JAVA_HOME` environment variable in the `%MIGTBX_HOME%/bin/setenv.bat` file to ensure you are running the preferred copy with the Migration Toolbox application.

## Non-fatal error occurs during extraction.

If only part of the automated migration succeeds (or fails), implement one of the solutions in [“Migration Toolbox doesn’t run correctly.” on page 111](#), then rerun the extraction or translation. If that doesn’t work, consider the following possibilities:

- [Are you doing a NetDynamics migration?](#)
- [Is manual migration warranted?](#)

## Are you doing a NetDynamics migration?

If a problem occurs with a NetDynamics migration, create a new project in the NetDynamics Studio and import the problematic objects. Simplify them until you can get this project to run through the appropriate tools. Introduce these files back into the original, now-migrated, project.

## Is manual migration warranted?

You may need to migrate the failed objects by hand. This is not as difficult as it may sound. The JATO framework was also designed for manual application authoring. Using the templates in the application package, follow the example of a migrated object of the same type. Documentation has been created to assist in creating new JATO objects manually.

## Fatal error occurs during extraction.

If a fatal error occurs, verify that the following items are not factors in the failure (listed in approximate order of likelihood):

1. Incorrect environment settings—Check the settings of your `%MIGTBX_HOME%/bin/setenv.bat` file and ensure they are appropriate for your machine.
2. Missing external classes
3. Incorrect tool property settings—Ensure that the Extraction Tool has valid property settings.
4. Using a non-existent runtime feature in a critical location, such as a class initializer or initialization of non-Spider threads to perform background tasks
5. Missing links directory or corrupted class files
6. Using the incorrect JDK version or platform
7. Conflicting class file versions in the boot classpath, such as those present in the JDK's extension directory

## Translation error occurs.

If you encounter an error during application translation, do the following:

- Verify that your application description file looks complete and that the XML is valid. Use a tool like XMLSpy or Internet Explorer to open the document and view it.
- Verify that the Translation Tool settings are correct.
- Verify that your environment settings in the %MIGTBX\_HOME%/bin/setenv.bat file are appropriate for your machine.
- Verify that you have a complete Migration Toolbox installation.

## Trouble migrating from WebLogic.

User is trying to port a WebLogic 6.1 web application to Sun ONE Application Server 7 using the Application Server Migration Toolbox. The job is nicely done.

But now two things occur when using a web application:

- it don't know if the webapp is connecting with the Oracle DB
- in debug mode the AppServ gave me this:

LONG message...

### *Possible Solution*

1. Copy the `struts.jar` file into `/WEB-INF/lib` and make the TLD and DTD files valuables on this directory.
2. Rebuild the EAR file with the `*.sh`.

It seems that the Migration Toolbox cannot handle JSP files that contain struts and other custom tags. These files do not get copied to the target area by the Migration Toolbox, and are listed as errors. I'm assuming that these files can just be moved manually, and the server then configured with the appropriate JAR files (`struts.jar`, and so on). I was just wondering if anyone had already got past this tool limitation, and if they'd be willing to share their experiences on migrating struts components into Sun ONE Application Server 7.

# Upgrade Problems

This chapter addresses problems that you may encounter while upgrading from one version of the Application Server to the Sun™ Open Net Environment (ONE) Application Server to another.

The following sections are contained in this chapter:

- [Upgrade Log](#)
- [Package-based installation consistently fails.](#)
- [User ID errors occur.](#)
- [Errors occur during the backup.](#)
- [Errors occur during reconfiguration.](#)
- [Upgrade interrupted; server is in inconsistent state.](#)
- [Doing a partial online upgrade of a cluster causes problems.](#)

Full instructions for upgrading can be found in the upgrading chapter of the *Sun ONE Application Server 7, Standard Edition Installation Guide*.

## Upgrade Log

Upgrade events and errors are captured in the Application Server upgrade log file (`upgrade.log`) during the active upgrade process. For commonly-encountered errors, the most likely action needed is included in this file for your convenience.

## Package-based installation consistently fails.

Examine the package installation logs to determine what is happening.

## User ID errors occur.

Errors specific to the user ID (for tarball installations) usually have to do with permissions not being set correctly.

### *Solution*

Verify that your permissions for the installation directory and the backup directory are correct.

## Errors occur during the backup.

The most likely error in this phase is related to permissions and space limitations.

### *Solution*

- Verify that your permissions for the directory specified are correct.
- Resolve any space problems, then restart the upgrade program where it left off.

## Errors occur during reconfiguration.

Problems in this phase are usually caused by a system crash or other interruption.

### *Solution*

After resolving any problems, restart the upgrade program where it left off.

## Upgrade interrupted; server is in inconsistent state.

You can restart an upgrade process for any reason by removing the `.audit_upgrade` file in the directory where the `setup` script is located, then restarting the upgrade. However, until you have completed the upgrade, the server instance you are upgrading is in an inconsistent state, so it is important that you complete any upgrade you have started.

### *Solution*

Complete the upgrade as quickly as possible to resolve the inconsistent state of the server.

## Doing a partial online upgrade of a cluster causes problems.

You must upgrade all application server instances in a cluster together. Otherwise there is a risk of version mismatch caused by a session failing over from one instance to another where the instances have different versions of components running.

---

|             |   |
|-------------|---|
| <b>NOTE</b> | You cannot perform an online upgrade of a cluster if the upgrade involves a change to the application database schema or the HADB schema. This requires a full reinstall. |
|-------------|---|

---

### *Solution*

Upgrade all instances in a cluster identically and in as short a time period as possible.

Additional information on upgrading a cluster online is contained in the *Sun ONE Application Server Administrator's Guide*.

Doing a partial online upgrade of a cluster causes problems.

# Summary of High Availability Commands

As a convenience, this appendix contains a summary of the command-line interface high-availability database (HADB) management commands for the Sun ONE Application Server 7, Enterprise Edition product.

- For detailed syntax on the `hadbm` commands, refer to the man pages.
- For instructions and guidelines on using the `hadbm` commands, refer to the *Sun ONE Application Server Administrator's Guide*. Instructions for using the `clsetup` command are contained in the *Installation Guide*.
- For a complete listing of the `asadmin` commands, refer to the appendix on Using the Command Line Interface in the *Administrator's Guide*.
- The `hadbm` command is located in the `install_dir/SUNWhadb/4/bin` directory.

The `hadbm` commands are run from the command-line using the following format:

```
hadbm <subcommand>
```

[Table A-1](#) shows the available subcommands.

**Table A-1** High-Availability Commands

|                           |  |
|---------------------------|--|
| <code>*clear</code>       | Reinitializes all the data space on all nodes, and clears all user data. |
| <code>list</code>         | Lists the existing HADBs.  |
| <code>nodestatus</code>   | Provides the key status information on all the nodes for the named HADB. |
| <code>refragment</code>   | Stores data on a newly-created node.                                     |
| <code>resourceinfo</code> | Provides information about the HADB resources.                           |
| <code>restart</code>      | Restarts the HADB  |
| <code>restartnode</code>  | Restarts the node in the specific start level.                           |

---

|           |   |
|-----------|---|
| set       | Sets the value of the specified configuration attributes to the named values. |
| start     | Starts the HADB after it has been stopped.                                    |
| startnode | Starts the node by running the startup procedure on the node.                 |
| status    | Provides information on the state of the HADB.                                |
| stop      | Stops the HADB  |
| stopnode  | Stops the node if the mirror node is up.                                      |

---

# Frequently Asked Questions (FAQs)

These frequently-asked-questions (FAQs) in this appendix are sorted into the following sections:

- [Installation and Initial Access](#)
- [Licensing](#)
- [Platform Support](#)
- [High-Availability](#)
- [Server Administration and Operation](#)
- [Application Debugging](#)
- [Upgrade/Migration](#)

## Installation and Initial Access

This section addresses the following questions:

- [How do I obtain version information?](#)
- [How do I access the sample applications?](#)

### How do I obtain version information?

In the Administration interface, click the server instance name in the left panel and select the General tab to see the Install Version information in the right panel.

At the command line, use the `asadmin version` command:

```
version [--user admin_user] [--password admin_password] [--host
localhost] [--port 4848] [--local=false] [--verbose=false]
```

You can also check the `/var/sadm/install/productregistry` XML file for the Sun ONE Application Server.

## How do I access the sample applications?

You can tour the features of the Sun ONE Application Server by running the sample applications. Access the start page here:

```
install_dir/samples/index.html
```

## Licensing

This section addresses the following questions:

- [How do I renew an expired evaluation license?](#)
- [Can I upgrade from an evaluation license to a full Sun ONE Application Server 7 license?](#)
- [In a media package, how do I get the license key?](#)
- [How can I get an evaluation license for Enterprise Edition?](#)

## How do I renew an expired evaluation license?

Update the license using the `asadmin install-license` command.

Instructions are contained in the Licensing appendix of the *Sun ONE Application Server Installation Guide*

## Can I upgrade from an evaluation license to a full Sun ONE Application Server 7 license?

Not if you selected the evaluation version that includes Sun ONE Studio. In this case, it is not possible to upgrade from Evaluation license to full license. You must uninstall the evaluation version using the uninstallation program, then install a new Standard Edition version.

Maybe, if you installed the evaluation version that does *not* include Sun ONE Studio (you must also have an unlimited license). Follow these steps:

1. Go to the Sun ONE Application Server 7 home directory.
2. Go to the `/bin` dir.
3. Run the `asadmin` command.  
This starts a new session with the prompt `asadmin>`
4. Enter the command `display-licence`.  
The licence details are displayed.
5. Check whether you have unlimited licence. If you don't, you cannot upgrade.
6. Enter the `install-licence` command.  
You will be prompted for your key.
7. Copy and paste the key from the following location:

<http://edist.central/>

## In a media package, how do I get the license key?

The software license key, also known as the activation code, is on the welcome card in the media package.

## How can I get an evaluation license for Enterprise Edition?

there isn't one

# Platform Support

Most questions about supported platforms for Sun ONE Application Server 7 can be found in the *Platform Summary* documents for the various releases here:

<http://docs.sun.com/db/prod/s1.asse>

This section addresses the following questions:

- [What's the difference between Platform Edition, Standard Edition, and Enterprise Edition?](#)
- [What Sun ONE Message Queue versions are supported?](#)
- [What J2SE versions are supported?](#)
- [What database drivers are supported?](#)
- [Is JADO supported for Sun ONE Application Server 7?](#)
- [Does Sun ONE Application Server 7 have a JCA connector to talk to SAP R3?](#)
- [Are custom authorization providers supported?](#)

## What's the difference between Platform Edition, Standard Edition, and Enterprise Edition?

The basic distinction between PE and SE is that SE allows remote administration and PE does not. Remote administration refers to the ability to manage and monitor a Sun ONE Application Server instance from another machine on the network. To administer a PE instance, you have to be running the graphical interface or the command-line interface on the machine where that PE instance is installed.

The basic difference between PE/SE and EE is that EE provides the High Availability features like load balancing (LB) and Session Persistence (SP) with existing PE/SE functionality. Please check the release notes of EE version for specific details of LB and SP.

## What Sun ONE Message Queue versions are supported?

Which versions of Sun ONE Message Queue are supported with which versions of Sun ONE Application Server?

- Sun ONE Application Server 7 PE/SE is bundled with Sun ONE Message Queue 3.0.1.
- Sun ONE Application Server 7 PE/SE, Update 1 is bundled with Sun ONE Message Queue 3.0.1 SP2.

- Sun ONE Application Server 7 EE is bundled with Sun ONE Message Queue 3.0.1

## What J2SE versions are supported?

- Sun ONE Application Server 7 PE/SE is bundled with J2SE 1.4.0\_02.
- Sun ONE Application Server 7 PE/SE, Update 1 is bundled with J2SE 1.4.1\_01.
- Sun ONE Application Server 7 EE is bundled with J2SE 1.4.0\_02

## What database drivers are supported?

The Application Server 7 product is designed to support connectivity to any database management system with a corresponding JDBC driver. Refer to the *Sun ONE Application Server Platform Summary* for a list of the components that Sun has tested and found to be acceptable for constructing J2EE-compatible database configurations.

Additional driver have been tested to meet the JDBC requirements of the J2EE 1.3 platform with the following JDBC Driver Certification Program:

<http://java.sun.com/products/jdbc/certification.html>

These drivers can be used for JDBC connectivity with the Application Server. Although Sun offers no product support for these drivers, we will support the use of these drivers with the Application Server.

## Is JADO supported for Sun ONE Application Server 7?

Answer: Java Data Objects (JADO) is used under the covers for the container-managed persistence (CMP) engine. JDO is not exposed/supported.

## Does Sun ONE Application Server 7 have a JCA connector to talk to SAP R3?

Our ISV partners for JCA connectors are Insevo, Attachmate, iWay, and Seagull. They have ported to Sun ONE Application Server 7.

## Are custom authorization providers supported?

As a J2EE 1.3 product, Sun ONE Application Server 7 does not support custom authorization providers. Custom authorization plug-in support is defined by JSR-115 (JACC) which was only introduced in J2EE 1.4. This will be supported in a future release.

## How is load balancing supported?

Definitely yes, on Enterprise Edition. The load balancer plug-in is a new component of this version and is fully documented in Enterprise Edition of the *Sun ONE Application Server Administrator's Guide*.

On PE and SE, the web connector on a web server that fronts PE and SE cannot do load balancing at all. However, with SE, a load distribution capability is enabled. In PE, the web connector (also called the reverse proxy) installed in a web server fronting a PE installation can send traffic to one and only one Sun ONE Application Server instance.

However in SE, the web connector can do load distribution. The connector can interrogate the URI in the request and determine which of several Sun ONE Application Server instances is to service the request. The mapping of the URI to an application server instance is a fixed 1:1 relationship. That is, requests sent to `http://www.abc.com/xyz/...` will always be distributed to the single Sun ONE Application Server instance that services the `xyz` URI.

For both PE and SE, you can have authentic load balancing by using a third-party product to front the web server or the actual application server instances themselves. If HttpSession is used in the application, you will want to turn on sticky load balancing based upon embedded tags or cookies. Before a session is established, the third-party load balancer could use techniques like round robin and weighted round robin load balancing to direct traffic to any of the Sun ONE Application Server instances. Thereafter, with sticky enabled, the load balancer will *stick* all subsequent processing to the application server instance that created the original session.

## Functional Support

- [Can remote EJB clients use something other than RMI over IIOP, or is that the only support protocol?](#)
- [Is sticky load balancing supported?](#)

- [Are multiple JVMs for hosting components supported?](#)

## Can remote EJB clients use something other than RMI over IIOP, or is that the only support protocol?

According to the J2EE spec, RMI/IIOP is the proposed way of remotely interacting directly with an enterprise bean inside an EJB container. However, you can wrap that bean as a web service and access it using SOAP/HTTP (SOAP over HTTP). You can also use a MOM approach using JMS destinations (queues and topics) as another way of interacting with the functionality of the bean.

## Is sticky load balancing supported?

Not for PE and SE. The Sun ONE Application Server reverse proxy plug-in (web connector) cannot channel HTTP requests based upon Session ID. However, if you use a load balancing facility from another vendor that will do the sticky load balancing, you can deploy it in front of your application servers (or web servers, if the Sun ONE Application Server reverse proxy plug-in is installed) to provide sticky load balancing.

Yes for EE. Sun ONE Application Server 7, Enterprise Edition load balancer uses a sticky round robin algorithm to load balance incoming HTTP and HTTPS requests. The load balancer plug-in uses the following methods to determine session stickiness:

- **Cookie-based method**—The load balancer plug-in uses a separate cookie to record the route information. The browser must support cookies.
- **Explicit URL rewriting method**—Sticky information is appended to the URL. Works even if the browser does not support cookies.

## Are multiple JVMs for hosting components supported?

With both PE and SE, you can create multiple application server instances for a given product installation. By default, the first server instance will be named `server1`. You can then use the Administration interface to create a second instance called, for example, `server2`. Once started, both `server1` and `server2` will have their own separate JVMs. However, a given instance (such as `server1`) has one and only one JVM. Instances do not share JVMs.

## High-Availability

This section addresses the following questions:

- [Where are the load balancer files?](#)
- [How do I enable the health checker?](#)

### Where are the load balancer files?

The load balancer configuration files, `sun-loadbalancer_1_0.dtd` and `loadbalancer.xml`, should be in the following locations:

- For Sun ONE Web Server—`webserver_install_dir/webserver_instance/config`
- For Apache Web Server—`webserver_install_dir/webserver_instance/conf`

### How do I enable the health checker?

To enable the health checker for the load balancer, edit the following properties in the `loadbalancer.xml` file:

- `url`—Specify the URL of the listener that the load balancer will check.
- `interval-in-seconds`—Specify the interval at which health checkw will occur. Default is 30 seconds.
- `timeout-in-seconds`—Specify the timeout interval for a response to be received from a healthy listener. Default is 10 seconds.

# Server Administration and Operation

This section addresses the following questions:

- What is the 'Apply Changes Required' message about?
- Why are multiple IIOP listeners useful and how many listeners are allowed?
- Does an application server instance have to be running in order to deploy to it?
- Does the Admin Server need to be running in order for me to run my application?
- Can I configure Sun ONE Application Server 7 to run as non-root?

## How do I start a single application server instance rather than all the instances in a domain?

To start *individual* instances in the default domain (domain1), use the `asadmin start-instance` command as follows:

```
start-instance server3
```

To start *all* instances in the default domain (domain 1), use the `asadmin start-appserver` command.

## What is the 'Apply Changes Required' message about?

What is happening when I receive the Apply Changes Required message in the Administration interface? I made my changes and pressed the Save button. Why do I have to also Apply Changes?

When you make your changes to a configuration setting for an instance and click Save, the update is written to a temporary file called `server.xml.changes` in a backup subdirectory under the `/config` directory. The actual `server.xml` file, stored in the main `/config` directory, represents the current configuration of the server. This file has yet to be updated. When you "Apply Changes" as prompted, the contents of the `server.xml.changes` file are applied to the `server.xml` file in the main `/config` directory. When the update is successful, the `server.xml.changes` file is deleted from the `/config/backup` directory.

This is a good feature because it allows you to make a number of configuration changes without needing to apply them to the running server one at a time. It is essentially a batch update feature.

---

|             |  |
|-------------|--|
| <b>NOTE</b> | The Administration interface tries to determine what configuration changes require a server restart. If you have changed several configuration settings that require a server instance restart, Apply Changes allows you to apply them all at once, requiring only one server restart. |
|-------------|--|

---

## Why are multiple IIOP listeners useful and how many listeners are allowed?

When configuring the ORB in the Admin Console, you can have multiple IIOP Listeners. The intention is to support an administrator's desire to have an application server instance support one port for "plain" IIOP and another port for IIOP/SSL. The Administration interface lets you configure as many IIOP Listeners as you want, but only two can be enabled (Listener Enabled toggled on) at the same time. If you try to have more than two enabled at once, the Administration interface displays an error.

## Does an application server instance have to be running in order to deploy to it?

No, the instance need not be running, but the Admin Server that controls the instance needs to be up and running. If the Admin Server is running, you can either use its Deploy feature in the Administration interface or the `asadmin deploy` command to deploy a module or an application to an instance regardless of whether or not the instance is up and running.

## Does the Admin Server need to be running in order for me to run my application?

No. You can exercise an application server instance without the Admin Server process running. However, to utilize the SNMP capabilities for monitoring the HTTP server, you need to have the Sun ONE Application Server administration process up and running due to its role in interacting with the SNMP master agent.

---

**TIP** After you add an HTTP listener to an application server instance, you must restart the SNMP monitoring subagent. If you do not, it's possible that information about the new listener will not be available through the SNMP monitoring subagent.

---

## Can I configure Sun ONE Application Server 7 to run as non-root?

Instructions for setting up administration for a non-root user are contained in the *Sun ONE Application Server Installation Guide*.

## Application Debugging

This section addresses the following questions:

- [Can I modify roles in the web or EJB deployment descriptors without restarting the application server?](#)

## Can I modify roles in the web or EJB deployment descriptors without restarting the application server?

Changes to application deployment descriptors require a redeploy (not necessarily a restart, if you are using dynamic redeployment).

However, changing application roles dynamically is not really the best approach. The application roles in J2EE are intended as design-time groupings. Instead, look into dynamically changing the mapping of particular users to these roles. You can do this by mapping the J2EE application role to a group (or groups) and altering the membership of these groups as needed.

## Upgrade/Migration

This section addresses the following questions:

- [Can I upgrade from other versions of Sun ONE Application Server 7 to Enterprise Edition?](#)

## Can I upgrade from other versions of Sun ONE Application Server 7 to Enterprise Edition?

Automatic upgrade from other versions of Sun ONE Application Server is not available for Enterprise Edition.

Only one version of Sun ONE Application Server can be on a single machine, so if you have an existing version 7 installation, you must first uninstall your existing version using the uninstallation program, then install Enterprise Edition.

# Index

## SYMBOLS

/etc/system file [51](#)  
 /var/sadm/install/productregistry [122](#)  
 /var/tmp/clsetup.log [28](#), [83](#)

## NUMERICS

404 error [57](#)

## A

ACC log service [24](#)  
 access log [82](#)  
 access problems  
   application [57](#)  
   web pages [57](#)  
 Admin Server [85](#)  
   port number, forgotten [50](#)  
 Admin Server not accessible [46](#)  
 Admin Server problems [46](#)  
 administration problems [81](#)  
 Apache Web Server [108](#)  
 Application Server won't start [50](#)  
 applications, sample [84](#)

applicaton migration [111](#)  
 asant [84](#)  
 Attachmate [125](#)  
 authentication/authorization [86](#)  
 authorization [126](#)  
 availability-enabled attribute [67](#)

## B

bin directory [84](#)

## C

Certificate Authority [87](#)  
 certificates [87](#)  
 CGI error [50](#)  
 cleanup after uninstall failure [31](#)  
 cleanup after uninstallation failure [31](#)  
 client certificate authentication [87](#)  
 clocks synchronization [65](#)  
 clsetup [35](#)  
 clsetup command [35](#)  
 cluster  
   online upgrade [117](#)

CMP [105](#)

co-location [36](#)

connection pool setting [75](#)

connection pool. [75](#)

cookies [127](#)

crash [46](#)

crash recovery [46](#)

custom authorization providers [126](#)

## D

data devices full [99](#)

database drivers

support for [125](#)

DataSource.getConnection [77](#)

debugging questions [131](#)

deployment [58](#)

deployment problems [49](#), [105](#)

diagnosis of a problem [19](#)

directories

bin [84](#)

DISPLAY environment variable [29](#)

distributable web application [66](#)

documentation [25](#)

dynamic reconfiguration [62](#)

## E

environment variables

DISPLAY [29](#)

MANPATH [84](#)

environment variables HADB settings [39](#)

error messages [22](#)

evaluation license [123](#)

## F

failover [64](#)

FAQs [121](#)

file descriptors limits [51](#)

file persistence type [66](#)

Forte for Java [16](#)

forum [26](#)

## H

HADB

history files [37](#), [82](#)

setting the environment [39](#), [90](#)

hadb create command failure [94](#)

HADB won't start [97](#)

hadbm command not working [90](#)

hadbm commands [119](#)

high-load situation [70](#)

history files for HADB [37](#), [82](#)

homogeneity [67](#)

HTTP session failover [64](#)

http-listener [44](#)

## I

inetd [96](#)

Insevo [125](#)

installation problems [27](#)

integration problems [107](#)

investigation of a problem [21](#)

isk read error [62](#)

iWay [125](#)

## J

j2ee-modules directory [49](#)

JADO [125](#)

- Java VM [68](#)
- JAVA\_HOME setting [91](#)
- JCA connectors [125](#)
- JDBC drivers [105](#)
- jdbc-connection-pool [77](#)
- JMS [108](#)
- JMS broker [52](#)
- JMS failed to start [53](#)
- JMS service startup failure [53](#)

## K

- kernel parameters, changing [51](#)
- Knowledge Base [26](#)

## L

- LDAP-based directory server [87](#)
- licensing questions [122](#)
- load balancer [33](#), [60](#), [128](#)
- load balancer plug-in [126](#)
- load balancing for PE and SE [126](#)
- loadbalancer.xml [33](#), [58](#)
- locally, server running [43](#)
- log files [23](#)
- log in problems [41](#)
- log levels [23](#)
- logs [82](#), [115](#)
  - installation [28](#)

## M

- man pages, accessing [84](#)
- MANPATH [84](#)
- MANPATH environment variable [84](#)
- max-pool-size [77](#)
- max-wait-time-in-millis [76](#)

- media package [123](#)
- memory persistence type [66](#)
- Message Queue broker issue [53](#)
- migration of applications [111](#)
- MSSQL [78](#)

## N

- networking problem [44](#)
- node failure [96](#)
- non-root [38](#), [52](#), [131](#)

## O

- online forum [26](#)
- online upgrade of cluster [117](#)
- onnection pool problems [74](#)

## P

- password, forgotten [50](#)
- PATH environment variable [84](#)
- persistence type [66](#)
- pkginfo [19](#)
- platforms [123](#)
- port [59](#)
- port number, forgotten [50](#)
- port problems [44](#), [47](#)
- ports
  - Admin Server [50](#)
- power failure and HADB [97](#)
- proxy problem [45](#)
- proxy server [46](#)

## Q

quiescence [60](#)

## R

refragmentation failed [99](#)  
reinstallation [29](#)  
remote installation [29](#)  
response-timeout-in-seconds [60](#), [76](#)  
rpm [19](#)  
RSH [35](#)

## S

sample applications [84](#)  
SAP R3 [125](#)  
Seagull [125](#)  
server.log [24](#), [82](#)  
session persistence [75](#)  
session store [67](#)  
shell resources, setting [51](#)  
showrev [19](#)  
silent installation [31](#)  
single sign-on [75](#)  
Solaris 9 [53](#)  
space on data devices [99](#)  
SSH [35](#)  
ssh [39](#)  
start page doesn't display [41](#)  
startup problems [41](#)  
sticky load balancing [126](#), [127](#)  
Sun customer support [17](#)  
Sun ONE Message Queue [124](#)  
    on Solaris [53](#)  
Sun ONE Message Queue broker [53](#)  
Sun ONE Studio [16](#), [106](#)  
SunSolve [26](#)  
synchronized clocks [65](#)

system resources, increasing [51](#)

## T

third-party products [126](#)  
thread dump, obtaining [24](#)  
transaction-isolation-level [70](#)  
trust database [86](#)

## U

ulimit [51](#)  
ungraceful shutdown [46](#), [97](#)  
uninstallation [28](#)  
uninstallation failure cleanup [31](#)  
uninstallation problems [27](#)  
upgrade [115](#), [122](#), [132](#)  
upgrade.log [115](#)  
user name, forgotten [50](#)

## V

version information [122](#)

## W

web page, can't access [57](#)  
web server [58](#)  
web.xml [66](#)  
webapp element [66](#)  
Weblogic [70](#)

## X

Xms/Xmx parameter [67](#)

Xss parameter [67](#)

