



# Sun GlassFish Enterprise Server 2.1 Developer's Guide



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

Part No: 820-4336-10  
January 2009

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

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

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

This distribution may include materials developed by third parties.

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

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

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

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

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

---

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

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

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

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

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

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

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

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

# Contents

---

<b>Preface</b> .....	19
<b>Part I Development Tasks and Tools</b> .....	25
<b>1 Setting Up a Development Environment</b> .....	27
Installing and Preparing the Server for Development .....	27
The GlassFish Project .....	28
Usage Profiles .....	29
High Availability Features .....	29
Development Tools .....	29
The <code>asadmin</code> Command .....	30
The Admin Console .....	30
The <code>asant</code> Utility .....	30
The <code>verifier</code> Tool .....	31
The NetBeans IDE .....	31
The Migration Tool .....	31
Debugging Tools .....	31
Profiling Tools .....	31
The Eclipse IDE .....	32
Sample Applications .....	32
<b>2 Class Loaders</b> .....	33
The Class Loader Hierarchy .....	33
Delegation .....	37
Using the Java Optional Package Mechanism .....	37
Using the Endorsed Standards Override Mechanism .....	37
Class Loader Universes .....	38

Application-Specific Class Loading .....	38
Circumventing Class Loader Isolation .....	39
Using the System Class Loader .....	40
Using the Common Class Loader .....	40
Sharing Libraries Across a Cluster .....	40
Packaging the Client JAR for One Application in Another Application .....	41
▼ To Package the Client JAR for One Application in Another Application .....	41
<b>3 The asant Utility .....</b>	<b>43</b>
Enterprise Server asant Tasks .....	44
The sun-appserv-deploy Task .....	44
The sun-appserv-undeploy Task .....	48
The sun-appserv-instance Task .....	51
The sun-appserv-component Task .....	54
The sun-appserv-admin Task .....	57
The sun-appserv-jspc Task .....	58
The sun-appserv-update Task .....	59
The wsgen Task .....	60
The wsimport Task .....	61
Reusable Subelements .....	63
The server Subelement .....	63
The component Subelement .....	65
The filesset Subelement .....	68
JBIs Tasks .....	68
<b>4 Debugging Applications .....</b>	<b>69</b>
Enabling Debugging .....	69
▼ To Set the Server to Automatically Start Up in Debug Mode .....	70
JPDA Options .....	70
Generating a Stack Trace for Debugging .....	71
Application Client Debugging .....	71
Sun GlassFish Message Queue Debugging .....	71
Enabling Verbose Mode .....	72
Enterprise Server Logging .....	72
Profiling Tools .....	72

---

The NetBeans Profiler .....	73
The HPROF Profiler .....	73
The JProbe Profiler .....	74
<b>Part II Developing Applications and Application Components .....</b>	<b>77</b>
<b>5 Securing Applications .....</b>	<b>79</b>
Security Goals .....	80
Enterprise Server Specific Security Features .....	80
Container Security .....	81
Declarative Security .....	81
Programmatic Security .....	82
Roles, Principals, and Principal to Role Mapping .....	82
Realm Configuration .....	84
Supported Realms .....	84
How to Configure a Realm .....	84
How to Set a Realm for an Application or Module .....	85
Creating a Custom Realm .....	85
JACC Support .....	87
Pluggable Audit Module Support .....	87
Configuring an Audit Module .....	87
The AuditModule Class .....	88
The server.policy File .....	89
Default Permissions .....	89
Changing Permissions for an Application .....	89
Enabling and Disabling the Security Manager .....	92
Configuring Message Security for Web Services .....	92
Message Security Providers .....	93
Message Security Responsibilities .....	95
Application-Specific Message Protection .....	96
Understanding and Running the Sample Application .....	99
Programmatic Login .....	102
Programmatic Login Precautions .....	102
Granting Programmatic Login Permission .....	103
The ProgrammaticLogin Class .....	103

---

User Authentication for Single Sign-on .....	104
<b>6 Developing Web Services .....</b>	<b>107</b>
Creating Portable Web Service Artifacts .....	108
Deploying a Web Service .....	108
Web Services Registry .....	109
The Web Service URI, WSDL File, and Test Page .....	110
JBI Runtime .....	111
Using the <code>jbi.xml</code> File .....	112
Using Application Server Descriptors .....	112
Using the Woodstox Parser .....	113
<b>7 Using the Java Persistence API .....</b>	<b>115</b>
Specifying the Database .....	116
Additional Database Properties .....	118
Configuring the Cache .....	118
Setting the Logging Level .....	118
Using Lazy Loading .....	119
Primary Key Generation Defaults .....	119
Automatic Schema Generation .....	120
Annotations .....	120
Supported Data Types .....	121
Generation Options .....	122
Query Hints .....	125
Changing the Persistence Provider .....	126
Restrictions and Optimizations .....	127
Extended Persistence Context Failover .....	127
Using <code>@OrderBy</code> with a Shared Session Cache .....	127
Using BLOB or CLOB Types with the Inet Oraxo JDBC Driver .....	128
Database Case Sensitivity .....	128
Sybase Finder Limitation .....	129
MySQL Database Restrictions .....	129

---

<b>8 Developing Web Applications</b> .....	133
Using Servlets .....	133
Invoking a Servlet With a URL .....	134
Servlet Output .....	134
Caching Servlet Results .....	135
About the Servlet Engine .....	138
Using JavaServer Pages .....	139
JSP Tag Libraries and Standard Portable Tags .....	140
JSP Caching .....	140
Options for Compiling JSP Files .....	144
Creating and Managing Sessions .....	144
Configuring Sessions .....	144
Session Managers .....	147
Advanced Web Application Features .....	152
Internationalization Issues .....	152
Virtual Servers .....	153
Default Web Modules .....	154
Class Loader Delegation .....	155
Using the default-web.xml File .....	155
Configuring Logging and Monitoring in the Web Container .....	156
Configuring Idempotent URL Requests .....	156
Header Management .....	157
Configuring Valves and Catalina Listeners .....	157
Alternate Document Roots .....	158
Redirecting URLs .....	159
Enabling Comet Support .....	160
Using a context.xml File .....	160
Enabling WebDav .....	160
Using mod_jk .....	162
<b>9 Using Enterprise JavaBeans Technology</b> .....	165
Summary of EJB 3.0 Changes .....	165
Value Added Features .....	166
Read-Only Beans .....	166
The pass-by-reference Element .....	167

Pooling and Caching .....	167
Bean-Level Container-Managed Transaction Timeouts .....	168
Priority Based Scheduling of Remote Bean Invocations .....	169
Immediate Flushing .....	169
EJB Timer Service .....	169
Using Session Beans .....	170
About the Session Bean Containers .....	171
Stateful Session Bean Failover .....	172
Session Bean Restrictions and Optimizations .....	177
Using Read-Only Beans .....	178
Read-Only Bean Characteristics and Life Cycle .....	178
Read-Only Bean Good Practices .....	179
Refreshing Read-Only Beans .....	179
Deploying Read-Only Beans .....	181
Using Message-Driven Beans .....	181
Message-Driven Bean Configuration .....	181
Message-Driven Bean Restrictions and Optimizations .....	183
Handling Transactions With Enterprise Beans .....	184
Flat Transactions .....	185
Global and Local Transactions .....	185
Commit Options .....	185
Administration and Monitoring .....	186
<b>10 Using Container-Managed Persistence .....</b>	<b>187</b>
Enterprise Server Support for CMP .....	187
CMP Mapping .....	188
Mapping Capabilities .....	188
The Mapping Deployment Descriptor File .....	188
Mapping Considerations .....	189
Automatic Schema Generation for CMP .....	192
Supported Data Types for CMP .....	193
Generation Options for CMP .....	195
Schema Capture .....	198
Automatic Database Schema Capture .....	198
Using the capture-schema Utility .....	198



Configuring the CMP Resource .....	199
Performance-Related Features .....	199
Version Column Consistency Checking .....	200
Relationship Prefetching .....	200
Read-Only Beans .....	201
Default Fetch Group Flags .....	202
Configuring Queries for 1.1 Finders .....	202
About JDOQL Queries .....	202
Query Filter Expression .....	203
Query Parameters .....	204
Query Variables .....	204
JDOQL Examples .....	205
CMP Restrictions and Optimizations .....	206
Disabling ORDER BY Validation .....	206
Setting the Heap Size on DB2 .....	207
Eager Loading of Field State .....	207
Restrictions on Remote Interfaces .....	207
PostgreSQL Case Insensitivity .....	207
No Support for lock-when-loaded on Sybase .....	208
Sybase Finder Limitation .....	208
Date and Time Fields .....	208
Set RECURSIVE_TRIGGERS to false on MSSQL .....	209
MySQL Database Restrictions .....	209
<b>11 Developing Java Clients .....</b>	<b>213</b>
Introducing the Application Client Container .....	213
ACC Security .....	213
ACC Naming .....	214
ACC Annotation .....	214
Java Web Start .....	214
Developing Clients Using the ACC .....	215
▼ To Access an EJB Component From an Application Client .....	215
▼ To Access a JMS Resource From an Application Client .....	217
Using Java Web Start .....	218
Running an Application Client Using the appClient Script .....	224

---

Using the package-applient Script .....	224
The client.policy File .....	224
Using RMI/IIOP Over SSL .....	224
Connecting to a Remote EJB Module Through a Firewall .....	226
<b>12 Developing Connectors .....</b>	<b>227</b>
Connector Support in the Enterprise Server .....	228
Connector Architecture for JMS and JDBC .....	228
Connector Configuration .....	228
Deploying and Configuring a Stand-Alone Connector Module .....	229
▼ To Deploy and Configure a Stand-Alone Connector Module .....	229
Redeploying a Stand-Alone Connector Module .....	230
Deploying and Configuring an Embedded Resource Adapter .....	230
Advanced Connector Configuration Options .....	231
Thread Pools .....	231
Security Maps .....	231
Overriding Configuration Properties .....	232
Testing a Connector Connection Pool .....	232
Handling Invalid Connections .....	233
Setting the Shutdown Timeout .....	233
Using Last Agent Optimization of Transactions .....	234
Inbound Communication Support .....	234
Configuring a Message Driven Bean to Use a Resource Adapter .....	235
<b>13 Developing Lifecycle Listeners .....</b>	<b>239</b>
Server Life Cycle Events .....	239
The LifecycleListener Interface .....	240
The LifecycleEvent Class .....	240
The Server Lifecycle Event Context .....	241
Deploying a Lifecycle Module .....	241
Considerations for Lifecycle Modules .....	242
<b>14 Developing Custom MBeans .....</b>	<b>243</b>
The MBean Life Cycle .....	244

MBean Class Loading .....	245
Creating, Deleting, and Listing MBeans .....	245
The <code>asadmin create-mbean</code> Command .....	245
The <code>asadmin delete-mbean</code> Command .....	246
The <code>asadmin list-mbeans</code> Command .....	246
The MBeanServer in the Enterprise Server .....	247
Enabling and Disabling MBeans .....	248
Handling MBean Attributes .....	248
<b>Part III Using Services and APIs .....</b>	<b>251</b>
<b>15 Using the JDBC API for Database Access .....</b>	<b>253</b>
General Steps for Creating a JDBC Resource .....	253
Integrating the JDBC Driver .....	254
Creating a Connection Pool .....	254
Testing a JDBC Connection Pool .....	255
Creating a JDBC Resource .....	255
Creating Applications That Use the JDBC API .....	255
Sharing Connections .....	255
Obtaining a Physical Connection From a Wrapped Connection .....	256
Marking Bad Connections .....	256
Using Non-Transactional Connections .....	257
Using JDBC Transaction Isolation Levels .....	258
Allowing Non-Component Callers .....	259
Restrictions and Optimizations .....	259
Disabling Stored Procedure Creation on Sybase .....	259
<b>16 Using the Transaction Service .....</b>	<b>261</b>
Transaction Resource Managers .....	261
Transaction Scope .....	262
Distributed Transaction Recovery .....	263
Configuring the Transaction Service .....	264
The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction .....	264
Transaction Logging .....	265

Storing Transaction Logs in a Database .....	265
Recovery Workarounds .....	266
<b>17 Using the Java Naming and Directory Interface .....</b>	<b>269</b>
Accessing the Naming Context .....	269
Global JNDI Names .....	270
Accessing EJB Components Using the CosNaming Naming Context .....	271
Accessing EJB Components in a Remote Application Server .....	271
Naming Environment for Lifecycle Modules .....	272
Configuring Resources .....	272
External JNDI Resources .....	273
Custom Resources .....	273
Using a Custom <code>jndi.properties</code> File .....	273
Mapping References .....	274
<b>18 Using the Java Message Service .....</b>	<b>275</b>
The JMS Provider .....	276
Message Queue Resource Adapter .....	277
Generic Resource Adapter .....	277
Administration of the JMS Service .....	277
Configuring the JMS Service .....	278
The Default JMS Host .....	278
Creating JMS Hosts .....	279
Checking Whether the JMS Provider Is Running .....	279
Creating Physical Destinations .....	279
Creating JMS Resources: Destinations and Connection Factories .....	280
Restarting the JMS Client After JMS Configuration .....	281
JMS Connection Features .....	281
Connection Pooling .....	281
Connection Failover .....	282
Load-Balanced Message Inflow .....	282
JMS Service High Availability .....	283
Transactions and Non-Persistent Messages .....	284
Authentication With <code>ConnectionFactory</code> .....	284
Message Queue <code>varhome</code> Directory .....	284

---

Delivering SOAP Messages Using the JMS API .....	285
▼ To Send SOAP Messages Using the JMS API .....	285
▼ To Receive SOAP Messages Using the JMS API .....	286
<b>19 Using the JavaMail API .....</b>	<b>289</b>
Introducing JavaMail .....	289
Creating a JavaMail Session .....	290
JavaMail Session Properties .....	290
Looking Up a JavaMail Session .....	290
Sending and Reading Messages Using JavaMail .....	291
▼ To Send a Message Using JavaMail .....	291
▼ To Read a Message Using JavaMail .....	292
<b>20 Using the Application Server Management Extensions .....</b>	<b>293</b>
About AMX .....	294
AMX MBeans .....	295
Configuration MBeans .....	296
Monitoring MBeans .....	296
Utility MBeans .....	296
Java EE Management MBeans .....	296
Other MBeans .....	297
MBean Notifications .....	297
Access to MBean Attributes .....	297
Dynamic Client Proxies .....	298
Connecting to the Domain Administration Server .....	298
Examining AMX Code Samples .....	299
The SampleMain Class .....	299
Connecting to the DAS .....	299
Starting an Enterprise Server .....	300
Deploying an Archive .....	301
Displaying the AMX MBean Hierarchy .....	301
Setting Monitoring States .....	301
Accessing AMX MBeans .....	301
Accessing and Displaying the Attributes of an AMX MBean .....	301
Listing AMX MBean Properties .....	301

Performing Queries .....	301
Monitoring Attribute Changes .....	302
Undeploying Modules .....	302
Stopping an Enterprise Server .....	302
Running the AMX Samples .....	302
▼ To Run the AMX Sample .....	302
<b>Index</b> .....	<b>305</b>

# Tables

---

TABLE 2-1	Sun GlassFish Enterprise Server Class Loaders .....	35
TABLE 3-1	The sun-appserv-deploy Subelements .....	45
TABLE 3-2	The sun-appserv-deploy Attributes .....	45
TABLE 3-3	The sun-appserv-undeploy Subelements .....	49
TABLE 3-4	The sun-appserv-undeploy Attributes .....	49
TABLE 3-5	The sun-appserv-instance Subelements .....	51
TABLE 3-6	The sun-appserv-instance Attributes .....	51
TABLE 3-7	The sun-appserv-component Subelements .....	55
TABLE 3-8	The sun-appserv-component Attributes .....	55
TABLE 3-9	The sun-appserv-admin Subelements .....	57
TABLE 3-10	The sun-appserv-admin Attributes .....	57
TABLE 3-11	The sun-appserv-jspc Attributes .....	58
TABLE 3-12	The sun-appserv-update Attributes .....	60
TABLE 3-13	The wsgen Attributes .....	60
TABLE 3-14	The wsimport Attributes .....	62
TABLE 3-15	The server Attributes .....	63
TABLE 3-16	The component Attributes .....	66
TABLE 7-1	Java Type to SQL Type Mappings .....	121
TABLE 7-2	Schema Generation Properties .....	123
TABLE 7-3	The asadmin deploy and asadmin deploydir Generation Options .....	124
TABLE 7-4	The asadmin undeploy Generation Options .....	125
TABLE 8-1	URL Fields for Servlets Within an Application .....	134
TABLE 8-2	The cache Attributes .....	142
TABLE 8-3	The flush Attributes .....	143
TABLE 8-4	Object Types Supported for Java EE Web Application Session State Failover .....	146
TABLE 9-1	Object Types Supported for Java EE Stateful Session Bean State Failover .....	173
TABLE 10-1	Java Type to JDBC Type Mappings for CMP .....	193
TABLE 10-2	Mappings of JDBC Types to Database Vendor Specific Types for CMP .....	194

TABLE 10-3	The sun-ejb-jar.xml Generation Elements .....	196
TABLE 10-4	The asadmin deploy and asadmin deploydir Generation Options for CMP .....	196
TABLE 10-5	The asadmin undeploy Generation Options for CMP .....	197
TABLE 15-1	Transaction Isolation Levels .....	258
TABLE 16-1	Schema for txn_log_table .....	266



# Figures

---

FIGURE 2-1	Class Loader Runtime Hierarchy .....	34
------------	--------------------------------------	----



# Preface

---

This *Developer's Guide* describes how to create and run Java™ Platform, Enterprise Edition (Java EE platform) applications that follow the open Java standards model for Java EE components and APIs in the Sun Java System Enterprise Server environment. Topics include developer tools, security, debugging, and creating lifecycle modules. This book is intended for use by software developers who create, assemble, and deploy Java EE applications using Sun GlassFish servers and software.

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

## Enterprise Server Documentation Set

The Uniform Resource Locator (URL) for Enterprise Server documentation is <http://docs.sun.com/coll/1343.8>. For an introduction to Enterprise Server, refer to the books in the order in which they are listed in the following table.

TABLE P-1 Books in the Enterprise Server Documentation Set

Book Title	Description
<i>Documentation Center</i>	Enterprise Server documentation topics organized by task and subject.
<i>Release Notes</i>	Late-breaking information about the software and the documentation. Includes a comprehensive, table-based summary of the supported hardware, operating system, Java Development Kit (JDK™), and database drivers.
<i>Quick Start Guide</i>	How to get started with the Enterprise Server product.
<i>Installation Guide</i>	Installing the software and its components.
<i>Application Deployment Guide</i>	Deployment of applications and application components to the Enterprise Server. Includes information about deployment descriptors.
<i>Developer's Guide</i>	Creating and implementing Java Platform, Enterprise Edition (Java EE platform) applications intended to run on the Enterprise Server that follow the open Java standards model for Java EE components and APIs. Includes information about developer tools, security, debugging, and creating lifecycle modules.

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

Book Title	Description
<i>Java EE 5 Tutorial</i>	Using Java EE 5 platform technologies and APIs to develop Java EE applications.
<i>Java WSIT Tutorial</i>	Developing web applications using the Web Service Interoperability Technologies (WSIT). Describes how, when, and why to use the WSIT technologies and the features and options that each technology supports.
<i>Administration Guide</i>	System administration for the Enterprise Server, including configuration, monitoring, security, resource management, and web services management.
<i>High Availability Administration Guide</i>	Setting up clusters, working with node agents, and using load balancers.
<i>Administration Reference</i>	Editing the Enterprise Server configuration file, <code>domain.xml</code> .
<i>Performance Tuning Guide</i>	Tuning the Enterprise Server to improve performance.
<i>Reference Manual</i>	Utility commands available with the Enterprise Server; written in man page style. Includes the <code>asadmin</code> command line interface.

## Related Documentation

For documentation about other stand-alone Sun GlassFish server products, go to the following:

- [Message Queue documentation \(http://docs.sun.com/coll/1343.4\)](http://docs.sun.com/coll/1343.4)
- [Identity Server documentation \(http://docs.sun.com/app/docs/prod/ident.mgmt#hic\)](http://docs.sun.com/app/docs/prod/ident.mgmt#hic)
- [Directory Server documentation \(http://docs.sun.com/coll/1224.1\)](http://docs.sun.com/coll/1224.1)
- [Web Server documentation \(http://docs.sun.com/coll/1308.3\)](http://docs.sun.com/coll/1308.3)

A Javadoc™ tool reference for packages provided with the Enterprise Server is located at <http://glassfish.dev.java.net/nonav/javaee5/api/index.html>. Additionally, the following resources might be useful:

- [The Java EE 5 Specifications \(http://java.sun.com/javaee/5/javatech.html\)](http://java.sun.com/javaee/5/javatech.html)
- [The Java EE Blueprints \(http://java.sun.com/reference/blueprints/index.html\)](http://java.sun.com/reference/blueprints/index.html)

For information on creating enterprise applications in the NetBeans™ Integrated Development Environment (IDE), see <http://www.netbeans.org/kb/55/index.html>.

For information about the Java DB database included with the Enterprise Server, see <http://developers.sun.com/javadb/>.

The GlassFish Samples project is a collection of sample applications that demonstrate a broad range of Java EE technologies. The GlassFish Samples are bundled with the Java EE Software Development Kit (SDK), and are also available from the GlassFish Samples project page at <https://glassfish-samples.dev.java.net/>.

## Default Paths and File Names

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

TABLE P-2 Default Paths and File Names

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

## Typographic Conventions

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

TABLE P-3 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>

TABLE P-3 Typographic Conventions (Continued)

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

## Symbol Conventions

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

TABLE P-4 Symbol Conventions

Symbol	Description	Example	Meaning
[ ]	Contains optional arguments and command options.	ls [-l]	The -l option is not required.
{   }	Contains a set of choices for a required command option.	-d {y n}	The -d option requires that you use either the y argument or the n argument.
`\${ }`	Indicates a variable reference.	`\${com.sun.javaRoot}`	References the value of the com.sun.javaRoot variable.
-	Joins simultaneous multiple keystrokes.	Control-A	Press the Control key while you press the A key.
+	Joins consecutive multiple keystrokes.	Ctrl+A+N	Press the Control key, release it, and then press the subsequent keys.
→	Indicates menu item selection in a graphical user interface.	File → New → Templates	From the File menu, choose New. From the New submenu, choose Templates.

## Documentation, Support, and Training

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

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

## Third-Party Web Site References

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

---

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

---

## Sun Welcomes Your Comments

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

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





PART I

Development Tasks and Tools



# Setting Up a Development Environment

---

This chapter gives guidelines for setting up an application development environment in the Sun Java™ System Enterprise Server. Setting up an environment for creating, assembling, deploying, and debugging your code involves installing the mainstream version of the Enterprise Server and making use of development tools. In addition, sample applications are available. These topics are covered in the following sections:

- “Installing and Preparing the Server for Development” on page 27
- “The GlassFish Project” on page 28
- “Usage Profiles” on page 29
- “High Availability Features” on page 29
- “Development Tools” on page 29
- “Sample Applications” on page 32

## Installing and Preparing the Server for Development

For the Sun Java Enterprise System, Enterprise Server installation is part of the system installation process. For more information, see <http://www.sun.com/software/javaenterprisesystem/index.html>.

For more information about stand-alone Enterprise Server installation, see the *Sun GlassFish Enterprise Server 2.1 Installation Guide*.

The following components are included in the full installation.

- JDK
- Enterprise Server core
  - Java 2 Platform, Standard Edition (Java SE) 6
  - Java EE 6 compliant application server
  - Admin Console
  - asadmin utility

- Other development and deployment tools
- Sun Java System Message Queue software
- The Java Business Integration runtime (JBI runtime)
- Java DB database, based on the [Derby database from Apache](http://db.apache.org/derby/manuals) (<http://db.apache.org/derby/manuals>)
- Load balancer plug-ins for web servers
- The High-Availability Database (HADB)

The NetBeans™ Integrated Development Environment (IDE) bundles the GlassFish edition of the Enterprise Server, so information about this IDE is provided as well.

After you have installed Enterprise Server, you can further optimize the server for development in these ways:

- Locate utility classes and libraries so they can be accessed by the proper class loaders. For more information, see “[Using the System Class Loader](#)” on page 40 or “[Using the Common Class Loader](#)” on page 40.
- Set up debugging. For more information, see [Chapter 4, “Debugging Applications.”](#)
- Configure the Java Virtual Machine (JVM™) software. For more information, see [Chapter 20, “Java Virtual Machine and Advanced Settings,”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

## The GlassFish Project

Enterprise Server is developed through the GlassFish™ project open-source community at <https://glassfish.dev.java.net/>. The GlassFish project provides a structured process for developing the Enterprise Server platform that makes the new features of Java EE 5 available faster, while maintaining the most important feature of Java EE: compatibility. It enables Java developers to access the Enterprise Server source code and to contribute to the development of the Enterprise Server. The GlassFish project is designed to encourage communication between Sun engineers and the community.

The Java ES edition of the Enterprise Server is based on the GlassFish source code, but provides additional value-added features such as access to a high-availability database (HADB) for session persistence and failover.

## Usage Profiles

When you install a domain, the usage profile you select determines the features that are available by default. Here is a summary of the usage profiles:

- developer profile - Provides a lightweight feature set optimized for developers, with one server instance and no clustering features.
- cluster profile - Provides the complete GlassFish feature set, including clustering features.
- enterprise profile - Provides access to features available by default only in the Java ES edition. In the GlassFish edition, this profile is only functional if you download the Java ES features separately and make global configuration changes. The Java ES features are HADB, Network Security Services (NSS), and optionally the Sun GlassFish database drivers.

For more information about usage profiles, see [“Usage Profiles” in \*Sun GlassFish Enterprise Server 2.1 Administration Guide\*](#).

## High Availability Features

High availability features such as load balancing and session failover are discussed in detail in the [Sun GlassFish Enterprise Server 2.1 High Availability Administration Guide](#). This book describes the following features in the following sections:

- For information about HTTP session persistence, see [“Distributed Sessions and Persistence” on page 145](#).
- For information about checkpointing of the stateful session bean state, see [“Stateful Session Bean Failover” on page 172](#).
- For information about failover and load balancing for Java clients, see [Chapter 11, “Developing Java Clients.”](#)
- For information about load balancing for message-driven beans, see [“Load-Balanced Message Inflow” on page 282](#).
- For information about JMS availability, see [“JMS Service High Availability” on page 283](#).

## Development Tools

The following general tools are provided with the Enterprise Server:

- [“The `asadmin` Command” on page 30](#)
- [“The Admin Console” on page 30](#)
- [“The `asant` Utility” on page 30](#)
- [“The `verifier` Tool” on page 31](#)

The following development tools are provided with the Enterprise Server or downloadable from Sun:

- “The NetBeans IDE” on page 31
- “The Migration Tool” on page 31

The following third-party tools might also be useful:

- “Debugging Tools” on page 31
- “Profiling Tools” on page 31
- “The Eclipse IDE” on page 32

## The `asadmin` Command

The `asadmin` command allows you to configure a local or remote server and perform both administrative and development tasks at the command line. For general information about `asadmin`, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

The `asadmin` command is located in the `as-install/bin` directory. Type `asadmin help` for a list of subcommands.

## The Admin Console

The Admin Console lets you configure the server and perform both administrative and development tasks using a web browser. For general information about the Admin Console, click the Help button in the Admin Console. This displays the Enterprise Server online help.

To access the Admin Console, type `http://host:4848` (developer profile) or `https://host:4848` (cluster and enterprise profiles) in your browser. The *host* is the name of the machine on which the Enterprise Server is running. By default, the *host* is `localhost`. For example:

```
http://localhost:4848
```

## The `asant` Utility

Apache Ant 1.6.5 is provided with the Enterprise Server and can be launched from the `bin` directory using the command `asant`. The Enterprise Server also provides server-specific tasks for administration and deployment; see [Chapter 3, “The `asant` Utility.”](#) The sample applications that can be used with the Enterprise Server use Ant `build.xml` files; see [“Sample Applications” on page 32.](#)

For more information about Ant, see the Apache Software Foundation web site at <http://ant.apache.org/>.

## The verifier Tool

The `verifier` tool checks a Java EE application file, including Java classes and deployment descriptors, for compliance with Java EE specifications. Java EE application files are Java archive (JAR), web archive (WAR), resource adapter archive (RAR), or enterprise archive (EAR) files. Use the `verifier` tool to check whether an application complies with the Java EE specification and to make applications portable across application servers. The `verifier` tool can be launched from the command line. For more information, see “[The verifier Utility](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## The NetBeans IDE

The NetBeans IDE allows you to create, assemble, and debug code from a single, easy-to-use interface. The GlassFish edition of the Enterprise Server is bundled with the NetBeans 5.5 IDE. To download the NetBeans IDE, see <http://www.netbeans.org>. This site also provides documentation on how to use the NetBeans IDE with the bundled Enterprise Server.

You can also use the Enterprise Server with the Sun Java Studio 8 software, which is built on the NetBeans IDE. For more information, see <http://developers.sun.com/prodtech/javatools/jsenterprise/>.

## The Migration Tool

The Migration Tool converts and reassembles Java EE applications and modules developed on other application servers. This tool also generates a report listing how many files are successfully and unsuccessfully migrated, with reasons for migration failure. For more information and to download the Migration Tool, see <http://java.sun.com/j2ee/tools/migration/index.html>.

## Debugging Tools

You can use several debugging tools with the Enterprise Server. For more information, see [Chapter 4, “Debugging Applications.”](#)

## Profiling Tools

You can use several profilers with the Enterprise Server. For more information, see “[Profiling Tools](#)” on page 72.

## The Eclipse IDE

A plug-in for the Eclipse IDE is available at <http://glassfishplugins.dev.java.net/>. This site also provides documentation on how to register the Enterprise Server and use Sun-specific deployment descriptors.

## Sample Applications

Sample applications that you can examine and deploy to the Enterprise Server are available. If you installed the Enterprise Server as part of installing the Java EE 5 SDK bundle from [Java EE 5 Downloads \(http://java.sun.com/javaee/5/downloads/\)](http://java.sun.com/javaee/5/downloads/), the samples may already be installed. You can download these samples separately from the [Code Samples \(http://java.sun.com/javaee/reference/code/index.jsp\)](http://java.sun.com/javaee/reference/code/index.jsp) page if you installed the Enterprise Server without them initially.

Most Enterprise Server samples have the following directory structure:

- The docs directory contains instructions for how to use the sample.
- The build.xml file defines asant targets for the sample. See [Chapter 3, “The asant Utility.”](#)
- The src/java directory under each component contains source code for the sample.
- The src/conf directory under each component contains the deployment descriptors.

With a few exceptions, sample applications follow the standard directory structure described here: <http://java.sun.com/blueprints/code/projectconventions.html>.

The *samples-install-dir*/bp-project/main.xml file defines properties common to all sample applications and implements targets needed to compile, assemble, deploy, and undeploy sample applications. In most sample applications, the build.xml file imports main.xml.

In addition to the Java EE 5 sample applications, samples are also available on the GlassFish web site at <https://glassfish-samples.dev.java.net/>.



# Class Loaders

---

Understanding Enterprise Server class loaders can help you determine where to place supporting JAR and resource files for your modules and applications. For general information about J2SE class loaders, see [Understanding Network Class Loaders](http://java.sun.com/developer/technicalArticles/Networking/classloaders/) (<http://java.sun.com/developer/technicalArticles/Networking/classloaders/>).

In a Java Virtual Machine (JVM), the class loaders dynamically load a specific Java class file needed for resolving a dependency. For example, when an instance of `java.util.Enumeration` needs to be created, one of the class loaders loads the relevant class into the environment. This section includes the following topics:

- “The Class Loader Hierarchy” on page 33
- “Delegation” on page 37
- “Using the Java Optional Package Mechanism” on page 37
- “Using the Endorsed Standards Override Mechanism” on page 37
- “Class Loader Universes” on page 38
- “Application-Specific Class Loading” on page 38
- “Circumventing Class Loader Isolation” on page 39

## The Class Loader Hierarchy

Class loaders in the Enterprise Server runtime follow a delegation hierarchy that is illustrated in the following figure and fully described in [Table 2-1](#).

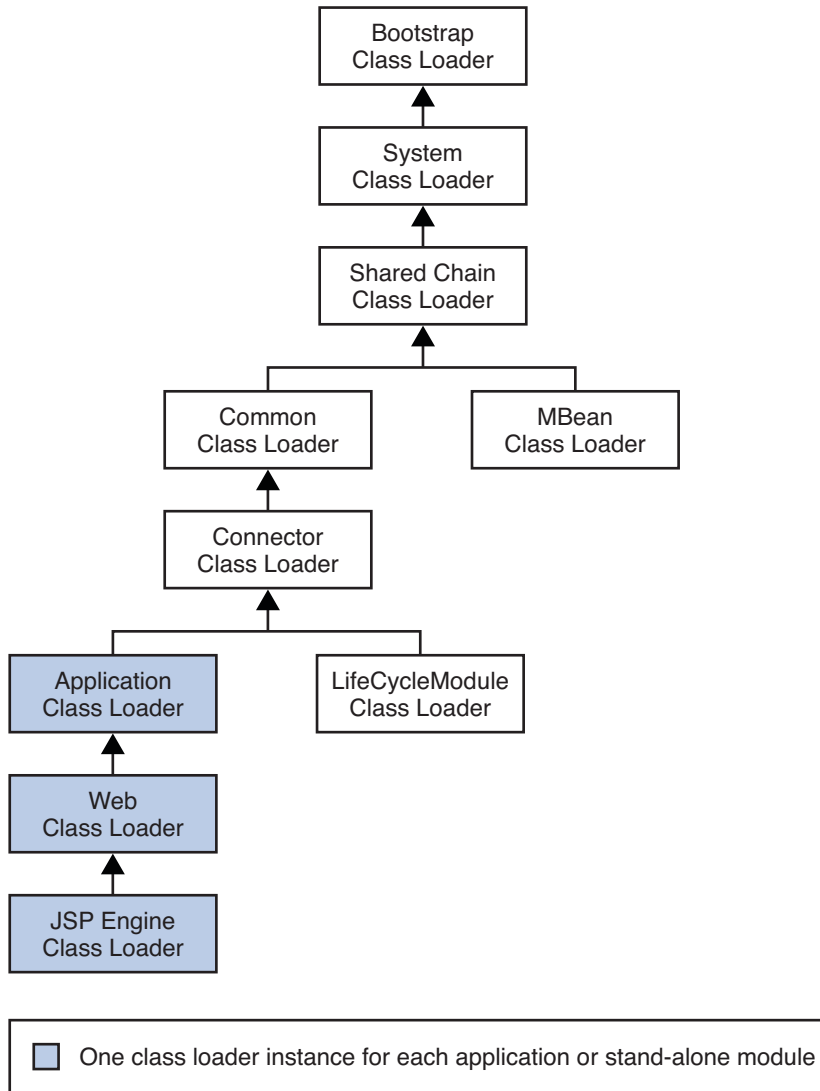


FIGURE 2-1 Class Loader Runtime Hierarchy

The following table describes the class loaders in the Enterprise Server.

TABLE 2-1 Sun GlassFish Enterprise Server Class Loaders

Class Loader	Description
Bootstrap	The Bootstrap class loader loads the basic runtime classes provided by the JVM, plus any classes from JAR files present in the system extensions directory. It is parent to the System class loader. To add JAR files to the system extensions, directory, see <a href="#">“Using the Java Optional Package Mechanism” on page 37</a> .
System	<p>The System class loader loads Enterprise Server launch classes. It is parent to the Shared Chain class loader. It is created based on the <code>system-classpath</code> attribute of the <code>java-config</code> element in the <code>domain.xml</code> file. In the developer profile, select the Enterprise Server component in the Admin Console and the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. Then select the Path Settings tab and edit the System Classpath field. See <a href="#">“Using the System Class Loader” on page 40</a> and <a href="#">“java-config” in Sun GlassFish Enterprise Server 2.1 Administration Reference</a>.</p> <p>Add the classes to the <code>system-classpath</code> attribute of the domain administration server (DAS) in addition to the <code>system-classpath</code> attribute on the server instances that use the classes. The default name for the DAS configuration is <code>server-config</code>.</p>
Shared Chain	<p>The Shared Chain class loader loads most of the core Enterprise Server classes. It is parent to the MBean class loader and the Common class loader. Classes specified by the <code>classpath-prefix</code> and <code>classpath-suffix</code> attributes of the <code>java-config</code> element in the <code>domain.xml</code> file are added to this class loader. In the developer profile, select the Enterprise Server component in the Admin Console and the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. Then select the Path Settings tab and edit the Classpath Prefix or Classpath Suffix field.</p> <p>The environment classpath is included if <code>env-classpath-ignored="false"</code> is set in the <code>java-config</code> element.</p> <p>Use <code>classpath-prefix</code> to place libraries ahead of Enterprise Server implementation classes in the shared chain. The <code>classpath-prefix</code> is ideal for placing development and diagnostic patches. To avoid overriding implementation classes, use <code>classpath-suffix</code> to place libraries after implementation classes in the shared chain.</p> <p>Add the classes to the <code>classpath-prefix</code> or <code>classpath-suffix</code> attribute of the DAS in addition to the corresponding attribute on the server instances that use the classes. The default name for the DAS configuration is <code>server-config</code>.</p>
MBean	The MBean class loader loads the MBean implementation classes. See <a href="#">“MBean Class Loading” on page 245</a> .
Common	The Common class loader loads classes in the <code>domain-dir/lib/classes</code> directory, followed by JAR files in the <code>domain-dir/lib</code> directory. It is parent to the Connector class loader. No special classpath settings are required. The existence of these directories is optional; if they do not exist, the Common class loader is not created. See <a href="#">“Using the Common Class Loader” on page 40</a> .

TABLE 2-1 Sun GlassFish Enterprise Server Class Loaders (Continued)

Class Loader	Description
Connector	The Connector class loader is a single class loader instance that loads individually deployed connector modules, which are shared across all applications. It is parent to the LifecycleModule class loader and the Application class loader.
LifeCycleModule	The LifecycleModule class loader is created once per lifecycle module. Each lifecycle-module element's classpath attribute is used to construct its own class loader. For more information on lifecycle modules, see <a href="#">Chapter 13, “Developing Lifecycle Listeners.”</a>
Application	<p>The Application class loader loads the classes in a specific enabled individually deployed module or Java EE application. One instance of this class loader is present in each class loader universe; see <a href="#">“Class Loader Universes” on page 38</a>. The Application class loader is created with a list of URLs that point to the locations of the classes it needs to load. It is parent to the Web class loader.</p> <p>The Application class loader loads classes in the following order:</p> <ol style="list-style-type: none"> <li>1. Classes specified by the library-directory element in the application.xml deployment descriptor or the --libraries option during deployment; see <a href="#">“Application-Specific Class Loading” on page 38</a></li> <li>2. Classes specified by the application's or module's location attribute in the domain.xml file, determined during deployment</li> <li>3. Classes in the classpaths of the application's sub-modules</li> <li>4. Classes in the application's or module's stubs directory</li> </ol> <p>The location attribute points to <i>domain-dir/applications/j2ee-apps/app-name</i> or <i>domain-dir/applications/j2ee-modules/module-name</i>.</p> <p>The stubs directory is <i>domain-dir/generated/ejb/j2ee-apps/app-name</i> or <i>domain-dir/generated/ejb/j2ee-modules/module-name</i>.</p>
Web	The Web class loader loads the servlets and other classes in a specific enabled web module or a Java EE application that contains a web module. This class loader is present in each class loader universe that contains a web module; see <a href="#">“Class Loader Universes” on page 38</a> . One instance is created for each web module. The Web class loader is created with a list of URLs that point to the locations of the classes it needs to load. The classes it loads are in WEB-INF/classes or WEB-INF/lib/*.jar. It is parent to the JSP Engine class loader.
JSP Engine	The JSP Engine class loader loads compiled JSP classes of enabled JSP files. This class loader is present in each class loader universe that contains a JSP page; see <a href="#">“Class Loader Universes” on page 38</a> . The JSP Engine class loader is created with a list of URLs that point to the locations of the classes it needs to load.

## Delegation

Note that the class loader hierarchy is not a Java inheritance hierarchy, but a delegation hierarchy. In the delegation design, a class loader delegates classloading to its parent before attempting to load a class itself. A class loader parent can be either the System class loader or another custom class loader. If the parent class loader cannot load a class, the class loader attempts to load the class itself. In effect, a class loader is responsible for loading only the classes not available to the parent. Classes loaded by a class loader higher in the hierarchy cannot refer to classes available lower in the hierarchy.

The Java Servlet specification recommends that the Web class loader look in the local class loader before delegating to its parent. You can make the Web class loader follow the delegation inversion model in the Servlet specification by setting `delegate="false"` in the `class-loader` element of the `sun-web.xml` file. It is safe to do this only for a web module that does not interact with any other modules. For details, see “[class-loader](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

The default value is `delegate="true"`, which causes the Web class loader to delegate in the same manner as the other class loaders. You must use `delegate="true"` for a web application that accesses EJB components or that acts as a web service client or endpoint. For details about `sun-web.xml`, see *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## Using the Java Optional Package Mechanism

Optional packages are packages of Java classes and associated native code that application developers can use to extend the functionality of the core platform.

To use the Java optional package mechanism, copy the JAR files into the `domain-dir/lib/ext` directory, then restart the server.

For more information, see [Optional Packages - An Overview](#) (<http://java.sun.com/javase/6/docs/technotes/guides/extensions/extensions.html>) and [Understanding Extension Class Loading](#) (<http://java.sun.com/docs/books/tutorial/ext/basics/load.html>).

## Using the Endorsed Standards Override Mechanism

Endorsed standards handle changes to classes and APIs that are bundled in the JDK but are subject to change by external bodies.

To use the endorsed standards override mechanism, copy the JAR files into the `domain-dir/lib/endorsed` directory, then restart the server.

For more information and the list of packages that can be overridden, see [Endorsed Standards Override Mechanism](http://java.sun.com/javase/6/docs/technotes/guides/standards/) (<http://java.sun.com/javase/6/docs/technotes/guides/standards/>).

## Class Loader Universes

Access to components within applications and modules installed on the server occurs within the context of isolated class loader universes, each of which has its own Application, EJB, Web, and JSP Engine class loaders.

- **Application Universe** – Each Java EE application has its own class loader universe, which loads the classes in all the modules in the application.
- **Individually Deployed Module Universe** – Each individually deployed EJB JAR, web WAR, or lifecycle module has its own class loader universe, which loads the classes in the module.

A resource such as a file that is accessed by a servlet, JSP, or EJB component must be in one of the following locations:

- A directory pointed to by the Libraries field or `--libraries` option used during deployment
- A directory pointed to by the `library-directory` element in the `application.xml` deployment descriptor
- A directory pointed to by the class loader's classpath; for example, the web class loader's classpath includes these directories:

```
module-name/WEB-INF/classes  
module-name/WEB-INF/lib
```

## Application-Specific Class Loading

You can specify application-specific library classes during deployment in one of the following ways:

- Use the Admin Console. Open the Applications component, then go to the page for the type of application or module. Select the Deploy button. Type the comma-separated paths in the Libraries field. For details, click the Help button in the Admin Console.
- Use the `asadmin deploy` command with the `--libraries` option and specify comma-separated paths. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

Application libraries are included in the Application class loader. Paths to libraries can be relative or absolute. A relative path is relative to `domain-dir/lib/applibs`. If the path is absolute, the path must be accessible to the domain administration server (DAS). The

---

Enterprise Server automatically synchronizes these libraries to all remote cluster instances when the cluster is restarted. However, libraries specified by absolute paths are not guaranteed to be synchronized.

---

**Tip** – You can use application-specific class loading to specify a different XML parser than the default Enterprise Server XML parser. For details, see [http://blogs.sun.com/sivakumart/entry/classloaders\\_in\\_glassfish\\_an\\_attempt](http://blogs.sun.com/sivakumart/entry/classloaders_in_glassfish_an_attempt).

You can also use application-specific class loading to access different versions of a library from different applications.

---

If multiple applications or modules refer to the same libraries, classes in those libraries are automatically shared. This can reduce the memory footprint and allow sharing of static information. However, applications or modules using application-specific libraries are not portable. Other ways to make libraries available are described in “[Circumventing Class Loader Isolation](#)” on page 39.

For general information about deployment, see the *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

---

**Note** – If you see an access control error message when you try to use a library, you may need to grant permission to the library in the server .policy file. For more information, see “[Changing Permissions for an Application](#)” on page 89.

---

## Circumventing Class Loader Isolation

Since each application or individually deployed module class loader universe is isolated, an application or module cannot load classes from another application or module. This prevents two similarly named classes in different applications from interfering with each other.

To circumvent this limitation for libraries, utility classes, or individually deployed modules accessed by more than one application, you can include the relevant path to the required classes in one of these ways:

- “[Using the System Class Loader](#)” on page 40
- “[Using the Common Class Loader](#)” on page 40
- “[Sharing Libraries Across a Cluster](#)” on page 40
- “[Packaging the Client JAR for One Application in Another Application](#)” on page 41

Using the System class loader or Common class loader requires a server restart and makes a library accessible to all applications or modules deployed on servers that share the same configuration.

## Using the System Class Loader

To use the System class loader, do one of the following, then restart the server:

- Use the Admin Console. In the developer profile, select the Enterprise Server component and select the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. Then select the Path Settings tab and edit the System Classpath field. For details, click the Help button in the Admin Console.
- Edit the `system-classpath` attribute of the `java-config` element in the `domain.xml` file. For details about `domain.xml`, see the [Sun GlassFish Enterprise Server 2.1 Administration Reference](#).

Using the System class loader makes an application or module accessible to all applications or modules deployed on servers that share the same configuration.

Add the classes to the `system-classpath` attribute of the DAS in addition to the `system-classpath` attribute on the server instances that use the classes. The default name for the DAS configuration is `server-config`.

## Using the Common Class Loader

To use the Common class loader, copy the JAR files into the `domain-dir/lib` directory or copy the `.class` files into the `domain-dir/lib/classes` directory, then restart the server.

Using the Common class loader makes an application or module accessible to all applications or modules deployed on servers that share the same configuration.

For example, using the Common class loader is the recommended way of adding JDBC drivers to the Enterprise Server. For a list of the JDBC drivers currently supported by the Enterprise Server, see the [Sun GlassFish Enterprise Server 2.1 Release Notes](#). For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

## Sharing Libraries Across a Cluster

To share libraries across a specific cluster, copy the JAR files to the `domain-dir/config/cluster-config-name/lib` directory. Then add the path to the JAR files to the System class loader as explained in “Using the System Class Loader” on page 40 or to the Shared Chain class loader as explained in Table 2–1.



---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

## Packaging the Client JAR for One Application in Another Application

By packaging the client JAR for one application in a second application, you allow an EJB or web component in the second application to call an EJB component in the first (dependent) application, without making either of them accessible to any other application or module.

As an alternative for a production environment, you can have the Common class loader load the client JAR of the dependent application as described in “Using the Common Class Loader” on page 40. Restart the server to make the dependent application accessible to all applications or modules deployed on servers that share the same configuration.

### ▼ To Package the Client JAR for One Application in Another Application

- 1 Deploy the dependent application.
- 2 Add the dependent application’s client JAR file to the calling application.
  - For a calling EJB component, add the client JAR file at the same level as the EJB component. Then add a `Class-Path` entry to the `MANIFEST.MF` file of the calling EJB component. The `Class-Path` entry has this syntax:

```
Class-Path: filepath1.jar filepath2.jar ...
```

Each *filepath* is relative to the directory or JAR file containing the `MANIFEST.MF` file. For details, see the Java EE specification.

- For a calling web component, add the client JAR file under the `WEB-INF/lib` directory.
- 3 If you need to package the client JAR with both the EJB and web components, set `delegate="true"` in the `class-loader` element of the `sun-web.xml` file.

This changes the Web class loader so that it follows the standard class loader delegation model and delegates to its parent before attempting to load a class itself.

For most applications, packaging the client JAR file with the calling EJB component is sufficient. You do not need to package the client JAR file with both the EJB and web components unless the web component is directly calling the EJB component in the dependent application.

**4 Deploy the calling application.**

The calling EJB or web component must specify in its `sun-ejb-jar.xml` or `sun-web.xml` file the JNDI name of the EJB component in the dependent application. Using an `ejb-link` mapping does not work when the EJB component being called resides in another application.

You do not need to restart the server.

## The asant Utility

---

Apache Ant 1.6.5 is provided with Enterprise Server and can be launched from the `bin` directory using the command `asant`. The Enterprise Server also provides server-specific tasks, which are described in this section.

Make sure you have done these things before using `asant`:

1. Include `as-install/bin` in the `PATH` environment variable (`/usr/sfw/bin` for Sun Java™ Enterprise System, or Java ES, on Solaris). The Ant script provided with the Enterprise Server, `asant`, is located in this directory. For details on how to use `asant`, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).
2. If you are executing platform-specific applications, such as the `exec` or `cvs` task, the `ANT_HOME` environment variable must be set to the Ant installation directory.
  - The `ANT_HOME` environment variable for Java ES on Solaris is `/usr/sfw` and must include the following paths.
    - `/usr/sfw/bin` – the Ant binaries (shell scripts)
    - `/usr/sfw/doc/ant` – HTML documentation
    - `/usr/sfw/lib/ant` – Java classes that implement Ant
  - The `ANT_HOME` environment variable for all other platforms is `as-install/lib`.
3. Set up your password file. The argument for the `passwordfile` option of each Ant task is a file. This file contains the password in the following format.

```
AS_ADMIN_PASSWORD=password
```

For more information about password files, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

This section covers the following `asant`-related topics:

- “Enterprise Server `asant` Tasks” on page 44
- “Reusable Subelements” on page 63
- “JBI Tasks” on page 68

For more information about Ant, see the Apache Software Foundation web site at <http://ant.apache.org/>.

For information about standard Ant tasks, see the Ant documentation at <http://ant.apache.org/manual/>.

---

**Note** – Variables in the examples in this chapter, such as `${asinstalldir}`, reference values defined in `build.xml` or properties files.

---

## Enterprise Server asant Tasks

Use the asant tasks provided by the Enterprise Server for assembling, deploying, and undeploying modules and applications, and for configuring the server. The tasks are as follows:

- “The sun-appserv-deploy Task” on page 44
- “The sun-appserv-undeploy Task” on page 48
- “The sun-appserv-instance Task” on page 51
- “The sun-appserv-component Task” on page 54
- “The sun-appserv-admin Task” on page 57
- “The sun-appserv-jspc Task” on page 58
- “The sun-appserv-update Task” on page 59
- “The wsgen Task” on page 60
- “The wsimport Task” on page 61

### The sun-appserv-deploy Task

Deploys any of the following to a local or remote Enterprise Server instance.

- Enterprise application (EAR file)
- Web application (WAR file)
- Enterprise Java Bean (EJB-JAR file)
- Enterprise connector (RAR file)
- Application client

#### Subelements of sun-appserv-deploy

The following table describes subelements for the sun-appserv-deploy task. These are objects upon which this task acts.

TABLE 3-1 The sun-appserv-deploy Subelements

Element	Description
“The server Subelement” on page 63	An Enterprise Server instance
“The component Subelement” on page 65	A component to be deployed
“The fileset Subelement” on page 68	A set of component files that match specified parameters

## Attributes of sun-appserv-deploy

The following table describes attributes for the sun-appserv-deploy task.

TABLE 3-2 The sun-appserv-deploy Attributes

Attribute	Default	Description
file	none	(optional if a component or fileset subelement is present, otherwise required) The component to deploy. If this attribute refers to a file, it must be a valid archive. If this attribute refers to a directory, it must contain a valid archive in which all components have been exploded. If upload is set to false, this must be an absolute path on the server machine.
name	file name without extension	(optional) The display name for the component being deployed.
force	true	(optional) If true, the component is overwritten if it already exists on the server. If false, sun-appserv-deploy fails if the component exists.
retrievestubs	client stubs not saved	(optional) The directory where client stubs are saved. This attribute is inherited by nested component elements.
precompilejsp	false	(optional) If true, all JSP files found in an enterprise application (.ear) or web application (.war) are precompiled. This attribute is ignored for other component types. This attribute is inherited by nested component elements.
verify	false	(optional) If true, syntax and semantics for all deployment descriptors are automatically verified for correctness. This attribute is inherited by nested component elements.
contextroot	file name without extension	(optional) The context root for a web module (WAR file). This attribute is ignored if the component is not a WAR file.

TABLE 3-2 The sun-appserv-deploy Attributes (Continued)

Attribute	Default	Description
dbvendorname	sun-ejb-jar.xml entry	<p>(optional) The name of the database vendor for which tables can be created. Allowed values are <code>java</code>, <code>db2</code>, <code>mssql</code>, <code>oracle</code>, <code>postgresql</code>, <code>pointbase</code>, <code>derby</code> (also for CloudScape), and <code>sybase</code>, case-insensitive.</p> <p>If not specified, the value of the <code>database-vendor-name</code> attribute in <code>sun-ejb-jar.xml</code> is used.</p> <p>If no value is specified, a connection is made to the resource specified by the <code>jndi-name</code> subelement of the <code>cmp-resource</code> element in the <code>sun-ejb-jar.xml</code> file, and the database vendor name is read. If the connection cannot be established, or if the value is not recognized, SQL-92 compliance is presumed.</p> <p>For details, see <a href="#">“Generation Options for CMP” on page 195</a>.</p>
createtables	sun-ejb-jar.xml entry	<p>(optional) If <code>true</code>, causes database tables to be created for beans that need them. If <code>false</code>, does not create tables. If not specified, the value of the <code>create-tables-at-deploy</code> attribute in <code>sun-ejb-jar.xml</code> is used.</p> <p>For details, see <a href="#">“Generation Options” on page 122</a> and <a href="#">“Generation Options for CMP” on page 195</a>.</p>
dropandcreatetables	sun-ejb-jar.xml entry	<p>(optional) If <code>true</code>, and if tables were automatically created when this application was last deployed, tables from the earlier deployment are dropped and fresh ones are created.</p> <p>If <code>true</code>, and if tables were <i>not</i> automatically created when this application was last deployed, no attempt is made to drop any tables. If tables with the same names as those that would have been automatically created are found, the deployment proceeds, but a warning indicates that tables could not be created.</p> <p>If <code>false</code>, settings of <code>create-tables-at-deploy</code> or <code>drop-tables-at-undeploy</code> in the <code>sun-ejb-jar.xml</code> file are overridden.</p> <p>For details, see <a href="#">“Generation Options” on page 122</a> and <a href="#">“Generation Options for CMP” on page 195</a>.</p>
uniquetablenames	sun-ejb-jar.xml entry	<p>(optional) If <code>true</code>, specifies that table names are unique within each application server domain. If not specified, the value of the <code>use-unique-table-names</code> property in <code>sun-ejb-jar.xml</code> is used.</p> <p>For details, see <a href="#">“Generation Options for CMP” on page 195</a>.</p>
enabled	true	(optional) If <code>true</code> , enables the component.
deploymentplan	none	(optional) A deployment plan is a JAR file containing Sun-specific descriptors. Use this attribute when deploying an EAR file that lacks Sun-specific descriptors.
availabilityenabled	false	(optional) If <code>true</code> , enables high availability features, including persistence of HTTP sessions and checkpointing of the stateful session bean state.
generatermistubs	false	(optional) If <code>true</code> , generates the static RMI-IIOP stubs and puts them in the client JAR file.

TABLE 3-2 The sun-appserv-deploy Attributes (Continued)

Attribute	Default	Description
upload	true	(optional) If true, the component is transferred to the server for deployment. If the component is being deployed on the local machine, set upload to false to reduce deployment time. If a directory is specified for deployment, upload must be false.
virtualservers	default virtual server only	(optional) A comma-separated list of virtual servers to be deployment targets. This attribute applies only to application (.ear) or web (.war) components and is ignored for other component types. This attribute is inherited by nested server elements.
user	admin	(optional) The user name used when logging into the application server administration instance. This attribute is inherited by nested server elements.
passwordfile	none	(optional) File containing passwords. The password from this file is retrieved for communication with the application server administration instance. This attribute is inherited by nested server elements.
host	localhost	(optional) Target server. When deploying to a remote server, use the fully qualified host name. This attribute is inherited by nested server elements.
port	4848	(optional) The administration port on the target server. This attribute is inherited by nested server elements.
target	name of default instance	(optional) Target application server instance. This attribute is inherited by nested server elements.
asinstalldir	see description	(optional) The installation directory for the local Enterprise Server installation, which is used to find the administrative classes. If not specified, the command checks if the asinstalldir parameter has been set. Otherwise, administrative classes must be in the system classpath.

## Examples of sun-appserv-deploy

Here is a simple application deployment script with many implied attributes:

```
<sun-appserv-deploy
file="${assemble}/simpleapp.ear"
passwordfile="${passwordfile}" />
```

Here is an equivalent script showing all the implied attributes:

```
<sun-appserv-deploy
file="${assemble}/simpleapp.ear"
name="simpleapp"
force="true"
precompilejsp="false"
verify="false"
upload="true"
user="admin"
passwordfile="${passwordfile}"
host="localhost"
```

```
port="4848"  
target="{default-instance-name}"  
asinstalldir="{asinstalldir}" />
```

This example deploys multiple components to the same Enterprise Server instance running on a remote server:

```
<sun-appserv-deploy passwordfile="{passwordfile}" host="greg.sun.com"  
  asinstalldir="/opt/sun" >  
  <component file="{assemble}/simpleapp.ear"/>  
  <component file="{assemble}/simpleservlet.war"  
    contextroot="test"/>  
  <component file="{assemble}/simplebean.jar"/>  
</sun-appserv-deploy>
```

This example deploys multiple components to two Enterprise Server instances running on remote servers. In this example, both servers are using the same admin password. If this were not the case, each password could be specified in the server element.

```
<sun-appserv-deploy passwordfile="{passwordfile}" asinstalldir="/opt/sun" >  
  <server host="greg.sun.com"/>  
  <server host="joe.sun.com"/>  
  <component file="{assemble}/simpleapp.ear"/>  
  <component file="{assemble}/simpleservlet.war"  
    contextroot="test"/>  
  <component file="{assemble}/simplebean.jar"/>  
</sun-appserv-deploy>
```

This example deploys the same components as the previous example because the three components match the `fileset` criteria, but note that it is not possible to set some component-specific attributes. All component-specific attributes (name and context root) use their default values.

```
<sun-appserv-deploy passwordfile="{passwordfile}" host="greg.sun.com"  
  asinstalldir="/opt/sun" >  
  <fileset dir="{assemble}" includes="**/*.?ar" />  
</sun-appserv-deploy>
```

## The sun-appserv-undeploy Task

Undeploys any of the following from a local or remote Enterprise Server instance.

- Enterprise application (EAR file)
- Web application (WAR file)
- Enterprise Java Bean (EJB-JAR file)
- Enterprise connector (RAR file)



- Application client

## Subelements of sun-appserv-undeploy

The following table describes subelements for the sun-appserv-undeploy task. These are objects upon which this task acts.

TABLE 3-3 The sun-appserv-undeploy Subelements

Element	Description
<a href="#">“The server Subelement” on page 63</a>	An Enterprise Server instance
<a href="#">“The component Subelement” on page 65</a>	A component to be deployed
<a href="#">“The filesset Subelement” on page 68</a>	A set of component files that match specified parameters

## Attributes of sun-appserv-undeploy

The following table describes attributes for the sun-appserv-undeploy task.

TABLE 3-4 The sun-appserv-undeploy Attributes

Attribute	Default	Description
name	file name without extension	(optional if a component or filesset subelement is present or the file attribute is specified, otherwise required) The display name for the component being undeployed.
file	none	(optional) The component to undeploy. If this attribute refers to a file, it must be a valid archive. If this attribute refers to a directory, it must contain a valid archive in which all components have been exploded.
droptables	sun-ejb-jar.xml entry	(optional) If true, causes database tables that were automatically created when the bean(s) were last deployed to be dropped when the bean(s) are undeployed. If false, does not drop tables.  If not specified, the value of the drop-tables-at-undeploy attribute in sun-ejb-jar.xml is used.  For details, see <a href="#">“Generation Options” on page 122</a> and <a href="#">“Generation Options for CMP” on page 195</a> .
cascade	false	(optional) If true, deletes all connection pools and connector resources associated with the resource adapter being undeployed.  If false, undeployment fails if any pools or resources are still associated with the resource adapter.  This attribute is applicable to connectors (resource adapters) and applications with connector modules.

TABLE 3-4 The sun-appserv-undeploy Attributes (Continued)

Attribute	Default	Description
user	admin	(optional) The user name used when logging into the application server administration instance. This attribute is inherited by nested server elements.
passwordfile	none	(optional) File containing passwords. The password from this file is retrieved for communication with the application server administration instance. This attribute is inherited by nested server elements.
host	localhost	(optional) Target server. When deploying to a remote server, use the fully qualified host name. This attribute is inherited by nested server elements.
port	4848	(optional) The administration port on the target server. This attribute is inherited by nested server elements.
target	name of default instance	(optional) Target application server instance. This attribute is inherited by nested server elements.
asinstalldir	see description	(optional) The installation directory for the local Enterprise Server installation, which is used to find the administrative classes. If not specified, the command checks to see if the asinstalldir parameter has been set. Otherwise, administrative classes must be in the system classpath.

## Examples of sun-appserv-undeploy

Here is a simple application undeployment script with many implied attributes:

```
<sun-appserv-undeploy name="simpleapp" passwordfile="${passwordfile}" />
```

Here is an equivalent script showing all the implied attributes:

```
<sun-appserv-undeploy
  name="simpleapp"
  user="admin"
  passwordfile="${passwordfile}"
  host="localhost"
  port="4848"
  target="${default-instance-name}"
  asinstalldir="${asinstalldir}" />
```

This example demonstrates using the archive files (EAR and WAR, in this case) for the undeployment, using the component name (for undeploying the EJB component in this example), and undeploying multiple components.

```
<sun-appserv-undeploy passwordfile="${passwordfile}">
  <component file="${assemble}/simpleapp.ear"/>
  <component file="${assemble}/simpleservlet.war"/>
  <component name="simplebean" />
</sun-appserv-undeploy>
```

As with the deployment process, components can be undeployed from multiple servers in a single command. This example shows the same three components being removed from two different instances of the Enterprise Server. In this example, the passwords for both instances are the same.

```
<sun-appserv-undeploy passwordfile="{passwordfile}">
  <server host="greg.sun.com"/>
  <server host="joe.sun.com"/>
  <component file="{assemble}/simpleapp.ear"/>
  <component file="{assemble}/simpleservlet.war"/>
  <component name="simplebean" />
</sun-appserv-undeploy>
```

## The sun-appserv-instance Task

Starts, stops, restarts, creates, or removes one or more application server instances.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see [“Usage Profiles” in Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

---

### Subelements of sun-appserv-instance

The following table describes subelements for the sun-appserv-instance task. These are objects upon which this task acts.

TABLE 3-5 The sun-appserv-instance Subelements

Element	Description
<a href="#">“The server Subelement” on page 63</a>	An Enterprise Server instance

### Attributes of sun-appserv-instance

The following table describes attributes for the sun-appserv-instance task.

TABLE 3-6 The sun-appserv-instance Attributes

Attribute	Default	Description
action	none	The control command for the target application server. Valid values are start, stop, create, and delete. A restart sends the stop command followed by the start command. The restart command is not supported on Windows.

TABLE 3-6 The sun-appserv-instance Attributes (Continued)

Attribute	Default	Description
debug	false	(optional) Deprecated. If <code>action</code> is set to <code>start</code> , specifies whether the server starts in debug mode. This attribute is ignored for other values of <code>action</code> . If <code>true</code> , the instance generates additional debugging output throughout its lifetime. This attribute is inherited by nested server elements.
nodeagent	none	(required if <code>action</code> is <code>create</code> , otherwise ignored) The name of the node agent on which the instance is being created.
cluster	none	(optional, applicable only if <code>action</code> is <code>create</code> ) The clustered instance to be created. The server's configuration is inherited from the named cluster.  The <code>config</code> and <code>cluster</code> attributes are mutually exclusive. If both are omitted, a stand-alone server instance is created.
config	none	(optional, applicable only if <code>action</code> is <code>create</code> ) The configuration for the new stand-alone instance.  The configuration must exist and must not be <code>default-config</code> (the cluster configuration template) or an already referenced stand-alone configuration (including the administration server configuration <code>server-config</code> ).  The <code>config</code> and <code>cluster</code> attributes are mutually exclusive. If both are omitted, a stand-alone server instance is created.
property	none	(optional, applicable only if <code>action</code> is <code>create</code> ) Defines system properties for the server instance. These properties override port settings in the server instance's configuration. The following properties are defined: <code>http-listener-1-port</code> , <code>http-listener-2-port</code> , <code>orb-listener-1-port</code> , <code>SSL-port</code> , <code>SSL_MUTUALAUTH-port</code> , <code>JMX_SYSTEM_CONNECTOR_port</code> .  System properties can be changed after instance creation using the system property commands. For details, see the <a href="#">Sun GlassFish Enterprise Server 2.1 Reference Manual</a> .
user	admin	(optional) The user name used when logging into the application server administration instance. This attribute is inherited by nested server elements.
passwordfile	none	(optional) File containing passwords. The password from this file is retrieved for communication with the application server administration instance. This attribute is inherited by nested server elements.
host	localhost	(optional) Target server. If it is a remote server, use the fully qualified host name. This attribute is inherited by nested server elements.
port	4848	(optional) The administration port on the target server. This attribute is inherited by nested server elements.
instance	name of default instance	(optional) Target application server instance. This attribute is inherited by nested server elements.
asinstalldir	see description	(optional) The installation directory for the local Enterprise Server installation, which is used to find the administrative classes. If not specified, the command checks to see if the <code>asinstalldir</code> parameter has been set. Otherwise, administrative classes must be in the system classpath.

## Examples of sun-appserv-instance

This example starts the local Enterprise Server instance:

```
<sun-appserv-instance action="start" passwordfile="${passwordfile}"
  instance="${default-instance-name}"/>
```

Here is an equivalent script showing all the implied attributes:

```
<sun-appserv-instance
  action="start"
  user="admin"
  passwordfile="${passwordfile}"
  host="localhost"
  port="4848"
  instance="${default-instance-name}"
  asinstalldir="${asinstalldir}" />
```

Multiple servers can be controlled using a single command. In this example, two servers are restarted, and in this case each server uses a different password:

```
<sun-appserv-instance action="restart"
  instance="${default-instance-name}"/>
<server host="greg.sun.com" passwordfile="${password.greg}"/>
<server host="joe.sun.com" passwordfile="${password.joe}"/>
</sun-appserv-instance>
```

This example creates a new Enterprise Server instance:

```
<sun-appserv-instance
  action="create" instanceport="8080"
  passwordfile="${passwordfile}"
  instance="development" />
```

Here is an equivalent script showing all the implied attributes:

```
<sun-appserv-instance
  action="create"
  instanceport="8080"
  user="admin"
  passwordfile="${passwordfile}"
  host="localhost"
  port="4848"
  instance="development"
  asinstalldir="${asinstalldir}" />
```

Instances can be created on multiple servers using a single command. This example creates a new instance named qa on two different servers. In this case, both servers use the same password.

```
<sun-appserv-instance
  action="create"
  instanceport="8080"
  instance="qa"
  passwordfile="{passwordfile}">
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
</sun-appserv-instance>
```

These instances can also be removed from their respective servers:

```
<sun-appserv-instance
  action="delete"
  instance="qa"
  passwordfile="{passwordfile}">
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
</sun-appserv-instance>
```

Different instance names and instance ports can also be specified using attributes of the server subelement:

```
<sun-appserv-instance action="create" passwordfile="{passwordfile}">
<server host="greg.sun.com" instanceport="8080" instance="qa"/>
<server host="joe.sun.com" instanceport="9090"
  instance="integration-test"/>
</sun-appserv-instance>
```

## The sun - appserv - component Task

Enables or disables the following Java EE component types that have been deployed to the Enterprise Server.

- Enterprise application (EAR file)
- Web application (WAR file)
- Enterprise Java Bean (EJB-JAR file)
- Enterprise connector (RAR file)
- Application client

You do not need to specify the archive to enable or disable a component: only the component name is required. You can use the component archive, however, because it implies the component name.

### Subelements of sun - appserv - component

The following table describes subelements for the sun - appserv - component task. These are objects upon which this task acts.

TABLE 3-7 The sun-appserv-component Subelements

Element	Description
<a href="#">“The server Subelement” on page 63</a>	An Enterprise Server instance
<a href="#">“The component Subelement” on page 65</a>	A component to be deployed
<a href="#">“The fileset Subelement” on page 68</a>	A set of component files that match specified parameters

## Attributes of sun-appserv-component

The following table describes attributes for the sun-appserv-component task.

TABLE 3-8 The sun-appserv-component Attributes

Attribute	Default	Description
action	none	The control command for the target application server. Valid values are <code>enable</code> and <code>disable</code> .
name	file name without extension	(optional if a component or <code>fileset</code> subelement is present or the <code>file</code> attribute is specified, otherwise required) The display name for the component being enabled or disabled.
file	none	(optional) The component to enable or disable. If this attribute refers to a file, it must be a valid archive. If this attribute refers to a directory, it must contain a valid archive in which all components have been exploded.
user	admin	(optional) The user name used when logging into the application server administration instance. This attribute is inherited by nested server elements.
passwordfile	none	(optional) File containing passwords. The password from this file is retrieved for communication with the application server administration instance. This attribute is inherited by nested server elements.
host	localhost	(optional) Target server. When enabling or disabling a remote server, use the fully qualified host name. This attribute is inherited by nested server elements.
port	4848	(optional) The administration port on the target server. This attribute is inherited by nested server elements.
target	name of default instance	(optional) Target application server instance. This attribute is inherited by nested server elements.
asinstalldir	see description	(optional) The installation directory for the local Enterprise Server installation, which is used to find the administrative classes. If not specified, the command checks to see if the <code>asinstalldir</code> parameter has been set. Otherwise, administrative classes must be in the system classpath.

## Examples of sun-appserv-component

Here is a simple example of disabling a component:

```
<sun-appserv-component
  action="disable"
  name="simpleapp"
  passwordfile="{passwordfile}" />
```

Here is a simple example of enabling a component:

```
<sun-appserv-component
  action="enable"
  name="simpleapp"
  passwordfile="{passwordfile}" />
```

Here is an equivalent script showing all the implied attributes:

```
<sun-appserv-component
  action="enable"
  name="simpleapp"
  user="admin"
  passwordfile="{passwordfile}"
  host="localhost"
  port="4848"
  target="{default-instance-name}"
  asinstalldir="{asinstalldir}" />
```

This example demonstrates disabling multiple components using the archive files (EAR and WAR, in this case) and using the component name (for an EJB component in this example).

```
<sun-appserv-component action="disable" passwordfile="{passwordfile}">
  <component file="{assemble}/simpleapp.ear"/>
  <component file="{assemble}/simpleservlet.war"/>
  <component name="simplebean" />
</sun-appserv-component>
```

Components can be enabled or disabled on multiple servers in a single task. This example shows the same three components being enabled on two different instances of the Enterprise Server. In this example, the passwords for both instances are the same.

```
<sun-appserv-component action="enable" passwordfile="{passwordfile}">
  <server host="greg.sun.com"/>
  <server host="joe.sun.com"/>
  <component file="{assemble}/simpleapp.ear"/>
  <component file="{assemble}/simpleservlet.war"/>
  <component name="simplebean" />
</sun-appserv-component>
```



## The sun-appserv-admin Task

Enables arbitrary administrative commands and scripts to be executed on the Enterprise Server. This is useful for cases where a specific Ant task has not been developed or a set of related commands are in a single script.

### Subelements of sun-appserv-admin

The following table describes subelements for the sun-appserv-admin task. These are objects upon which this task acts.

TABLE 3-9 The sun-appserv-admin Subelements

Element	Description
<a href="#">“The server Subelement” on page 63</a>	An Enterprise Server instance

### Attributes of sun-appserv-admin

The following table describes attributes for the sun-appserv-admin task.

TABLE 3-10 The sun-appserv-admin Attributes

Attribute	Default	Description
command	none	(exactly one of these is required: <code>command</code> or <code>explicitcommand</code> ) The command to execute. If the <code>user</code> , <code>passwordfile</code> , <code>host</code> , <code>port</code> , or <code>target</code> attributes are also specified, they are automatically inserted into the command before execution. If any of these options are specified in the command string, the corresponding attribute values are ignored.
explicitcommand	none	(exactly one of these is required: <code>command</code> or <code>explicitcommand</code> ) The exact command to execute. No command processing is done, and all other attributes are ignored.
user	admin	(optional) The user name used when logging into the application server administration instance. This attribute is inherited by nested server elements.
passwordfile	none	(optional) File containing passwords. The password from this file is retrieved for communication with the application server administration instance. This attribute is inherited by nested server elements.
host	localhost	(optional) Target server. If it is a remote server, use the fully qualified host name. This attribute is inherited by nested server elements.
port	4848	(optional) The administration port on the target server. This attribute is inherited by nested server elements.
asinstalldir	see description	(optional) The installation directory for the local Enterprise Server installation, which is used to find the administrative classes. If not specified, the command checks if the <code>asinstalldir</code> parameter has been set. Otherwise, administrative classes must be in the system classpath.

## Examples of sun-appserv-admin

Here is an example of executing the create-jms-dest command:

```
<sun-appserv-admin command="create-jms-dest --desttype topic">
```

Here is an example of using explicitcommand to execute the create-jms-dest command:

```
<sun-appserv-admin explicitcommand="create-jms-dest --user adminuser --host localhost
--port 4848 --desttype topic --target server1 simpleJmsDest">
```

## The sun-appserv-jspc Task

Precompiles JSP source code into Enterprise Server compatible Java code for initial invocation by Enterprise Server. Use this task to speed up access to JSP files or to check the syntax of JSP source code. You can feed the resulting Java code to the javac task to generate class files for the JSP files.

### Attributes of sun-appserv-jspc

The following table describes attributes for the sun-appserv-jspc task.

TABLE 3-11 The sun-appserv-jspc Attributes

Attribute	Default	Description
destdir	none	The destination directory for the generated Java source files.
srcdir	none	(exactly one of these is required: srcdir or webapp) The source directory where the JSP files are located.
webapp	none	(exactly one of these is required: srcdir or webapp) The directory containing the web application. All JSP files within the directory are recursively parsed. The base directory must have a WEB-INF subdirectory beneath it. When webapp is used, sun-appserv-jspc hands off all dependency checking to the compiler.
verbose	2	(optional) The verbosity integer to be passed to the compiler.
classpath	none	(optional) The classpath for running the JSP compiler.
classpathref	none	(optional) A reference to the JSP compiler classpath.
uribase	/	(optional) The URI context of relative URI references in the JSP files. If this context does not exist, it is derived from the location of the JSP file relative to the declared or derived value of uriroot. Only pages translated from an explicitly declared JSP file are affected.

TABLE 3-11 The sun-appserv-jspc Attributes (Continued)

Attribute	Default	Description
urifoot	see description	(optional) The root directory of the web application, against which URI files are resolved. If this directory is not specified, the first JSP file is used to derive it: each parent directory of the first JSP file is searched for a WEB-INF directory, and the directory closest to the JSP file that has one is used. If no WEB-INF directory is found, the directory from which sun-appserv-jspc was called is used. Only pages translated from an explicitly declared JSP file (including tag libraries) are affected.
package	none	(optional) The destination package for the generated Java classes.
asinstalldir	see description	(optional) The installation directory for the local Enterprise Server installation, which is used to find the administrative classes. If not specified, the command checks if the asinstalldir parameter has been set. Otherwise, administrative classes must be in the system classpath.

### Example of sun-appserv-jspc

The following example uses the webapp attribute to generate Java source files from JSP files. The sun-appserv-jspc task is immediately followed by a javac task, which compiles the generated Java files into class files. The classpath value in the javac task must be all on one line with no spaces.

```
<sun-appserv-jspc
  destdir="${assemble.war}/generated"
  webapp="${assemble.war}"
  classpath="${assemble.war}/WEB-INF/classes"
  asinstalldir="${asinstalldir}" />
<javac
  srcdir="${assemble.war}/WEB-INF/generated"
  destdir="${assemble.war}/WEB-INF/generated"
  debug="on"
  classpath="${assemble.war}/WEB-INF/classes:${asinstalldir}/lib/
    appserv-rt.jar:${asinstalldir}/lib/appserv-ext.jar">
  <include name="**/*.java"/>
</javac>
```

## The sun-appserv-update Task

Enables deployed applications (EAR files) and modules (EJB JAR, RAR, and WAR files) to be updated and reloaded for fast iterative development. This task copies modified class files, XML files, and other contents of the archive files to the appropriate subdirectory of the *domain-dir/applications* directory, then touches the `.reload` file to cause dynamic reloading to occur.

This is a local task and must be executed on the same machine as the Enterprise Server.

For more information about dynamic reloading, see the [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

## Attributes of sun-appserv-update

The following table describes attributes for the sun-appserv-update task.

TABLE 3-12 The sun-appserv-update Attributes

Attribute	Default	Description
file	none	The component to update, which must be a valid archive.
domain	domain1	(optional) The domain in which the application has been previously deployed.

## Example of sun-appserv-update

The following example updates the Java EE application `foo.ear`, which is deployed to the default domain, `domain1`.

```
<sun-appserv-update file="foo.ear"/>
```

## The wsgen Task

Generates JAX-WS portable artifacts used in JAX-WS web services. Reads a web service endpoint class and generates all the required artifacts for web service deployment and invocation.

## Attributes of wsgen

The following table describes attributes for the wsgen task.

TABLE 3-13 The wsgen Attributes

Attribute	Default	Description
sei	none	Specifies the name of the service endpoint interface (SEI) class.
destdir	current directory	(optional) Specifies where to place the output generated classes.
classpath	system classpath	(optional) Specifies where to find the input class files. Same as <code>cp</code> attribute.
cp	system classpath	(optional) Specifies where to find the input class files. Same as <code>classpath</code> attribute.
resourcedestdir	current directory	(optional) Specifies where to place generated resource files such as WSDL files. Used only if the <code>genwsdl</code> attribute is set to <code>true</code> .

TABLE 3-13 The `wsgen` Attributes (Continued)

Attribute	Default	Description
<code>sourcedestdir</code>	current directory	(optional) Specifies where to place generated source files.
<code>keep</code>	<code>false</code>	(optional) If <code>true</code> , keeps generated files.
<code>verbose</code>	<code>false</code>	(optional) If <code>true</code> , outputs compiler messages.
<code>genwsdl</code>	<code>true</code>	(optional) If <code>true</code> , generates a WSDL file.
<code>protocol</code>	<code>soap1.1</code>	(optional) Specifies the protocol to use in the <code>wsdl:binding</code> . Used only if the <code>genwsdl</code> attribute is set to <code>true</code> .  Allowed values are <code>soap1.1</code> or <code>Xsoap1.2</code> . <code>Xsoap1.2</code> is not standard and is only used if the <code>extension</code> attribute is set to <code>true</code> .
<code>servicename</code>	<code>none</code>	(optional) Specifies a particular <code>wsdl:service</code> name for the generated WSDL file. Used only if the <code>genwsdl</code> attribute is set to <code>true</code> . For example:  <code>servicename="{http://mynamespace/}MyService"</code>
<code>portname</code>	<code>none</code>	(optional) Specifies a particular <code>wsdl:port</code> name for the generated WSDL. Used only if the <code>genwsdl</code> attribute is set to <code>true</code> . For example:  <code>portname="{http://mynamespace/}MyPort"</code>
<code>extension</code>	<code>false</code>	(optional) If <code>true</code> , allows vendor extensions not in the specification. Use of extensions may result in applications that are not portable and may not interoperate with other implementations.

## Example of `wsgen`

The following example generates portable artifacts for `fromjava.server.AddNumbersImpl`, uses `compile.classpath` as the classpath, and writes the WSDL file to `${wsdl.dir}`.

```
<wsgen
  sourcedestdir="${wsdl.dir}"
  sei="fromjava.server.AddNumbersImpl">
  <classpath refid="compile.classpath"/>
</wsgen>
```

## The `wsimport` Task

Generates JAX-WS portable artifacts for a given WSDL file. Portable artifacts include service endpoint interfaces (SEIs), services, exception classes mapped from the `wsdl:fault` and `soap:headerfault` tags, asynchronous response beans derived from the `wsdl:message` tag, and JAXB generated value types. After generation, these artifacts can be packaged in a WAR file with the WSDL and schema documents along with the endpoint implementation and then deployed.

## Attributes of `wsimport`

The following table describes attributes for the `wsimport` task.

TABLE 3-14 The `wsimport` Attributes

Attribute	Default	Description
<code>wSDL</code>	<code>none</code>	Specifies the name of the WSDL file.
<code>destDir</code>	<code>current directory</code>	(optional) Specifies where to place the output generated classes.
<code>sourcedestDir</code>	<code>current directory</code>	(optional) Specifies where to place generated source files. Used only if the <code>keep</code> attribute is set to <code>true</code> .
<code>keep</code>	<code>false</code>	(optional) If <code>true</code> , keeps generated files.
<code>verbose</code>	<code>false</code>	(optional) If <code>true</code> , outputs compiler messages.
<code>binding</code>	<code>none</code>	(optional) Specifies external JAX-WS or JAXB binding files. JAX-WS and JAXB binding files can customize things like package names and bean names. More information on JAX-WS and JAXB binding files can be found in the customization documentation included with this release.
<code>extension</code>	<code>false</code>	(optional) If <code>true</code> , allows vendor extensions not in the specification. Use of extensions may result in applications that are not portable and may not interoperate with other implementations.
<code>wSDLLocation</code>	<code>none</code>	(optional) Specifies the value of <code>@WebService.wSDLLocation</code> and <code>@WebServiceClient.wSDLLocation</code> annotation elements for the generated SEI and Service interface. This should be set to the URI of the web service WSDL file.
<code>catalog</code>	<code>none</code>	(optional) Specifies the catalog file to resolve external entity references. Supported formats are TR9401, XCatalog, and OASIS XML Catalog. Additionally, the Ant <code>xmlCatalog</code> type can be used to resolve entities.
<code>package</code>	<code>none</code>	(optional) Specifies the target package, overriding any WSDL and schema binding customization for package name, and the default package name algorithm defined in the JAX-WS specification.

## Example of `wsimport`

The following example generates client-side artifacts for `AddNumbers.wSDL` and stores `.class` files in the `${build.classes.home}` directory using the `custom.xml` customization file.

```
<wsimport
  destDir="${build.classes.home}"
  wSDL="AddNumbers.wSDL"
  binding="custom.xml">
</wsimport>
```

# Reusable Subelements

Reusable subelements of the Ant tasks for the Enterprise Server are as follows. These are objects upon which the Ant tasks act.

- “The server Subelement” on page 63
- “The component Subelement” on page 65
- “The filesset Subelement” on page 68

## The server Subelement

Specifies an Enterprise Server instance. Allows a single task to act on multiple server instances. The server attributes override corresponding attributes in the parent task; therefore, the parent task attributes function as default values.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

### Attributes of server

The following table describes attributes for the server element.

TABLE 3–15 The server Attributes

Attribute	Default	Description
user	admin	(optional) The user name used when logging into the Enterprise Server domain administration server (DAS).
passwordfile	none	(optional) File containing passwords. The password from this file is retrieved for communication with the Enterprise Server DAS.
host	localhost	(optional) Target server. When targeting a remote server, use the fully qualified host name.
port	4848	(optional) The administration port on the target server.
instance	name of default instance	(optional) Target application server instance.
instanceport	none	(applies to “The sun-appserv-instance Task” on page 51 only) Deprecated.
nodeagent	none	(applies to “The sun-appserv-instance Task” on page 51 only, required if action is create, otherwise ignored) The name of the node agent on which the instance is being created.

TABLE 3-15 The server Attributes (Continued)

Attribute	Default	Description
debug	false	(applies to “The sun-appserv-instance Task” on page 51 only, optional) Deprecated. If action is set to start, specifies whether the server starts in debug mode. This attribute is ignored for other values of action. If true, the instance generates additional debugging output throughout its lifetime.
upload	true	(applies to “The sun-appserv-deploy Task” on page 44 only, optional) If true, the component is transferred to the server for deployment. If the component is being deployed on the local machine, set upload to false to reduce deployment time.
virtualservers	default virtual server only	(applies to “The sun-appserv-deploy Task” on page 44 only, optional) A comma-separated list of virtual servers to be deployment targets. This attribute applies only to application (.ear) or web (.war) components and is ignored for other component types.

## Examples of server

You can control multiple servers using a single task. In this example, two servers are started, each using a different password. Only the second server is started in debug mode.

```
<sun-appserv-instance action="start">
  <server host="greg.sun.com" passwordfile="{password.greg}"/>
  <server host="joe.sun.com" passwordfile="{password.joe}"
    debug="true"/>
</sun-appserv-instance>
```

You can create instances on multiple servers using a single task. This example creates a new instance named qa on two different servers. Both servers use the same password.

```
<sun-appserv-instance action="create" instanceport="8080"
  instance="qa" passwordfile="{passwordfile}>
  <server host="greg.sun.com"/>
  <server host="joe.sun.com"/>
</sun-appserv-instance>
```

These instances can also be removed from their respective servers:

```
<sun-appserv-instance action="delete" instance="qa"
  passwordfile="{passwordfile}>
  <server host="greg.sun.com"/>
  <server host="joe.sun.com"/>
</sun-appserv-instance>
```

You can specify different instance names and instance ports using attributes of the nested server element:

```
<sun-appserv-instance action="create" passwordfile="{passwordfile}>
  <server host="greg.sun.com" instanceport="8080" instance="qa"/>
```



```
<server host="joe.sun.com" instanceport="9090"
  instance="integration-test"/>
</sun-appserv-instance>
```

You can deploy multiple components to multiple servers (see the [“The component Subelement” on page 65](#)). This example deploys each component to two Enterprise Server instances running on remote servers. Both servers use the same password.

```
<sun-appserv-deploy passwordfile="${passwordfile}"
  asinstalldir="/opt/slas8" >
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simpleservlet.war"
  contextroot="test"/>
<component file="${assemble}/simplebean.jar"/>
</sun-appserv-deploy>
```

You can also undeploy multiple components from multiple servers. This example shows the same three components being removed from two different instances. Both servers use the same password.

```
<sun-appserv-undeploy passwordfile="${passwordfile}">
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simpleservlet.war"/>
<component name="simplebean" />
</sun-appserv-undeploy>
```

You can enable or disable components on multiple servers. This example shows the same three components being enabled on two different instances. Both servers use the same password.

```
<sun-appserv-component action="enable" passwordfile="${passwordfile}">
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simpleservlet.war"/>
<component name="simplebean" />
</sun-appserv-component>
```

## The component Subelement

Specifies a Java EE component. Allows a single task to act on multiple components. The component attributes override corresponding attributes in the parent task; therefore, the parent task attributes function as default values.

## Attributes of component

The following table describes attributes for the component element.

TABLE 3-16 The component Attributes

Attribute	Default	Description
file	none	(optional if the parent task is “ <a href="#">The sun-appserv-undeploy Task</a> ” on page 48 or “ <a href="#">The sun-appserv-component Task</a> ” on page 54) The target component. If this attribute refers to a file, it must be a valid archive. If this attribute refers to a directory, it must contain a valid archive in which all components have been exploded. If <code>upload</code> is set to <code>false</code> , this must be an absolute path on the server machine.
name	file name without extension	(optional) The display name for the component.
force	true	(applies to “ <a href="#">The sun-appserv-deploy Task</a> ” on page 44 only, optional) If <code>true</code> , the component is overwritten if it already exists on the server. If <code>false</code> , the containing element’s operation fails if the component exists.
precompilejsp	false	(applies to “ <a href="#">The sun-appserv-deploy Task</a> ” on page 44 only, optional) If <code>true</code> , all JSP files found in an enterprise application ( <code>.ear</code> ) or web application ( <code>.war</code> ) are precompiled. This attribute is ignored for other component types.
retrievestubs	client stubs not saved	(applies to “ <a href="#">The sun-appserv-deploy Task</a> ” on page 44 only, optional) The directory where client stubs are saved.
contextroot	file name without extension	(applies to “ <a href="#">The sun-appserv-deploy Task</a> ” on page 44 only, optional) The context root for a web module (WAR file). This attribute is ignored if the component is not a WAR file.
verify	false	(applies to “ <a href="#">The sun-appserv-deploy Task</a> ” on page 44 only, optional) If <code>true</code> , syntax and semantics for all deployment descriptors is automatically verified for correctness.

## Examples of component

You can deploy multiple components using a single task. This example deploys each component to the same Enterprise Server instance running on a remote server.

```
<sun-appserv-deploy passwordfile="{passwordfile}" host="greg.sun.com"
  asinstalldir="/opt/slas8" >
<component file="{assemble}/simpleapp.ear"/>
<component file="{assemble}/simpleervlet.war"
  contextroot="test"/>
<component file="{assemble}/simplebean.jar"/>
</sun-appserv-deploy>
```

You can also undeploy multiple components using a single task. This example demonstrates using the archive files (EAR and WAR, in this case) and the component name (for the EJB component).

```
<sun-appserv-undeploy passwordfile="${passwordfile}">
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simplereservlet.war"/>
<component name="simplebean" />
</sun-appserv-undeploy>
```

You can deploy multiple components to multiple servers. This example deploys each component to two instances running on remote servers. Both servers use the same password.

```
<sun-appserv-deploy passwordfile="${passwordfile}" asinstalldir="/opt/slas8" >
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simplereservlet.war"
    contextroot="test"/>
<component file="${assemble}/simplebean.jar"/>
</sun-appserv-deploy>
```

You can also undeploy multiple components to multiple servers. This example shows the same three components being removed from two different instances. Both servers use the same password.

```
<sun-appserv-undeploy passwordfile="${passwordfile}">
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simplereservlet.war"/>
<component name="simplebean" />
</sun-appserv-undeploy>
```

You can enable or disable multiple components. This example demonstrates disabling multiple components using the archive files (EAR and WAR, in this case) and the component name (for the EJB component).

```
<sun-appserv-component action="disable" passwordfile="${passwordfile}">
<component file="${assemble}/simpleapp.ear"/>
<component file="${assemble}/simplereservlet.war"/>
<component name="simplebean" />
</sun-appserv-component>
```

You can enable or disable multiple components on multiple servers. This example shows the same three components being enabled on two different instances. Both servers use the same password.

```
<sun-appserv-component action="enable" passwordfile="{passwordfile}">
<server host="greg.sun.com"/>
<server host="joe.sun.com"/>
<component file="{assemble}/simpleapp.ear"/>
<component file="{assemble}/simpleservlet.war"/>
<component name="simplebean" />
</sun-appserv-component>
```

## The fileset Subelement

Selects component files that match specified parameters. When `fileset` is included as a subelement, the name and context root attributes of the containing element must use their default values for each file in the `fileset`. For more information, see <http://ant.apache.org/manual/CoreTypes/fileset.html>.

## JBI Tasks

The `asant` utility supports the Java Business Integration (JBI) Ant tasks. The Ant Tasks Reference is included with the Enterprise Server at *as-install/jbi/doc/antdoc/*.

For more information about JBI in the Enterprise Server, see “[JBI Runtime](#)” on page 111.

# Debugging Applications

---

This chapter gives guidelines for debugging applications in the Sun GlassFish Enterprise Server. It includes the following sections:

- “Enabling Debugging” on page 69
- “JPDA Options” on page 70
- “Generating a Stack Trace for Debugging” on page 71
- “Application Client Debugging” on page 71
- “Sun GlassFish Message Queue Debugging” on page 71
- “Enabling Verbose Mode” on page 72
- “Enterprise Server Logging” on page 72
- “Profiling Tools” on page 72

## Enabling Debugging

When you enable debugging, you enable both local and remote debugging. To start the server in debug mode, use the `--debug` option as follows:

```
asadmin start-domain --user adminuser --debug [domain-name]
```

You can then attach to the server from the Java Debugger (`jdb`) at its default Java Platform Debugger Architecture (JPDA) port, which is 9009. For example, for UNIX® systems:

```
jdb -attach 9009
```

For Windows:

```
jdb -connect com.sun.jdi.SocketAttach:port=9009
```

For more information about the `jdb` debugger, see the following links:

- Java Platform Debugger Architecture - The Java Debugger:  
<http://java.sun.com/products/jpda/doc/soljdb.html>

- Java Platform Debugger Architecture - Connecting with JDB:  
<http://java.sun.com/products/jpda/doc/conninv.html#JDB>

Enterprise Server debugging is based on the JPDA. For more information, see “[JPDA Options](#)” on page 70.

You can attach to the Enterprise Server using any JPDA compliant debugger, including that of [NetBeans](http://www.netbeans.org) (<http://www.netbeans.org>), Sun Java Studio, JBuilder, Eclipse, and so on.

You can enable debugging even when the application server is started without the `--debug` option. This is useful if you start the application server from the Windows Start Menu, or if you want to make sure that debugging is always turned on.

## ▼ To Set the Server to Automatically Start Up in Debug Mode

- 1 Use the Admin Console. In the developer profile, select the Enterprise Server component and the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration.
- 2 Check the Debug Enabled box.
- 3 To specify a different port (from 9009, the default) to use when attaching the JVM to a debugger, specify `address=port-number` in the Debug Options field.
- 4 To add JPDA options, add any desired JPDA debugging options in Debug Options. See “[JPDA Options](#)” on page 70.

**See Also** For details, click the Help button in the Admin Console from the JVM Settings page.

## JPDA Options

The default JPDA options in Enterprise Server are as follows:

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

For Windows, you can change `dt_socket` to `dt_shmem`.

If you substitute `suspend=y`, the JVM starts in suspended mode and stays suspended until a debugger attaches to it. This is helpful if you want to start debugging as soon as the JVM starts.

To specify a different port (from 9009, the default) to use when attaching the JVM to a debugger, specify `address=port-number`.

You can include additional options. A list of JPDA debugging options is available at <http://java.sun.com/products/jpda/doc/conninv.html#Invocation>.

## Generating a Stack Trace for Debugging

To generate a Java stack trace for debugging, use the `asadmin generate-jvm-report --type=thread` command. The stack trace goes to the `domain-dir/logs/server.log` file and also appears on the command prompt screen. For more information about the `asadmin generate-jvm-report` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Application Client Debugging

When the `appclient` script executes the `java` command to run the Application Client Container (ACC), which in turn runs the client, it includes on the command line the value of the `VMARGS` environment variable. You can set this variable to any suitable value. The following example debugging setup is for Windows systems:

```
VMARGS=-Xdebug -Xrunjdpw:transport=dt_socket,server=y,suspend=y,address=8118
```

The following example debugging setup is for UNIX-based systems:

```
set VMARGS=-Xdebug -Xrunjdpw:transport=dt_socket,server=y,suspend=y,address=8118
```

For debugging an application client, you should set `suspend` to `y` so you can connect the debugger to the client before any code has actually executed. Otherwise, the client may start running and execute past the point you want to examine.

You should use different ports for the server and client if you are debugging both concurrently. For details about setting the port, see “JPDA Options” on page 70.

For information about the `appclient` script, see *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Sun GlassFish Message Queue Debugging

Sun GlassFish Message Queue has a broker logger, which can be useful for debugging Java Message Service (JMS) applications, including message-driven bean applications. You can adjust the logger’s verbosity, and you can send the logger output to the broker’s console using the broker’s `-tty` option. For more information, see the *Sun Java System Message Queue 4.3 Administration Guide*.

## Enabling Verbose Mode

To have the server logs and messages printed to `System.out` on your command prompt screen, you can start the server in verbose mode. This makes it easy to do simple debugging using print statements, without having to view the `server.log` file every time.

To start the server in verbose mode, use the `--verbose` option as follows:

```
asadmin start-domain --user adminuser --verbose [domain-name]
```

On Windows platforms, you must perform an extra preparation step if you want to use Ctrl-Break to print a thread dump. In the `as-install/asenv.bat` file, change `AS_NATIVE_LAUNCHER="false"` to `AS_NATIVE_LAUNCHER="true"`.

When the server is in verbose mode, messages are logged to the console or terminal window in addition to the log file. In addition, pressing Ctrl-C stops the server and pressing Ctrl-\ (on UNIX platforms) or Ctrl-Break (on Windows platforms) prints a thread dump. On UNIX platforms, you can also print a thread dump using the `jstack` command (see <http://java.sun.com/javase/6/docs/technotes/tools/share/jstack.html>) or the command `kill -QUIT process_id`.

## Enterprise Server Logging

You can use the Enterprise Server's log files to help debug your applications. Use the Admin Console. In the developer profile, select the Enterprise Server component. In the cluster profile, select the Stand-Alone Instances component, and select the instance from the table. Then click the View Log Files button in the General Information page.

To change logging settings in the developer profile, select the Logging tab. In the cluster profile, select Logger Settings under the relevant configuration.

For details about logging, click the Help button in the Admin Console.

## Profiling Tools

You can use a profiler to perform remote profiling on the Enterprise Server to discover bottlenecks in server-side performance. This section describes how to configure these profilers for use with the Enterprise Server:

- “The NetBeans Profiler” on page 73
- “The HPROF Profiler” on page 73
- “The JProbe Profiler” on page 74



Information about comprehensive monitoring and management support in the Java™ 2 Platform, Standard Edition (J2SE™ platform) is available at <http://java.sun.com/javase/6/docs/technotes/guides/management/index.html>.

## The NetBeans Profiler

For information on how to use the NetBeans profiler, see <http://www.netbeans.org> and [http://blogs.sun.com/roller/page/bhavani?entry=analyzing\\_the\\_performance\\_of\\_java](http://blogs.sun.com/roller/page/bhavani?entry=analyzing_the_performance_of_java).

## The HPROF Profiler

The Heap and CPU Profiling Agent (HPROF) is a simple profiler agent shipped with the Java 2 SDK. It is a dynamically linked library that interacts with the Java Virtual Machine Profiler Interface (JVMPi) and writes out profiling information either to a file or to a socket in ASCII or binary format.

HPROF can monitor CPU usage, heap allocation statistics, and contention profiles. In addition, it can also report complete heap dumps and states of all the monitors and threads in the Java virtual machine. For more details on the HPROF profiler, see the technical article at <http://java.sun.com/developer/technicalArticles/Programming/HPROF.html>.

After HPROF is enabled using the following instructions, its libraries are loaded into the server process.

### ▼ To Use HPROF Profiling on UNIX

- 1 Use the Admin Console. In the developer profile, select the Enterprise Server component and the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. Then select the Profiler tab.**
- 2 Edit the following fields:**
  - Profiler Name – hprof
  - Profiler Enabled – true
  - Classpath – (leave blank)
  - Native Library Path – (leave blank)
  - JVM Option – Select Add, type the HPROF JVM option in the Value field, then check its box. The syntax of the HPROF JVM option is as follows:

```
-Xrunhprof[:helP] [[:param=value,param2=value2, ...]
```

Here is an example of *params* you can use:

```
-Xrunhprof:file=log.txt,thread=y,depth=3
```

The `file` parameter determines where the stack dump is written.

Using `help` lists parameters that can be passed to HPROF. The output is as follows:

```
Hprof usage: -Xrunhprof[:help][[:<option>=<value>, ...]
```

Option Name and Value	Description	Default
-----	-----	-----
heap=dump sites all	heap profiling	all
cpu=samples old	CPU usage	off
format=a b	ascii or binary output	a
file=<file>	write data to file (.txt for ascii)	java.hprof
net=<host>:<port>	send data over a socket	write to file
depth=<size>	stack trace depth	4
cutoff=<value>	output cutoff point	0.0001
lineno=y n	line number in traces?	y
thread=y n	thread in traces?	n
doe=y n	dump on exit?	y

---

**Note** – Do not use `help` in the JVM Option field. This parameter prints text to the standard output and then exits.

The help output refers to the parameters as options, but they are not the same thing as JVM options.

---

### 3 Restart the Enterprise Server.

This writes an HPROF stack dump to the file you specified using the `file` HPROF parameter.

## The JProbe Profiler

Information about JProbe™ from Sitraka is available at <http://www.quest.com/jprobe/>.

After JProbe is installed using the following instructions, its libraries are loaded into the server process.

### ▼ To Enable Remote Profiling With JProbe

#### 1 Install JProbe 3.0.1.1.

For details, see the JProbe documentation.

## 2 Configure Enterprise Server using the Admin Console:

a. In the developer profile, select the Enterprise Server component and the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. Then select the Profiler tab.

b. Edit the following fields before selecting Save and restarting the server:

- Profiler Name – `jprobe`
- Profiler Enabled – `true`
- Classpath – (leave blank)
- Native Library Path – `JProbe-dir/profiler`
- JVM Option – For each of these options, select Add, type the option in the Value field, then check its box
  - Xbootclasspath/p:`JProbe-dir/profiler/jpagent.jar`
  - Xrunjprobeagent
  - Xnoclassgc

---

**Note** – If any of the configuration options are missing or incorrect, the profiler might experience problems that affect the performance of the Enterprise Server.

---

When the server starts up with this configuration, you can attach the profiler.

## 3 Set the following environment variable:

`JPROBE_ARGS_0=-jp_input=JPL-file-path`

See [Step 6](#) for instructions on how to create the JPL file.

## 4 Start the server instance.

5 Launch the `jpprofiler` and attach to Remote Session. The default port is 4444.

6 Create the JPL file using the JProbe Launch Pad. Here are the required settings:

a. Select Server Side for the type of application.

b. On the Program tab, provide the following details:

- Target Server – *other-server*
- Server home Directory – *as-install*
- Server class File – `com.sun.enterprise.server.J2EERunner`
- Working Directory – *as-install*

- Classpath – *as-install/lib/appserv-rt.jar*
- Source File Path – *source-code-dir* (in case you want to get the line level details)
- Server class arguments – (optional)
- Main Package – *com.sun.enterprise.server*

You must also set VM, Attach, and Coverage tabs appropriately. For further details, see the JProbe documentation. After you have created the JPL file, use this as an input to `JPROBE_ARGS_0`.

PART II

Developing Applications and Application  
Components



# Securing Applications

---

This chapter describes how to write secure Java EE applications, which contain components that perform user authentication and access authorization for the business logic of Java EE components.

For information about administrative security for the Enterprise Server, see [Chapter 9](#), “Configuring Security,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

For general information about Java EE security, see “Chapter 29: Introduction to Security in Java EE” in the [Java EE 5 Tutorial](#) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

This chapter contains the following sections:

- “Security Goals” on page 80
- “Enterprise Server Specific Security Features” on page 80
- “Container Security” on page 81
- “Roles, Principals, and Principal to Role Mapping” on page 82
- “Realm Configuration” on page 84
- “JACC Support” on page 87
- “Pluggable Audit Module Support” on page 87
- “The `server.policy` File” on page 89
- “Configuring Message Security for Web Services” on page 92
- “Programmatic Login” on page 102
- “User Authentication for Single Sign-on” on page 104

## Security Goals

In an enterprise computing environment, there are many security risks. The goal of the Sun GlassFish Enterprise Server is to provide highly secure, interoperable, and distributed component computing based on the Java EE security model. Security goals include:

- Full compliance with the Java EE security model. This includes EJB and servlet role-based authorization.
- Support for single sign-on across all Enterprise Server applications within a single security domain.
- Support for web services message security.
- Security support for application clients.
- Support for several underlying authentication realms, such as simple file and Lightweight Directory Access Protocol (LDAP). Certificate authentication is also supported for Secure Socket Layer (SSL) client authentication. For Solaris, OS platform authentication is supported in addition to these.
- Support for declarative security through Enterprise Server specific XML-based role mapping.
- Support for Java Authorization Contract for Containers (JACC) pluggable authorization as included in the Java EE specification and defined by [Java Specification Request \(JSR\) 115](http://www.jcp.org/en/jsr/detail?id=115) (<http://www.jcp.org/en/jsr/detail?id=115>).
- Support for Java™ Authentication Service Provider Interface for Containers as included in the Java EE specification and defined by [JSR 196](http://www.jcp.org/en/jsr/detail?id=196) (<http://www.jcp.org/en/jsr/detail?id=196>).
- Support for Web Services Interoperability Technologies (WSIT) as described in [The WSIT Tutorial](https://wsit-docs.dev.java.net/releases/m5/) (<https://wsit-docs.dev.java.net/releases/m5/>).
- Support for P-asserted identity authentication as defined in [RFC \(Request for Comments\) 3325](http://www.ietf.org/rfc/rfc3325.txt) (<http://www.ietf.org/rfc/rfc3325.txt>).

## Enterprise Server Specific Security Features

The Enterprise Server supports the Java EE security model, as well as the following features which are specific to the Enterprise Server:

- Message security; see [“Configuring Message Security for Web Services” on page 92](#)
- Single sign-on across all Enterprise Server applications within a single security domain; see [“User Authentication for Single Sign-on” on page 104](#)
- Programmatic login; see [“Programmatic Login” on page 102](#)



# Container Security

The component containers are responsible for providing Java EE application security. The container provides two security forms:

- “[Declarative Security](#)” on page 81
- “[Programmatic Security](#)” on page 82

*Annotations* (also called metadata) enable a declarative style of programming, and so encompass both the declarative and programmatic security concepts. Users can specify information about security within a class file using annotations. When the application is deployed, this information can either be used by or overridden by the application or module deployment descriptor.

## Declarative Security

*Declarative security* means that the security mechanism for an application is declared and handled externally to the application. Deployment descriptors describe the Java EE application’s security structure, including security roles, access control, and authentication requirements.

The Enterprise Server supports the deployment descriptors specified by Java EE and has additional security elements included in its own deployment descriptors. Declarative security is the application deployer’s responsibility. For more information about Sun-specific deployment descriptors, see the *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

There are two levels of declarative security, as follows:

- “[Application Level Security](#)” on page 81
- “[Component Level Security](#)” on page 82

## Application Level Security

For an application, roles used by any application container must be defined in `@DeclareRoles` annotations in the code or `role-name` elements in the application deployment descriptor (`application.xml`). The role names are scoped to the EJB XML deployment descriptors (`ejb-jar.xml` and `sun-ejb-jar.xml` files) and to the servlet XML deployment descriptors (`web.xml` and `sun-web.xml` files). For an individually deployed web or EJB module, you define roles using `@DeclareRoles` annotations or `role-name` elements in the Java EE deployment descriptor files `web.xml` or `ejb-jar.xml`.

To map roles to principals and groups, define matching `security-role-mapping` elements in the `sun-application.xml`, `sun-ejb-jar.xml`, or `sun-web.xml` file for each `role-name` used by the application. For more information, see “[Roles, Principals, and Principal to Role Mapping](#)” on page 82.

## Component Level Security

Component level security encompasses web components and EJB components.

A secure web container authenticates users and authorizes access to a servlet or JSP by using the security policy laid out in the servlet XML deployment descriptors (`web.xml` and `sun-web.xml` files).

The EJB container is responsible for authorizing access to a bean method by using the security policy laid out in the EJB XML deployment descriptors (`ejb-jar.xml` and `sun-ejb-jar.xml` files).

## Programmatic Security

*Programmatic security* involves an EJB component or servlet using method calls to the security API, as specified by the Java EE security model, to make business logic decisions based on the caller or remote user's security role. Programmatic security should only be used when declarative security alone is insufficient to meet the application's security model.

The Java EE specification defines programmatic security as consisting of two methods of the EJB `EJBContext` interface and two methods of the servlet `HttpServletRequest` interface. The Enterprise Server supports these interfaces as specified in the specification.

For more information on programmatic security, see the following:

- The Java EE Specification
- [“Programmatic Login” on page 102](#)

## Roles, Principals, and Principal to Role Mapping

For applications, you define roles in `@DeclareRoles` annotations or the Java EE deployment descriptor file `application.xml`. You define the corresponding role mappings in the Enterprise Server deployment descriptor file `sun-application.xml`. For individually deployed web or EJB modules, you define roles in `@DeclareRoles` annotations or the Java EE deployment descriptor files `web.xml` or `ejb-jar.xml`. You define the corresponding role mappings in the Enterprise Server deployment descriptor files `sun-web.xml` or `sun-ejb-jar.xml`.

For more information regarding Java EE deployment descriptors, see the Java EE Specification. For more information regarding Enterprise Server deployment descriptors, see [Appendix A, “Deployment Descriptor Files,” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#).

Each `security-role-mapping` element in the `sun-application.xml`, `sun-web.xml`, or `sun-ejb-jar.xml` file maps a role name permitted by the application or module to principals and groups. For example, a `sun-web.xml` file for an individually deployed web module might contain the following:

```

<sun-web-app>
  <security-role-mapping>
    <role-name>manager</role-name>
    <principal-name>jgarcia</principal-name>
    <principal-name>mwebster</principal-name>
    <group-name>team-leads</group-name>
  </security-role-mapping>
  <security-role-mapping>
    <role-name>administrator</role-name>
    <principal-name>dsmith</principal-name>
  </security-role-mapping>
</sun-web-app>

```

A role can be mapped to either specific principals or to groups (or both). The principal or group names used must be valid principals or groups in the realm for the application or module. Note that the `role-name` in this example must match the `@DeclareRoles` annotations or the `role-name` in the `security-role` element of the corresponding `web.xml` file.

You can also specify a custom principal implementation class. This provides more flexibility in how principals can be assigned to roles. A user's JAAS login module now can authenticate its custom principal, and the authenticated custom principal can further participate in the Enterprise Server authorization process. For example:

```

<security-role-mapping>
  <role-name>administrator</role-name>
  <principal-name class-name="CustomPrincipalImplClass">
    dsmith
  </principal-name>
</security-role-mapping>

```

You can specify a default principal and a default principal to role mapping, each of which applies to the entire Enterprise Server instance. The default principal to role mapping maps group principals to named roles. Web modules that omit the `run-as` element in `web.xml` use the default principal. Applications and modules that omit the `security-role-mapping` element use the default principal to role mapping. These defaults are part of the Security Service, which you can access in the following ways:

- In the Admin Console, select the Security component under the relevant configuration. For details, click the Help button in the Admin Console.
- Use the `asadmin set` command. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#). For example, you can set the default principal as follows.

```

asadmin set --user adminuser server1.security-service.default-principal=dsmith
asadmin set --user adminuser server1.security-service.default-principal-password=secret

```

You can set the default principal to role mapping as follows.

```
asadmin set --user adminuser server1.security-service.activate-default-principal-to-role-mapping=true
asadmin set --user adminuser server1.security-service.mapped-principal-class=CustomPrincipalImplClass
```

## Realm Configuration

This section covers the following topics:

- [“Supported Realms” on page 84](#)
- [“How to Configure a Realm” on page 84](#)
- [“How to Set a Realm for an Application or Module” on page 85](#)
- [“Creating a Custom Realm” on page 85](#)

## Supported Realms

The following realms are supported in the Enterprise Server:

- `file` – Stores user information in a file. This is the default realm when you first install the Enterprise Server.
- `ldap` – Stores user information in an LDAP directory.
- `jdbc` – Stores user information in a database.

In the JDBC realm, the server gets user credentials from a database. The Application Server uses the database information and the enabled JDBC realm option in the configuration file.

For digest authentication, a JDBC realm should be created with `jdbcDigestRealm` as the JAAS context. The realm must be referenced in a `realm-name` element in the `web.xml` file as is standard practice.

- `certificate` – Sets up the user identity in the Enterprise Server security context, and populates it with user data obtained from cryptographically verified client certificates.
- `solaris` – Allows authentication using Solaris `username+password` data. This realm is only supported on the Solaris operating system, version 9 and above.

For information about configuring realms, see [“How to Configure a Realm” on page 84](#).

## How to Configure a Realm

You can configure a realm in one of these ways:

- In the Admin Console, open the Security component under the relevant configuration and go to the Realms page. For details, click the Help button in the Admin Console.
- Use the `asadmin create-auth-realm` command to configure realms on local servers. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## How to Set a Realm for an Application or Module

The following deployment descriptor elements have optional `realm` or `realm-name` data subelements or attributes that override the domain's default realm:

- `sun-application` element in `sun-application.xml`
- `login-config` element in `web.xml`
- `as-context` element in `sun-ejb-jar.xml`
- `client-container` element in `sun-acc.xml`
- `client-credential` element in `sun-acc.xml`

If modules within an application specify realms, these are ignored. If present, the realm defined in `sun-application.xml` is used, otherwise the domain's default realm is used.

For example, a realm is specified in `sun-application.xml` as follows:

```
<sun-application>
    ...
    <realm>ldap</realm>
</sun-application>
```

For more information about the deployment descriptor files and elements, see [Appendix A, “Deployment Descriptor Files,”](#) in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## Creating a Custom Realm

You can create a custom realm by providing a custom Java Authentication and Authorization Service (JAAS) login module class and a custom realm class. Note that client-side JAAS login modules are not suitable for use with the Enterprise Server.

JAAS is a set of APIs that enable services to authenticate and enforce access controls upon users. JAAS provides a pluggable and extensible framework for programmatic user authentication and authorization. JAAS is a core API and an underlying technology for Java EE security mechanisms. For more information about JAAS, refer to the JAAS specification for Java SDK, available at <http://java.sun.com/products/jaas/>.

For general information about realms and login modules, see “Chapter 29: Introduction to Security in Java EE” in the [Java EE 5 Tutorial](#) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

For Javadoc tool pages relevant to custom realms, go to <http://glassfish.dev.java.net/nonav/javaee5/api/index.html> and click on the `com.sun.appserv.security` package.

Custom login modules must extend the `com.sun.appserv.security.AppservPasswordLoginModule` class. This class implements `javax.security.auth.spi.LoginModule`. Custom login modules must not implement `LoginModule` directly.

Custom login modules must provide an implementation for one abstract method defined in `AppservPasswordLoginModule`:

```
abstract protected void authenticateUser() throws LoginException
```

This method performs the actual authentication. The custom login module must not implement any of the other methods, such as `login()`, `logout()`, `abort()`, `commit()`, or `initialize()`. Default implementations are provided in `AppservPasswordLoginModule` which hook into the Enterprise Server infrastructure.

The custom login module can access the following protected object fields, which it inherits from `AppservPasswordLoginModule`. These contain the user name and password of the user to be authenticated:

```
protected String _username;
protected String _password;
```

The `authenticateUser()` method must end with the following sequence:

```
String[] grpList;
// populate grpList with the set of groups to which
// _username belongs in this realm, if any
commitUserAuthentication(grpList);
```

Custom realms must extend the `com.sun.appserv.security.AppservRealm` class and implement the following methods:

```
public void init(Properties props) throws BadRealmException,
    NoSuchRealmException
```

This method is invoked during server startup when the realm is initially loaded. The `props` argument contains the properties defined for this realm in `domain.xml`. The realm can do any initialization it needs in this method. If the method returns without throwing an exception, the Enterprise Server assumes that the realm is ready to service authentication requests. If an exception is thrown, the realm is disabled.

```
public String getAuthType()
```

This method returns a descriptive string representing the type of authentication done by this realm.

```
public abstract Enumeration getGroupNames(String username) throws
    InvalidOperationException, NoSuchUserException
```

This method returns an Enumeration (of String objects) enumerating the groups (if any) to which the given username belongs in this realm.

## JACC Support

JACC (Java Authorization Contract for Containers) is part of the Java EE specification and defined by JSR 115 (<http://www.jcp.org/en/jsr/detail?id=115>). JACC defines an interface for pluggable authorization providers. Specifically, JACC is used to plug in the Java policy provider used by the container to perform Java EE caller access decisions. The Java policy provider performs Java policy decisions during application execution. This provides third parties with a mechanism to develop and plug in modules that are responsible for answering authorization decisions during Java EE application execution. The interfaces and rules used for developing JACC providers are defined in the JACC 1.0 specification.

The Enterprise Server provides a simple file-based JACC-compliant authorization engine as a default JACC provider. To configure an alternate provider using the Admin Console, open the Security component under the relevant configuration, and select the JACC Providers component. For details, click the Help button in the Admin Console.

## Pluggable Audit Module Support

Audit modules collect and store information on incoming requests (servlets, EJB components) and outgoing responses. You can create a custom audit module. This section covers the following topics:

- “Configuring an Audit Module” on page 87
- “The AuditModule Class” on page 88

For additional information about audit modules, see [Audit Callbacks](http://developers.sun.com/prodtech/appserver/reference/techart/ws_mgmt3.html#8.2) ([http://developers.sun.com/prodtech/appserver/reference/techart/ws\\_mgmt3.html#8.2](http://developers.sun.com/prodtech/appserver/reference/techart/ws_mgmt3.html#8.2)).

## Configuring an Audit Module

To configure an audit module, you can perform one of the following tasks:

- To specify an audit module using the Admin Console, open the Security component under the relevant configuration, and select the Audit Modules component. For details, click the Help button in the Admin Console.
- You can use the `asadmin create-audit-module` command to configure an audit module. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## The AuditModule Class

You can create a custom audit module by implementing a class that extends `com.sun.appserv.security.AuditModule`.

For Javadoc tool pages relevant to audit modules, go to <http://glassfish.dev.java.net/nonav/javaee5/api/index.html> and click on the `com.sun.appserv.security` package.

The `AuditModule` class provides default “no-op” implementations for each of the following methods, which your custom class can override.

```
public void init(Properties props)
```

The preceding method is invoked during server startup when the audit module is initially loaded. The `props` argument contains the properties defined for this module in `domain.xml`. The module can do any initialization it needs in this method. If the method returns without throwing an exception, the Enterprise Server assumes the module realm is ready to service audit requests. If an exception is thrown, the module is disabled.

```
public void authentication(String user, String realm, boolean success)
```

This method is invoked when an authentication request has been processed by a realm for the given user. The success flag indicates whether the authorization was granted or denied.

```
public void webInvocation(String user, HttpServletRequest req, String type, boolean success)
```

This method is invoked when a web container call has been processed by authorization. The success flag indicates whether the authorization was granted or denied. The `req` object is the standard `HttpServletRequest` object for this request. The `type` string is one of `hasUserDataPermission` or `hasResourcePermission` (see [JSR 115](http://www.jcp.org/en/jsr/detail?id=115) (<http://www.jcp.org/en/jsr/detail?id=115>)).

```
public void ejbInvocation(String user, String ejb, String method, boolean success)
```

This method is invoked when an EJB container call has been processed by authorization. The success flag indicates whether the authorization was granted or denied. The `ejb` and `method` strings describe the EJB component and its method that is being invoked.

```
public void webServiceInvocation(String uri, String endpoint, boolean success)
```

This method is invoked during validation of a web service request in which the endpoint is a servlet. The `uri` is the URL representation of the web service endpoint. The `endpoint` is the name of the endpoint representation. The success flag indicates whether the authorization was granted or denied.

```
public void ejbAsWebServiceInvocation(String endpoint, boolean success)
```



This method is invoked during validation of a web service request in which the endpoint is a stateless session bean. The endpoint is the name of the endpoint representation. The success flag indicates whether the authorization was granted or denied.

## The server.policy File

Each Enterprise Server domain has its own global J2SE policy file, located in *domain-dir/config*. The file is named `server.policy`.

The Enterprise Server is a Java EE compliant application server. As such, it follows the requirements of the Java EE specification, including the presence of the security manager (the Java component that enforces the policy) and a limited permission set for Java EE application code.

This section covers the following topics:

- [“Default Permissions” on page 89](#)
- [“Changing Permissions for an Application” on page 89](#)
- [“Enabling and Disabling the Security Manager” on page 92](#)

## Default Permissions

Internal server code is granted all permissions. These are covered by the `AllPermission` grant blocks to various parts of the server infrastructure code. Do not modify these entries.

Application permissions are granted in the default grant block. These permissions apply to all code not part of the internal server code listed previously. The Enterprise Server does not distinguish between EJB and web module permissions. All code is granted the minimal set of web component permissions (which is a superset of the EJB minimal set). Do not modify these entries.

A few permissions above the minimal set are also granted in the default `server.policy` file. These are necessary due to various internal dependencies of the server implementation. Java EE application developers must not rely on these additional permissions. In some cases, deleting these permissions might be appropriate. For example, one additional permission is granted specifically for using connectors. If connectors are not used in a particular domain, you should remove this permission, because it is not otherwise necessary.

## Changing Permissions for an Application

The default policy for each domain limits the permissions of Java EE deployed applications to the minimal set of permissions required for these applications to operate correctly. Do not add

extra permissions to the default set (the grant block with no codebase, which applies to all code). Instead, add a new grant block with a codebase specific to the applications requiring the extra permissions, and only add the minimally necessary permissions in that block.

If you develop multiple applications that require more than this default set of permissions, you can add the custom permissions that your applications need. The `com.sun.aas.instanceRoot` variable refers to the *domain-dir*. For example:

```
grant codeBase "file:${com.sun.aas.instanceRoot}/applications/j2ee-apps/-" {
  ...
}
```

You can add permissions to stub code with the following grant block:

```
grant codeBase "file:${com.sun.aas.instanceRoot}/generated/-" {
  ...
}
```

In general, you should add extra permissions only to the applications or modules that require them, not to all applications deployed to a domain. For example:

```
grant codeBase "file:${com.sun.aas.instanceRoot}/applications/j2ee-apps/MyApp/-" {
  ...
}
```

For a module:

```
grant codeBase "file:${com.sun.aas.instanceRoot}/applications/j2ee-modules/MyModule/-" {
  ...
}
```

An alternative way to add permissions to a specific application or module is to edit the `granted.policy` file for that application or module. The `granted.policy` file is located in the *domain-dir/generated/policy/app-or-module-name* directory. In this case, you add permissions to the default grant block. Do not delete permissions from this file.

When the application server policy subsystem determines that a permission should not be granted, it logs a `server.policy` message specifying the permission that was not granted and the protection domains, with indicated code source and principals that failed the protection check. For example, here is the first part of a typical message:

```
[#|2005-12-17T16:16:32.671-0200|INFO|sun-appserver-pe9.1|
javax.enterprise.system.core.security|_ThreadID=14;_ThreadName=Thread-31;|
JACC Policy Provider: PolicyWrapper.implies, context(null)-
permission((java.util.PropertyPermission java.security.manager write))
domain that failed(ProtectionDomain
(file:/E:/glassfish/domains/domain1/applications/j2ee-modules/cejug-clfds/ ... )
...
```

Granting the following permission eliminates the message:

```
grant codeBase "file:${com.sun.aas.instanceRoot}/applications/j2ee-modules/cejug-clfds/-" {  
    permission java.util.PropertyPermission "java.security.manager", "write";  
}
```

---

**Note** – Do not add `java.security.AllPermission` to the `server.policy` file for application code. Doing so completely defeats the purpose of the security manager, yet you still get the performance overhead associated with it.

---

As noted in the Java EE specification, an application should provide documentation of the additional permissions it needs. If an application requires extra permissions but does not document the set it needs, contact the application author for details.

As a last resort, you can iteratively determine the permission set an application needs by observing `AccessControlException` occurrences in the server log.

If this is not sufficient, you can add the `-Djava.security.debug=failure` JVM option to the domain. Use the following `asadmin create-jvm-options` command, then restart the server:

```
asadmin create-jvm-options --user adminuser -Djava.security.debug=failure
```

For more information about the `asadmin create-jvm-options` command, see the *Sun GlassFish Enterprise Server 2.1 Administration Reference*.

You can use the J2SE standard `policytool` or any text editor to edit the `server.policy` file. For more information, see

<http://java.sun.com/docs/books/tutorial/security1.2/tour2/index.html>.

For detailed information about policy file syntax, see <http://java.sun.com/javase/6/docs/technotes/guides/security/PolicyFiles.html#FileSyntax>.

For information about using system properties in the `server.policy` file, see

[http://java.sun.com/](http://java.sun.com/javase/6/docs/technotes/guides/security/PolicyFiles.html#PropertyExp)

[javase/6/docs/technotes/guides/security/PolicyFiles.html#PropertyExp](http://java.sun.com/javase/6/docs/technotes/guides/security/PolicyFiles.html#PropertyExp). For

information about Enterprise Server system properties, see “`system-property`” in *Sun GlassFish Enterprise Server 2.1 Administration Reference*.

For detailed information about the permissions you can set in the `server.policy` file, see

<http://java.sun.com/javase/6/docs/technotes/guides/security/permissions.html>.

The Javadoc for the `Permission` class is at

<http://java.sun.com/javase/6/docs/api/java/security/Permission.html>.

## Enabling and Disabling the Security Manager

The security manager is disabled in the developer and cluster profiles and enabled in the enterprise profile by default.

In a production environment, you may be able to safely disable the security manager if all of the following are true:

- Performance is critical
- Deployment to the production server is carefully controlled
- Only trusted applications are deployed
- Applications don't need policy enforcement

Disabling the security manager may improve performance significantly for some types of applications. To disable the security manager, do one of the following:

- To use the Admin Console, open the Security component under the relevant configuration, and uncheck the Security Manager Enabled box. Then restart the server. For details, click the Help button in the Admin Console.
- Use the following `asadmin delete-jvm-options` command, then restart the server:

```
asadmin delete-jvm-options --user adminuser -Djava.security.manager
```

To re-enable the security manager, use the corresponding `create-jvm-options` command. For more information about the `create-jvm-options` and `asadmin delete-jvm-options` commands, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## Configuring Message Security for Web Services

In *message security*, security information is applied at the message layer and travels along with the web services message. Web Services Security (WSS) is the use of XML Encryption and XML Digital Signatures to secure messages. WSS profiles the use of various security tokens including X.509 certificates, Security Assertion Markup Language (SAML) assertions, and username/password tokens to achieve this.

Message layer security differs from transport layer security in that it can be used to decouple message protection from message transport so that messages remain protected after transmission, regardless of how many hops they travel.

---

**Note** – In this release of the Enterprise Server, message layer annotations are not supported.

---

For more information about message security, see the following:

- The Java EE 5 Tutorial (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>) chapter titled “Chapter 29: Introduction to Security in Java EE”
- Chapter 10, “Configuring Message Security,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*
- JSR 196 (<http://www.jcp.org/en/jsr/detail?id=196>), Java Authentication Service Provider Interface for Containers
- The Liberty Alliance Project specifications at <http://www.projectliberty.org/resources/specifications.php>
- The Oasis Web Services Security (WSS) specification at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf>
- The Web Services Interoperability Organization (WS-I) Basic Security Profile (BSP) specification at <http://www.ws-i.org/Profiles/BasicSecurityProfile-1.0.html>
- The XML and Web Services Security page at <https://xwss.dev.java.net/>
- The WSIT page at <https://wsit.dev.java.net/>

The following web services security topics are discussed in this section:

- “Message Security Providers” on page 93
- “Message Security Responsibilities” on page 95
- “Application-Specific Message Protection” on page 96
- “Understanding and Running the Sample Application” on page 99

## Message Security Providers

When you first install the Enterprise Server, the providers `XWS_ClientProvider` and `XWS_ServerProvider` are configured but disabled. You can enable them in one of the following ways:

- To enable the message security providers using the Admin Console, open the Security component under the relevant configuration, select the Message Security component, and select SOAP. Then select `XWS_ServerProvider` from the Default Provider list and `XWS_ClientProvider` from the Default Client Provider list. For details, click the Help button in the Admin Console.
- You can enable the message security providers using the following commands.

```
asadmin set --user adminuser
server-config.security-service.message-security-config.SOAP.default_provider=XWS_ServerProvider
asadmin set --user adminuser
server-config.security-service.message-security-config.SOAP.default_client_provider=XWS_ClientProvider
```

For more information about the `asadmin set` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

The example described in “[Understanding and Running the Sample Application](#)” on page 99 uses the `ClientProvider` and `ServerProvider` providers, which are enabled when the `asant` targets are run. You don’t need to enable these on the Enterprise Server prior to running the example.

If you install the Access Manager, you have these additional provider choices:

- `AMClientProvider` and `AMServerProvider` – These providers secure web services and Simple Object Access Protocol (SOAP) messages using either WS-I BSP or Liberty ID-WSF tokens. These providers are used automatically if they are configured as the default providers. If you wish to override any provider settings, you can configure these providers in `message-security-binding` elements in the `sun-web.xml`, `sun-ejb-jar.xml`, and `sun-application-client.xml` deployment descriptor files.
- `AMHttpProvider` – This provider handles the initial end user authentication for securing web services using Liberty ID-WSF tokens and redirects requests to the Access Manager for single sign-on. To use this provider, specify it in the `httpServlet-security-provider` attribute of the `sun-web-app` element in the `sun-web.xml` file.

Liberty specifications can be viewed at

<http://www.projectliberty.org/resources/specifications.php>. The WS-I BSP specification can be viewed at

<http://www.ws-i.org/Profiles/BasicSecurityProfile-1.0.html>.

For more information about the Sun-specific deployment descriptor files, see the *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

For information about configuring these providers in the Enterprise Server, see Chapter 10, “[Configuring Message Security](#),” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*. For additional information about overriding provider settings, see “[Application-Specific Message Protection](#)” on page 96.

You can create new message security providers in one of the following ways:

- To create a message security provider using the Admin Console, open the Security component under the relevant configuration, and select the Message Security component. For details, click the Help button in the Admin Console.
- You can use the `asadmin create-message-security-provider` command to create a message security provider. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

In addition, you can set a few optional provider properties. For more information, see the property descriptions under “[provider-config](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Reference*.

## Message Security Responsibilities

In the Enterprise Server, the system administrator and application deployer roles are expected to take primary responsibility for configuring message security. In some situations, the application developer may also contribute, although in the typical case either of the other roles may secure an existing application without changing its implementation and without involving the developer. The responsibilities of the various roles are defined in the following sections:

- [“Application Developer” on page 95](#)
- [“Application Deployer” on page 95](#)
- [“System Administrator” on page 96](#)

### Application Developer

The application developer can turn on message security, but is not responsible for doing so. Message security can be set up by the system administrator so that all web services are secured, or set up by the application deployer when the provider or protection policy bound to the application must be different from that bound to the container.

The application developer is responsible for the following:

- Determining if an application-specific message protection policy is required by the application. If so, ensuring that the required policy is specified at application assembly which may be accomplished by communicating with the application deployer.
- Determining if message security is necessary at the Enterprise Server level. If so, ensuring that this need is communicated to the system administrator, or taking care of implementing message security at the Enterprise Server level.

### Application Deployer

The application deployer is responsible for the following:

- Specifying (at application assembly) any required application-specific message protection policies if such policies have not already been specified by upstream roles (the developer or assembler)
- Modifying Sun-specific deployment descriptors to specify application-specific message protection policies information (message-security-binding elements) to web service endpoint and service references

These security tasks are discussed in [“Application-Specific Message Protection” on page 96](#). A sample application using message security is discussed in [“Understanding and Running the Sample Application” on page 99](#).

## System Administrator

The system administrator is responsible for the following:

- Configuring message security providers on the Enterprise Server.
- Managing user databases.
- Managing keystore and truststore files.
- Installing the sample. This is only done if the `xms` sample application is used to demonstrate the use of message layer web services security.

A system administrator uses the Admin Console to manage server security settings and uses a command line tool to manage certificate databases. Certificates and private keys are stored in key stores and are managed with `keytool`. If Network Security Services (NSS) is installed and you have selected the enterprise profile, certificates and private keys are stored in an NSS database, where they are managed using `certutil`. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*. System administrator tasks are discussed in Chapter 10, “Configuring Message Security,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

## Application-Specific Message Protection

When the Enterprise Server provided configuration is insufficient for your security needs, and you want to override the default protection, you can apply *application-specific message security* to a web service.

Application-specific security is implemented by adding the message security binding to the web service endpoint, whether it is an EJB or servlet web service endpoint. Modify Sun-specific XML files to add the message binding information.

Message security can also be specified using a WSIT security policy in the WSDL file. For details, see the WSIT page at <https://wsit.dev.java.net/>.

For more information about message security providers, see “Message Security Providers” on page 93.

For more details on message security binding for EJB web services, servlet web services, and clients, see the XML file descriptions in Appendix A, “Deployment Descriptor Files,” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

- For `sun-ejb-jar.xml`, see “The `sun-ejb-jar.xml` File” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.
- For `sun-web.xml`, see “The `sun-web.xml` File” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.
- For `sun-application-client.xml`, see “The `sun-application-client.xml` file” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.



This section contains the following topics:

- [“Using a Signature to Enable Message Protection for All Methods” on page 97](#)
- [“Configuring Message Protection for a Specific Method Based on Digital Signatures” on page 98](#)

## Using a Signature to Enable Message Protection for All Methods

To enable message protection for all methods using digital signature, update the `message-security-binding` element for the EJB web service endpoint in the application’s `sun-ejb-jar.xml` file. In this file, add `request-protection` and `response-protection` elements, which are analogous to the `request-policy` and `response-policy` elements discussed in [Chapter 10, “Configuring Message Security,” in \*Sun GlassFish Enterprise Server 2.1 Administration Guide\*](#). To apply the same protection mechanisms for all methods, leave the `method-name` element blank. [“Configuring Message Protection for a Specific Method Based on Digital Signatures” on page 98](#) discusses listing specific methods or using wildcard characters.

This section uses the sample application discussed in [“Understanding and Running the Sample Application” on page 99](#) to apply application-level message security to show only the differences necessary for protecting web services using various mechanisms.

### ▼ To Enable Message Protection for All Methods Using Digital Signature

- 1 **In a text editor, open the application’s `sun-ejb-jar.xml` file.**

For the `xms` example, this file is located in the directory `app-dir/xms-ejb/src/conf`, where `app-dir` is defined in [“To Set Up the Sample Application” on page 100](#).

- 2 **Modify the `sun-ejb-jar.xml` file by adding the `message-security-binding` element as shown:**

```
<sun-ejb-jar>
  <enterprise-beans>
    <unique-id>1</unique-id>
    <ejb>
      <ejb-name>HelloWorld</ejb-name>
      <jndi-name>HelloWorld</jndi-name>
      <webservice-endpoint>
        <port-component-name>HelloIF</port-component-name>
        <endpoint-address-uri>service/HelloWorld</endpoint-address-uri>
        <message-security-binding auth-layer="SOAP">
          <message-security>
            <request-protection auth-source="content" />
            <response-protection auth-source="content"/>
          </message-security>
        </message-security-binding>
      </webservice-endpoint>
    </ejb>
```

```

    </enterprise-beans>
</sun-ejb-jar>

```

- 3 **Compile, deploy, and run the application as described in “To Run the Sample Application” on page 100.**

## Configuring Message Protection for a Specific Method Based on Digital Signatures

To enable message protection for a specific method, or for a set of methods that can be identified using a wildcard value, follow these steps. As in the example discussed in “Using a Signature to Enable Message Protection for All Methods” on page 97, to enable message protection for a specific method, update the message-security-binding element for the EJB web service endpoint in the application’s sun-ejb-jar.xml file. To this file, add request-protection and response-protection elements, which are analogous to the request-policy and response-policy elements discussed in Chapter 10, “Configuring Message Security,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*. The administration guide includes a table listing the set and order of security operations for different request and response policy configurations.

This section uses the sample application discussed in “Understanding and Running the Sample Application” on page 99 to apply application-level message security to show only the differences necessary for protecting web services using various mechanisms.

### ▼ To Enable Message Protection for a Particular Method or Set of Methods Using Digital Signature

- 1 **In a text editor, open the application’s sun-ejb-jar.xml file.**

For the xms example, this file is located in the directory *app-dir/xms-ejb/src/conf*, where *app-dir* is defined in “To Set Up the Sample Application” on page 100.

- 2 **Modify the sun-ejb-jar.xml file by adding the message-security-binding element as shown:**

```

<sun-ejb-jar>
  <enterprise-beans>
    <unique-id>1</unique-id>
    <ejb>
      <ejb-name>HelloWorld</ejb-name>
      <jndi-name>HelloWorld</jndi-name>
      <webservice-endpoint>
        <port-component-name>HelloIF</port-component-name>
        <endpoint-address-uri>service/HelloWorld</endpoint-address-uri>
        <message-security-binding auth-layer="SOAP">
          <message-security>

```

```

    <message>
      <java-method>
        <method-name>ejbCreate</method-name>
      </java-method>
    </message>
  <message>
    <java-method>
      <method-name>sayHello</method-name>
    </java-method>
  </message>
  <request-protection auth-source="content" />
  <response-protection auth-source="content"/>
</message-security>
</message-security-binding>
</webservice-endpoint>
</ejb>
</enterprise-beans>
</sun-ejb-jar>

```

- 3 **Compile, deploy, and run the application as described in “To Run the Sample Application” on page 100.**

## Understanding and Running the Sample Application

This section discusses the WSS sample application. This sample application is installed on your system only if you installed the J2EE 1.4 samples. If you have not installed these samples, see [“To Set Up the Sample Application” on page 100](#).

The objective of this sample application is to demonstrate how a web service can be secured with WSS. The web service in the `xms` example is a simple web service implemented using a Java EE EJB endpoint and a web service endpoint implemented using a servlet. In this example, a service endpoint interface is defined with one operation, `sayHello`, which takes a string then sends a response with `Hello` prefixed to the given string. You can view the WSDL file for the service endpoint interface at `app-dir/xms-ejb/src/conf/HelloWorld.wsdl`, where `app-dir` is defined in [“To Set Up the Sample Application” on page 100](#).

In this application, the client looks up the service using the JNDI name `java:comp/env/service/HelloWorld` and gets the port information using a static stub to invoke the operation using a given name. For the name `Duke`, the client gets the response `Hello Duke!`

This example shows how to use message security for web services at the Enterprise Server level. For information about using message security at the application level, see [“Application-Specific Message Protection” on page 96](#). The WSS message security mechanisms implement message-level authentication (for example, XML digital signature and encryption) of SOAP web services invocations using the X.509 and username/password profiles of the OASIS

WS-Security standard, which can be viewed from the following URL:  
<http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf>.

This section includes the following topics:

- “To Set Up the Sample Application” on page 100
- “To Run the Sample Application” on page 100

## ▼ To Set Up the Sample Application

**Before You Begin** To have access to this sample application, you must have previously installed the J2EE 1.4 samples. If the samples are not installed, follow the steps in the following section.

After you follow these steps, the sample application is located in the directory *as-install/j2ee14-samples/samples/webservices/security/ejb/apps/xms/* or in a directory of your choice. For easy reference throughout the rest of this section, this directory is referred to as simply *app-dir*.

- 1 **Go to the J2EE 1.4 download URL** (<http://java.sun.com/j2ee/1.4/download.html>) **in your browser.**
- 2 **Click on the Download button for the Samples Bundle.**
- 3 **Click on Accept License Agreement.**
- 4 **Click on the J2EE SDK Samples link.**
- 5 **Choose a location for the `j2eesdk-1_4_03-samples.zip` file.**  
Saving the file to *as-install* is recommended.

- 6 **Unzip the file.**

Unzipping to the *as-install/j2ee14-samples* directory is recommended. For example, you can use the following command.

```
unzip j2eesdk-1_4_03-samples.zip -d j2ee14-samples
```

## ▼ To Run the Sample Application

- 1 **Make sure that the Enterprise Server is running.**

Message security providers are set up when the *asant* targets are run, so you do not need to configure these on the Enterprise Server prior to running this example.

- 2 If you are not running HTTP on the default port of 8080, change the WSDL file for the example to reflect the change, and change the `common.properties` file to reflect the change as well.**

The WSDL file for this example is located at `app-dir/xms-ejb/src/conf/HelloWorld.wsdl`.

The port number is in the following section:

```
<service name="HelloWorld">
  <port name="HelloIFPort" binding="tns:HelloIFBinding">
    <soap:address location="http://localhost:8080/service/HelloWorld"/>
  </port>
</service>
```

Verify that the properties in the `as-install/samples/common.properties` file are set properly for your installation and environment. If you need a more detailed description of this file, refer to the “Configuration” section for the web services security applications at `as-install/j2ee14-samples/samples/webservices/security/docs/common.html#Logging`.

- 3 Change to the `app-dir` directory.**
- 4 Run the following `asant` targets to compile, deploy, and run the example application:**

- a. To compile samples:**

```
asant
```

- b. To deploy samples:**

```
asant deploy
```

- c. To run samples:**

```
asant run
```

If the sample has compiled and deployed properly, you see the following response on your screen after the application has run:

```
run:[echo] Running the xms program:[exec] Established message level security :
Hello Duke!
```

- 5 To undeploy the sample, run the following `asant` target:**

```
asant undeploy
```

All of the web services security examples use the same web service name (`HelloWorld`) and web service ports. These examples show only the differences necessary for protecting web services using various mechanisms. Make sure to undeploy an application when you have completed running it. If you do not, you receive an `Already in Use` error and deployment failures when you try to deploy another web services example application.

# Programmatic Login

Programmatic login allows a deployed Java EE application or module to invoke a login method. If the login is successful, a `SecurityContext` is established as if the client had authenticated using any of the conventional Java EE mechanisms. Programmatic login is supported for servlet and EJB components on the server side, and for stand-alone or application clients on the client side. Programmatic login is useful for an application having special needs that cannot be accommodated by any of the Java EE standard authentication mechanisms.

---

**Note** – Programmatic login is specific to the Enterprise Server and not portable to other application servers.

---

This section contains the following topics:

- [“Programmatic Login Precautions” on page 102](#)
- [“Granting Programmatic Login Permission” on page 103](#)
- [“The ProgrammaticLogin Class” on page 103](#)

## Programmatic Login Precautions

The Enterprise Server is not involved in how the login information (user, password) is obtained by the deployed application. Programmatic login places the burden on the application developer with respect to assuring that the resulting system meets security requirements. If the application code reads the authentication information across the network, the application determines whether to trust the user.

Programmatic login allows the application developer to bypass the application server-supported authentication mechanisms and feed authentication data directly to the security service. While flexible, this capability should not be used without some understanding of security issues.

Since this mechanism bypasses the container-managed authentication process and sequence, the application developer must be very careful in making sure that authentication is established before accessing any restricted resources or methods. It is also the application developer’s responsibility to verify the status of the login attempt and to alter the behavior of the application accordingly.

The programmatic login state does not necessarily persist in sessions or participate in single sign-on.

Lazy authentication is not supported for programmatic login. If an access check is reached and the deployed application has not properly authenticated using the programmatic login method, access is denied immediately and the application might fail if not coded to account for this occurrence. One way to account for this occurrence is to catch the access control or security exception, perform a programmatic login, and repeat the request.

## Granting Programmatic Login Permission

The `ProgrammaticLoginPermission` permission is required to invoke the programmatic login mechanism for an application if the security manager is enabled. For information about the security manager, see “[The `server.policy` File](#)” on page 89. This permission is not granted by default to deployed applications because this is not a standard Java EE mechanism.

To grant the required permission to the application, add the following to the `domain-dir/config/server.policy` file:

```
grant codeBase "file:jar-file-path" {
    permission com.sun.appserv.security.ProgrammaticLoginPermission
        "login";
};
```

The *jar-file-path* is the path to the application's JAR file.

## The ProgrammaticLogin Class

The `com.sun.appserv.security.ProgrammaticLogin` class enables a user to perform login programmatically.

For Javadoc tool pages relevant to programmatic login, go to <http://glassfish.dev.java.net/nonav/javaee5/api/index.html> and click on the `com.sun.appserv.security` package.

The `ProgrammaticLogin` class has four login methods, two for servlets or JSP files and two for EJB components.

The login methods for servlets or JSP files have the following signatures:

```
public java.lang.Boolean login(String user, String password,
    javax.servlet.http.HttpServletRequest request,
    javax.servlet.http.HttpServletResponse response)

public java.lang.Boolean login(String user, String password,
    String realm, javax.servlet.http.HttpServletRequest request,
    javax.servlet.http.HttpServletResponse response, boolean errors)
    throws java.lang.Exception
```

The login methods for EJB components have the following signatures:

```
public java.lang.Boolean login(String user, String password)

public java.lang.Boolean login(String user, String password,
    String realm, boolean errors) throws java.lang.Exception
```

All of these login methods accomplish the following:

- Perform the authentication
- Return `true` if login succeeded, `false` if login failed

The login occurs on the realm specified unless it is null, in which case the domain's default realm is used. The methods with no `realm` parameter use the domain's default realm.

If the `errors` flag is set to `true`, any exceptions encountered during the login are propagated to the caller. If set to `false`, exceptions are thrown.

On the client side, `realm` and `errors` parameters are ignored and the actual login does not occur until a resource requiring a login is accessed. A `java.rmi.AccessException` with `COBRA_NO_PERMISSION` occurs if the actual login fails.

The logout methods for servlets or JSP files have the following signatures:

```
public java.lang.Boolean logout(HttpServletRequest request,
                               HttpServletResponse response)
```

```
public java.lang.Boolean logout(HttpServletRequest request,
                               HttpServletResponse response, boolean errors)
    throws java.lang.Exception
```

The logout methods for EJB components have the following signatures:

```
public java.lang.Boolean logout()
```

```
public java.lang.Boolean logout(boolean errors)
    throws java.lang.Exception
```

All of these logout methods return `true` if logout succeeded, `false` if logout failed.

If the `errors` flag is set to `true`, any exceptions encountered during the logout are propagated to the caller. If set to `false`, exceptions are thrown.

## User Authentication for Single Sign-on

The single sign-on feature of the Enterprise Server allows multiple web applications deployed to the same virtual server to share the user authentication state. With single sign-on enabled, users who log in to one web application become implicitly logged into other web applications on the same virtual server that require the same authentication information. Otherwise, users would have to log in separately to each web application whose protected resources they tried to access.

A sample application using the single sign-on scenario could be a consolidated airline booking service that searches all airlines and provides links to different airline web sites. After the user signs on to the consolidated booking service, the user information can be used by each individual airline site without requiring another sign-on.



Single sign-on operates according to the following rules:

- Single sign-on applies to web applications configured for the same realm and virtual server. The realm is defined by the `realm-name` element in the `web.xml` file. For information about virtual servers, see [Chapter 13, “Configuring the HTTP Service,”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.
- As long as users access only unprotected resources in any of the web applications on a virtual server, they are not challenged to authenticate themselves.
- As soon as a user accesses a protected resource in any web application associated with a virtual server, the user is challenged to authenticate himself or herself, using the login method defined for the web application currently being accessed.
- After authentication, the roles associated with this user are used for access control decisions across all associated web applications, without challenging the user to authenticate to each application individually.
- When the user logs out of one web application (for example, by invalidating the corresponding session), the user’s sessions in all web applications are invalidated. Any subsequent attempt to access a protected resource in any application requires the user to authenticate again.

The single sign-on feature utilizes HTTP cookies to transmit a token that associates each request with the saved user identity, so it can only be used in client environments that support cookies.

To configure single sign-on, set the following properties in the `virtual-server` element of the `domain.xml` file:

- `sso-enabled` - If `false`, single sign-on is disabled for this virtual server, and users must authenticate separately to every application on the virtual server. The default is `true`.
- `sso-max-inactive-seconds` - Specifies the time after which a user’s single sign-on record becomes eligible for purging if no client activity is received. Since single sign-on applies across several applications on the same virtual server, access to any of the applications keeps the single sign-on record active. The default value is 5 minutes (300 seconds). Higher values provide longer single sign-on persistence for the users at the expense of more memory use on the server.
- `sso-reap-interval-seconds` - Specifies the interval between purges of expired single sign-on records. The default value is 60.

Here is an example configuration with all default values:

```
<virtual-server id="server" ... >
  ...
  <property name="sso-enabled" value="true"/>
  <property name="sso-max-inactive-seconds" value="300"/>
  <property name="sso-reap-interval-seconds" value="60"/>
</virtual-server>
```



# Developing Web Services

---

This chapter describes Enterprise Server support for web services. Java™ API for XML-Based Web Services (JAX-WS) version 2.0 is supported. Java API for XML-Based Remote Procedure Calls (JAX-RPC) version 1.1 is supported for backward compatibility. This chapter contains the following sections:

- “Creating Portable Web Service Artifacts” on page 108
- “Deploying a Web Service” on page 108
- “Web Services Registry” on page 109
- “The Web Service URI, WSDL File, and Test Page” on page 110
- “JBI Runtime” on page 111
- “Using the Woodstox Parser” on page 113

“Part Two: Web Services” in the [Java EE 5 Tutorial](http://java.sun.com/javaee/5/docs/tutorial/doc/index.html) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>) shows how to deploy simple web services to the Enterprise Server. “Chapter 20: Java API for XML Registries” explains how to set up a registry and create clients that access the registry.

For additional information about JAX-WS and web services, see [Java Specification Request \(JSR\) 224](http://jcp.org/aboutJava/communityprocess/pfd/jsr224/index.html) (<http://jcp.org/aboutJava/communityprocess/pfd/jsr224/index.html>) and [JSR 109](http://jcp.org/en/jsr/detail?id=109) (<http://jcp.org/en/jsr/detail?id=109>).

For information about web services security, see “[Configuring Message Security for Web Services](#)” on page 92.

For information about web services administration, monitoring, logging, and registries, see Chapter 14, “[Managing Web Services](#),” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

The Fast Infoset standard specifies a binary format based on the XML Information Set. This format is an efficient alternative to XML. For information about using Fast Infoset, see the following links:

- [Java Web Services Developer Pack 1.6 Release Notes](http://java.sun.com/webservices/docs/1.6/ReleaseNotes.html) (<http://java.sun.com/webservices/docs/1.6/ReleaseNotes.html>)
- [Fast Infoset in Java Web Services Developer Pack, Version 1.6](http://java.sun.com/webservices/docs/1.6/jaxrpc/fastinfoset/manual.html) (<http://java.sun.com/webservices/docs/1.6/jaxrpc/fastinfoset/manual.html>)
- [Fast Infoset Project](http://fi.dev.java.net) (<http://fi.dev.java.net>)

## Creating Portable Web Service Artifacts

For a tutorial that shows how to use the `wsimport` and `wsgen` commands, see “Part Two: Web Services” in the [Java EE 5 Tutorial](http://java.sun.com/javaee/5/docs/tutorial/doc/index.html) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>). For reference information on these commands, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Deploying a Web Service

You deploy a web service endpoint to the Enterprise Server just as you would any servlet, stateless session bean (SLSB), or application. After you deploy the web service, the next step is to publish it. For more information about publishing a web service, see “Web Services Registry” on page 109.

You can use the autodeployment feature to deploy a simple [JSR 181](http://jcp.org/en/jsr/detail?id=181) (<http://jcp.org/en/jsr/detail?id=181>) annotated file. You can compile and deploy in one step, as in the following example:

```
javac -cp javaee.jar -d domain-dir/autodeploy MyWSDemo.java
```

---

**Note** – For complex services with dependent classes, user specified WSDL files, or other advanced features, autodeployment of an annotated file is not sufficient.

---

The Sun-specific deployment descriptor files `sun-web.xml` and `sun-ejb-jar.xml` provide optional web service enhancements in their `webservice-endpoint` and `webservice-description` elements, including a `debugging-enabled` subelement that enables the creation of a test page. The test page feature is enabled by default and described in “[The Web Service URI, WSDL File, and Test Page](#)” on page 110.

For more information about deployment, autodeployment, and deployment descriptors, see the *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*. For more information about the `asadmin deploy` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

# Web Services Registry

You deploy a registry to the Enterprise Server just as you would any connector module, except that if you are using the Admin Console, you must select a Registry Type value. After deployment, you can configure a registry in one of the following ways:

- In the Admin Console, open the Web Services component, and select the Registry tab. For details, click the Help button in the Admin Console.
- To configure a registry using the command line, use the following commands.

- Set the registry type to `com.sun.appserv.registry.ebxml` or `com.sun.appserv.registry.uddi`. Use a backslash before each period as an escape character. For example:

```
asadmin create-resource-adapter-config --user adminuser
--property com.sun.appserv.registry.ebxml=true MyReg
```

- Set any properties needed by the registry. For an ebXML registry, set the `LifeCycleManagerURL` and `QueryManagerURL` properties. In the following example, the system property `REG_URL` is set to `http://siroe.com:6789/soar/registry/soap`.

```
asadmin create-connector-connection-pool --user adminuser --raname MyReg
--connectiondefinition javax.xml.registry.ConnectionFactory --property
LifeCycleManagerURL=${REG_URL}:QueryManagerURL=${REG_URL} MyRegCP
```

- Set a JNDI name for the registry resource. For example:

```
asadmin create-connector-resource --user adminuser --poolname MyRegCP jndi-MyReg
```

For details on these commands, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

After you deploy a web service, you can publish it to a registry in one of the following ways:

- In the Admin Console, open the Web Services component, select the web service in the listing on the General tab, and select the Publish tab. For details, click the Help button in the Admin Console.
- Use the `asadmin publish-to-registry` command. For example:

```
asadmin publish-to-registry --user adminuser --registryjndinames jndi-MyReg --webservicename my-ws#simple
```

For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

The Sun Java Enterprise System (Java ES) includes a Sun-specific ebXML registry. For more information about the Java ES registry and registries in general, see “Chapter 20: Java API for XML Registries” in the [Java EE 5 Tutorial](#) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

A connector module that accesses UDDI registries is provided with the Enterprise Server in the *as-install/lib/install/applications/jaxr-ra* directory.

You can also use the JWSDP registry available at <http://java.sun.com/webservices/jwsdp/index.jsp> or the SOA registry available at <http://www.sun.com/products/soa/index.jsp>.

## The Web Service URI, WSDL File, and Test Page

Clients can run a deployed web service by accessing its service endpoint address URI, which has the following format:

```
http://host:port/context-root/servlet-mapping-url-pattern
```

The *context-root* is defined in the *application.xml* or *web.xml* file, and can be overridden in the *sun-application.xml* or *sun-web.xml* file. The *servlet-mapping-url-pattern* is defined in the *web.xml* file.

In the following example, the *context-root* is *my-ws* and the *servlet-mapping-url-pattern* is */simple*:

```
http://localhost:8080/my-ws/simple
```

You can view the WSDL file of the deployed service in a browser by adding *?WSDL* to the end of the URI. For example:

```
http://localhost:8080/my-ws/simple?WSDL
```

For debugging, you can run a test page for the deployed service in a browser by adding *?Tester* to the end of the URL. For example:

```
http://localhost:8080/my-ws/simple?Tester
```

You can also test a service using the Admin Console. Open the Web Services component, select the web service in the listing on the General tab, and select Test. For details, click the Help button in the Admin Console.

---

**Note** – The test page works only for WS-I compliant web services. This means that the tester servlet does not work for services with WSDL files that use RPC/encoded binding.

---

Generation of the test page is enabled by default. You can disable the test page for a web service by setting the value of the *debugging-enabled* element in the *sun-web.xml* and *sun-ejb-jar.xml* deployment descriptor to *false*. For more information, see the [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

# JBI Runtime

The Java Business Integration runtime (JBI runtime) provides a distributed infrastructure used for enterprise integration. It consists of a set of binding components and service engines, which integrate various types of information technology assets. The binding components and service engines are interconnected with a normalized message router. Binding components and service engines adapt information technology assets to a standard services model, based on XML message exchange using standardized message exchange patterns. The JBI runtime provides services for transforming and routing messages, as well as the ability to centrally administer the distributed system.

This JBI runtime incorporates the [JSR 208](http://jcp.org/en/jsr/detail?id=208) (<http://jcp.org/en/jsr/detail?id=208>) specification for JBI and other open standards. The JBI runtime allows you to integrate web services and enterprise applications as loosely coupled composite applications within a Service-Oriented Architecture (SOA).

The distribution of the JBI runtime includes a Java EE service engine, an HTTP SOAP binding component, a WSDL shared library, and Ant tasks described in “[JBI Tasks](#)” on page 68. For information about JBI administration in the Enterprise Server, see the [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

Additional components, tools, and documentation are available for download. Refer to [Project Open ESB](#) (<https://open-esb.dev.java.net/>) for more information on the additional components, tools, and documentation that are available.

The Java EE Service Engine acts as a bridge between the Java EE and JBI runtime environments for web service providers and web service consumers. The Java EE Service Engine provides better performance than a SOAP over HTTP binding component due to in-process communication between components and additional protocols provided by JBI binding components such as JMS, SMTP, and File.

The JSR 208 specification allows transactions to be propagated to other components using a message exchange property specified in the `JTA_TRANSACTION_PROPERTY_NAME` field. The Java EE Service Engine uses this property to set and get a transaction object from the JBI message exchange. It then uses the transaction object to take part in a transaction. This means a Java EE application or module can take part in a transaction started by a JBI application. Conversely, a JBI application can take part in a transaction started by a Java EE application or module.

Similarly, the JSR 208 specification allows a security subject to be propagated as a message exchange property named `javax.jbi.security.subject`. Thus a security subject can be propagated from a Java EE application or module to a JBI application or the reverse.

To deploy a Java EE application or module as a JBI service unit, use the Admin Console or the `asadmin deploy-jbi-service-assembly` command. For more information about the `asadmin deploy-jbi-service-assembly` command, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## Using the `jbi.xml` File

Section 6.3.1 of the JSR 208 specification describes the `jbi.xml` file. This is a deployment descriptor, located in the `META-INF` directory. To deploy a Java EE application or module as a JBI service unit, you need only specify a small subset of elements in the `jbi.xml` file. Here is an example provider:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<jbi version="1.0" xmlns="http://java.sun.com/xml/ns/jbi" xmlns:ns0="http://ejbws.jbi.misc/">
  <services binding-component="false">
    <provides endpoint-name="MiscPort" interface-name="ns0:Misc" service-name="ns0:MiscService"/>
  </services>
</jbi>
```

Here is an example consumer:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<jbi version="1.0" xmlns="http://java.sun.com/xml/ns/jbi" xmlns:ns0="http://message.hello.jbi/">
  <services binding-component="false">
    <consumes endpoint-name="MsgPort" interface-name="ns0:Msg" service-name="ns0:MsgService"/>
  </services>
</jbi>
```

The Java EE Service Engine enables the endpoints described in the `provides` section of the `jbi.xml` file in the JBI runtime. Similarly, the Java EE Service Engine routes invocations of the endpoints described in the `consumes` section from the Java EE web service consumer to the JBI runtime.

## Using Application Server Descriptors

To determine whether a web service endpoint is enabled in the JBI runtime environment, you can set a `jbi-enabled` attribute in the Enterprise Server. This attribute is set to `false` (disabled) by default. To enable an endpoint for JBI, set the attribute to `true` using the `asadmin set` command. For example, if an endpoint is bundled as a WAR file named `my-ws.war` with an endpoint named `simple`, use the following command:

```
asadmin set --user adminuser server.applications.web-module.my-ws.web-service-endpoint.simple.jbi-enabled=true
```

Determining whether requests from a web service consumer are routed through the Java EE Service Engine is unnecessary and deprecated, but supported for backward compatibility. You can set a `stub-property` named `jbi-enabled` in the consumer's `sun-web.xml` or `sun-ejb-jar.xml` file. This property is set to `true` (enabled) by default. Here is an example of the `sun-web.xml` file:



```

<sun-web-app>
  <service-ref>
    <service-ref-name>sun-web.serviceref/calculator</service-ref-name>
    <port-info>
      <wsdl-port>
        <namespaceURI>http://example.web.service/Calculator</namespaceURI>
        <localpart>CalculatorPort</localpart>
      </wsdl-port>
      <service-endpoint-interface>service.web.example.calculator.Calculator</service-endpoint-interface>
      <stub-property name="jbi-enabled" value="true"/>
    </port-info>
  </service-ref>
</sun-web-app>

```

For more information about the `sun-web.xml` and `sun-ejb-jar.xml` deployment descriptor files, see the [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

## Using the Woodstox Parser

The default XML parser in the Enterprise Server is the Sun GlassFish XML Parser (SJSXP). Using the Woodstox parser, which is bundled with the Enterprise Server, may improve performance. Woodstox and SJSXP both provide implementations of the StAX API. To enable the Woodstox parser, set the following system properties for the default `server-config` configuration in the `domain.xml` file, then restart the server:

```

<config name=server-config>
  ...
  <system-property name="javax.xml.stream.XMLEventFactory"
    value="com.ctc.wstx.stax.WstxEventFactory"/>
  <system-property name="javax.xml.stream.XMLInputFactory"
    value="com.ctc.wstx.stax.WstxInputFactory"/>
  <system-property name="javax.xml.stream.XMLOutputFactory"
    value="com.ctc.wstx.stax.WstxOutputFactory"/>
</config>

```

In addition, set these properties for any other configurations referenced by server instances or clusters on which you want to use the Woodstox parser. For more information about the `domain.xml` file and system properties, see the [Sun GlassFish Enterprise Server 2.1 Administration Reference](#).

---

**Note** – If you are using a stand-alone client, you must set these same properties for the client on the java command line as follows:

```
-Djavax.xml.stream.XMLInputFactory=com.ctc.wstx.stax.WstxInputFactory  
-Djavax.xml.stream.XMLOutputFactory=com.ctc.wstx.stax.WstxOutputFactory  
-Djavax.xml.stream.XMLEventFactory=com.ctc.wstx.stax.WstxEventFactory
```

Setting these properties is not necessary if you are using an application client, which is recommended and supported.

---

For more information about the Woodstox parser, see <http://woodstox.codehaus.org/>. For more information about the StAX API, see *Chapter 17: Streaming API for XML* in the *Java EE 5 Tutorial* (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

## Using the Java Persistence API

---

Sun GlassFish Enterprise Server support for the Java Persistence API includes all required features described in the Java Persistence Specification. Although officially part of the Enterprise JavaBeans Specification v3.0, also known as [JSR 220](http://jcp.org/en/jsr/detail?id=220) (<http://jcp.org/en/jsr/detail?id=220>), the Java Persistence API can also be used with non-EJB components outside the EJB container.

The Java Persistence API provides an object/relational mapping facility to Java developers for managing relational data in Java applications. For basic information about the Java Persistence API, see “Part Four: Persistence” in the [Java EE 5 Tutorial](http://java.sun.com/javaee/5/docs/tutorial/doc/index.html) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

This chapter contains Enterprise Server specific information on using the Java Persistence API in the following topics:

- “Specifying the Database” on page 116
- “Additional Database Properties” on page 118
- “Configuring the Cache” on page 118
- “Setting the Logging Level” on page 118
- “Using Lazy Loading” on page 119
- “Primary Key Generation Defaults” on page 119
- “Automatic Schema Generation” on page 120
- “Query Hints” on page 125
- “Changing the Persistence Provider” on page 126
- “Restrictions and Optimizations” on page 127

---

**Note** – The default persistence provider in the Enterprise Server is based on Oracle's TopLink Essentials Java Persistence API implementation. All configuration options in TopLink Essentials are available to applications that use the Enterprise Server's default persistence provider.

---

## Specifying the Database

The Enterprise Server uses the bundled Java DB (Derby) database by default. If the `transaction-type` element is omitted or specified as `JTA` and both the `jta-data-source` and `non-jta-data-source` elements are omitted in the `persistence.xml` file, Java DB is used as a JTA data source. If `transaction-type` is specified as `RESOURCE_LOCAL` and both `jta-data-source` and `non-jta-data-source` are omitted, Java DB is used as a non-JTA data source.

To use a non-default database, either specify a value for the `jta-data-source` element, or set the `transaction-type` element to `RESOURCE_LOCAL` and specify a value for the `non-jta-data-source` element.

If you are using the default persistence provider, the provider attempts to automatically detect the database based on the connection metadata. You can specify the optional `toplink.platform.class.name` property to guarantee that the database is correct. For example:

```
<?xml version="1.0" encoding="UTF-8"?>
  <persistence xmlns="http://java.sun.com/xml/ns/persistence">
    <persistence-unit name="em1">
      <jta-data-source>jdbc/MyDB2DB</jta-data-source>
      <properties>
        <property name="toplink.platform.class.name"
          value="oracle.toplink.essentials.platform.database.DB2Platform"/>
      </properties>
    </persistence-unit>
  </persistence>
```

The following `toplink.platform.class.name` property values are allowed. Supported platforms have been tested with the Enterprise Server and are found to be Java EE compatible.

```
//Supported platforms
oracle.toplink.essentials.platform.database.DerbyPlatform
oracle.toplink.essentials.platform.database.oracle.OraclePlatform
oracle.toplink.essentials.platform.database.SQLServerPlatform
oracle.toplink.essentials.platform.database.DB2Platform
oracle.toplink.essentials.platform.database.SybasePlatform
oracle.toplink.essentials.platform.database.CloudscapePlatform
oracle.toplink.essentials.platform.database.MySQL4Platform
oracle.toplink.essentials.platform.database.PointBasePlatform
oracle.toplink.essentials.platform.database.PostgreSQLPlatform

//Others available
oracle.toplink.essentials.platform.database.InformixPlatform
oracle.toplink.essentials.platform.database.TimesTenPlatform
oracle.toplink.essentials.platform.database.AttunityPlatform
```

```

oracle.toplink.essentials.platform.database.HSQLPlatform
oracle.toplink.essentials.platform.database.SQLAnywherePlatform
oracle.toplink.essentials.platform.database.DBasePlatform
oracle.toplink.essentials.platform.database.DB2MainFramePlatform
oracle.toplink.essentials.platform.database.AccessPlatform

```

To use the Java Persistence API outside the EJB container (in Java SE mode), do not specify the `jta-data-source` or `non-jta-data-source` elements if the `DataSource` is not available. Instead, specify the `provider` element and any additional properties required by the JDBC driver or the database. For example:

```

<?xml version="1.0" encoding="UTF-8"?>
  <persistence xmlns="http://java.sun.com/xml/ns/persistence">
    <persistence-unit name="em2">
      <provider>oracle.toplink.essentials.ejb.cmp3.EntityManagerFactoryProvider</provider>
      <transaction-type>RESOURCE_LOCAL</transaction-type>
      <non-jta-data-source>jdbc/MyDB2DB</non-jta-data-source>
      <properties>
        <property name="toplink.platform.class.name"
          value="oracle.toplink.essentials.platform.database.DB2Platform"/>
        <!-- JDBC connection properties -->
        <property name="toplink.jdbc.driver" value="org.apache.derby.jdbc.ClientDriver"/>
        <property name="toplink.jdbc.url"
value="jdbc:derby://localhost:1527/testdb;retrieveMessagesFromServerOnGetMessage=true;create=true;"/>
        <property name="toplink.jdbc.user" value="APP"/>
        <property name="toplink.jdbc.password" value="APP"/>
      </properties>
    </persistence-unit>
  </persistence>

```

For more information about `toplink` properties, see [“Additional Database Properties”](#) on page 118.

For a list of the JDBC drivers currently supported by the Enterprise Server, see the *Sun GlassFish Enterprise Server 2.1 Release Notes*. For configurations of supported and other drivers, see [“Configurations for Specific JDBC Drivers”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

To change the persistence provider, see [“Changing the Persistence Provider”](#) on page 126.

## Additional Database Properties

If you are using the default persistence provider, you can specify in the `persistence.xml` file the database properties listed at *Persistence Unit Extensions* in [TopLink JPA Extensions Reference](http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html) (<http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html>).

For schema generation properties, see “[Generation Options](#)” on page 122. For query hints, see “[Query Hints](#)” on page 125.

## Configuring the Cache

If you are using the default persistence provider, you can configure whether caching occurs, the type of caching, the size of the cache, and whether client sessions share the cache. Caching properties for the default persistence provider are described in detail at *Extensions for Caching* in [TopLink JPA Extensions Reference](http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html) (<http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html>).

## Setting the Logging Level

One of the default persistence provider's database properties that you can set in the `persistence.xml` file is `toplink.logging.level`. For example, setting the logging level to `FINE` or higher logs all SQL statements. For details about this property, see *Extensions for Logging* in [TopLink JPA Extensions Reference](http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html) (<http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html>).

You can also set the TopLink Essentials logging level globally in the Application Server in any of the following ways:

- Set a `module-log-levels` property using the `asadmin` command. For example:

```
asadmin set --user adminuser "server.log-service.module-log-levels.property.oracle\toplink\essentials"=FINE
```

- Set a JVM option using the `asadmin` command. For example:

```
asadmin create-jvm-options --user adminuser -Dtoplink.logging.level=FINE
```

- Set a `module-log-levels` property using the Admin Console. In the developer profile, select the Application Server component and the Logging tab. In the cluster profile, select the Logger Settings component under the relevant configuration. Select the Log Levels tab. Then scroll down to Additional Module Log Level Properties, select Add Property, type `oracle.toplink.essentials` in the Name field, and type the desired logging level in the Value field.

Setting the logging level to `OFF` disables TopLink Essentials logging. A logging level set in the `persistence.xml` file takes precedence over the global logging level.

You can set the logging level for Java Persistence in general using the Admin Console. In the developer profile, select the Application Server component and the Logging tab. In the cluster profile, select the Logger Settings component under the relevant configuration. Select the Log Levels tab. Then set the logging level for Persistence. Setting the logging level to OFF disables Java Persistence logging.

## Using Lazy Loading

The default persistence provider treats only `OneToOne`, `ManyToOne`, `OneToMany`, and `ManyToMany` mappings specially when they are annotated as `LAZY`. `OneToMany` and `ManyToMany` mappings are loaded lazily by default in compliance with the Java Persistence Specification. Other mappings are always loaded eagerly. For `OneToOne` and `ManyToOne` mappings, value holder indirection is used. For `OneToMany` and `ManyToMany` mappings, transparent indirection is used.

For basic information about lazy loading, see *Lazy Loading* in *TopLink JPA Extensions Reference* (<http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html>). For details about indirection, see *Indirection* in *Mapping Concepts* ([http://www.oracle.com/technology/products/ias/toplink/doc/10131/main/\\_html/mapun002.htm](http://www.oracle.com/technology/products/ias/toplink/doc/10131/main/_html/mapun002.htm)).

## Primary Key Generation Defaults

In the descriptions of the `@GeneratedValue`, `@SequenceGenerator`, and `@TableGenerator` annotations in the Java Persistence Specification, certain defaults are noted as specific to the persistence provider. The default persistence provider's primary key generation defaults are listed here.

`@GeneratedValue` defaults are as follows:

- Using `strategy=AUTO` (or no `strategy`) creates a `@TableGenerator` named `SEQ_GEN` with default settings. Specifying a generator has no effect.
- Using `strategy=TABLE` without specifying a generator creates a `@TableGenerator` named `SEQ_GEN_TABLE` with default settings. Specifying a generator but no `@TableGenerator` creates and names a `@TableGenerator` with default settings.
- Using `strategy=IDENTITY` or `strategy=SEQUENCE` produces the same results, which are database-specific.
  - For Oracle databases, not specifying a generator creates a `@SequenceGenerator` named `SEQ_GEN_SEQUENCE` with default settings. Specifying a generator but no `@SequenceGenerator` creates and names a `@SequenceGenerator` with default settings.
  - For PostgreSQL databases, a `SERIAL` column named `entity-table_pk-column_SEQ` is created.

- For MySQL databases, an `AUTO_INCREMENT` column is created.
- For other supported databases, an `IDENTITY` column is created.

The `@SequenceGenerator` annotation has one default specific to the default provider. The default `sequenceName` is the specified name.

`@TableGenerator` defaults are as follows:

- The default `tableName` is `SEQUENCE`.
- The default `pkColumnName` is `SEQ_NAME`.
- The default `valueColumnName` is `SEQ_COUNT`.
- The default `pkColumnValue` is the specified name, or the default name if no name is specified.

## Automatic Schema Generation

The automatic schema generation feature of the Enterprise Server defines database tables based on the fields or properties in entities and the relationships between the fields or properties. This insulates developers from many of the database related aspects of development, allowing them to focus on entity development. The resulting schema is usable as-is or can be given to a database administrator for tuning with respect to performance, security, and so on. This section covers the following topics:

- [“Annotations” on page 120](#)
- [“Supported Data Types” on page 121](#)
- [“Generation Options” on page 122](#)

---

**Note** – Automatic schema generation is supported on an all-or-none basis: it expects that no tables exist in the database before it is executed. It is not intended to be used as a tool to generate extra tables or constraints.

Deployment won't fail if all tables are not created, and undeployment won't fail if not all tables are dropped. Instead, an error is written to the server log. This is done to allow you to investigate the problem and fix it manually. You should not rely on the partially created database schema to be correct for running the application.

---

## Annotations

The following annotations are used in automatic schema generation: `@AssociationOverride`, `@AssociationOverrides`, `@AttributeOverride`, `@AttributeOverrides`, `@Column`, `@DiscriminatorColumn`, `@DiscriminatorValue`, `@Embedded`, `@EmbeddedId`, `@GeneratedValue`, `@Id`, `@IdClass`, `@JoinColumn`, `@JoinColumns`, `@JoinTable`, `@Lob`, `@ManyToMany`, `@ManyToOne`, `@OneToMany`, `@OneToOne`, `@PrimaryKeyJoinColumn`, `@PrimaryKeyJoinColumns`,



`@SecondaryTable`, `@SecondaryTables`, `@SequenceGenerator`, `@Table`, `@TableGenerator`, `@UniqueConstraint`, and `@Version`. For information about these annotations, see the Java Persistence Specification.

For `@Column` annotations, the `insertable` and `updatable` elements are not used in automatic schema generation.

For `@OneToMany` and `@ManyToOne` annotations, no `ForeignKeyConstraint` is created in the resulting DDL files.

## Supported Data Types

The following table shows mappings of Java types to SQL types when the default persistence provider and automatic schema generation are used.

TABLE 7-1 Java Type to SQL Type Mappings

Java Type	Java DB, Derby, CloudScape	Oracle	DB2	Sybase	MS-SQL Server	MySQL Server
<code>boolean</code> , <code>java.lang.Boolean</code>	SMALLINT	NUMBER(1)	SMALLINT	BIT	BIT	TINYINT(1)
<code>int</code> , <code>java.lang.Integer</code>	INTEGER	NUMBER(10)	INTEGER	INTEGER	INTEGER	INTEGER
<code>long</code> , <code>java.lang.Long</code>	BIGINT	NUMBER(19)	INTEGER	NUMERIC(19)	NUMERIC(19)	BIGINT
<code>float</code> , <code>java.lang.Float</code>	FLOAT	NUMBER(19,4)	FLOAT	FLOAT(16)	FLOAT(16)	FLOAT
<code>double</code> , <code>java.lang.Double</code>	FLOAT	NUMBER(19,4)	FLOAT	FLOAT(32)	FLOAT(32)	DOUBLE
<code>short</code> , <code>java.lang.Short</code>	SMALLINT	NUMBER(5)	SMALLINT	SMALLINT	SMALLINT	SMALLINT
<code>byte</code> , <code>java.lang.Byte</code>	SMALLINT	NUMBER(3)	SMALLINT	SMALLINT	SMALLINT	SMALLINT
<code>java.lang.Number</code>	DECIMAL	NUMBER(38)	DECIMAL(15)	NUMERIC(38)	NUMERIC(28)	DECIMAL(38)
<code>java.math.BigInteger</code>	BIGINT	NUMBER(38)	BIGINT	NUMERIC(38)	NUMERIC(28)	BIGINT
<code>java.math.BigDecimal</code>	DECIMAL	NUMBER(38)	DECIMAL(15)	NUMERIC(38)	NUMERIC(28)	DECIMAL(38)
<code>java.lang.String</code>	VARCHAR(255)	VARCHAR(255)	VARCHAR(255)	VARCHAR(255)	VARCHAR(255)	VARCHAR(255)
<code>char</code> , <code>java.lang.Character</code>	CHAR(1)	CHAR(1)	CHAR(1)	CHAR(1)	CHAR(1)	CHAR(1)
<code>byte[]</code> , <code>java.lang.Byte[]</code> , <code>java.sql.Blob</code>	BLOB(64000)	LONG RAW	BLOB(64000)	IMAGE	IMAGE	BLOB(64000)
<code>char[]</code> , <code>java.lang.Character[]</code> , <code>java.sql.Clob</code>	CLOB(64000)	LONG	CLOB(64000)	TEXT	TEXT	TEXT(64000)

TABLE 7-1 Java Type to SQL Type Mappings (Continued)

Java Type	Java DB, Derby, CloudScape	Oracle	DB2	Sybase	MS-SQL Server	MySQL Server
<code>java.sql.Date</code>	DATE	DATE	DATE	DATETIME	DATETIME	DATE
<code>java.sql.Time</code>	TIME	DATE	TIME	DATETIME	DATETIME	TIME
<code>java.sql.Timestamp</code>	TIMESTAMP	DATE	TIMESTAMP	DATETIME	DATETIME	DATETIME

## Generation Options

Schema generation properties or `asadmin` command line options can control automatic schema generation by the following:

- Creating tables during deployment
- Dropping tables during undeployment
- Dropping and creating tables during redeployment
- Generating the DDL files

---

**Note** – Before using these options, make sure you have a properly configured database. See [“Specifying the Database” on page 116](#).

---

The following optional schema generation properties control the automatic creation of database tables at deployment. You can specify them in the `persistence.xml` file.

TABLE 7-2 Schema Generation Properties

Property	Default	Description
<code>toplink.ddl-generation</code>	<code>none</code>	<p>Specifies whether tables and DDL files are created during deployment, and whether tables are dropped first if they already exist. Allowed values are <code>create-tables</code>, <code>drop-and-create-tables</code>, and <code>none</code>.</p> <p>If <code>create-tables</code> is specified, database tables are created for entities that need them.</p> <p>If <code>drop-and-create-tables</code> is specified, and if tables were automatically created when this application was last deployed, tables from the earlier deployment are dropped and fresh ones are created. If tables were <i>not</i> automatically created when this application was last deployed, no attempt is made to drop any tables. If tables with the same names as those that would have been automatically created are found, the deployment proceeds, but a warning is thrown to indicate that tables could not be created.</p> <p>If <code>none</code> is specified, no tables are created or dropped.</p> <p>The <code>asadmin</code> generation options listed in <a href="#">Table 7-3</a> and <a href="#">Table 7-4</a> override the value of this property.</p> <p>If you are using persistence outside the EJB container and would like to create the DDL files without creating tables, additionally define a Java system property <code>INTERACT_WITH_DB</code> and set its value to <code>false</code>.</p>
<code>toplink.create-ddl-jdbc-file-name</code>	<code>createDDL.jdbc</code>	Specifies the name of the JDBC file that contains the DDL statements required to create the required objects (tables, sequences, and constraints) in the database.
<code>toplink.drop-ddl-jdbc-file-name</code>	<code>dropDDL.jdbc</code>	Specifies the name of the JDBC file that contains the DDL statements required to drop the required objects (tables, sequences, and constraints) from the database.
<code>toplink.application-location</code>	<code>.</code> for the current working directory	<p>Specifies the location where the DDL files are written.</p> <p>For persistence within the EJB container, if this property is not set, DDL files are written to one of the following locations, for applications and modules, respectively:</p> <p><i>domain-dir/generated/ejb/j2ee-apps/app-name</i></p> <p><i>domain-dir/generated/ejb/j2ee-modules/mod-name</i></p>

TABLE 7-2 Schema Generation Properties (Continued)

Property	Default	Description
<code>toplink.ddl-generation.output-mode</code>	both	<p>Specifies the DDL generation target if you are in Java SE mode, outside the EJB container. Values are as follows:</p> <ul style="list-style-type: none"> <li>■ <code>both</code> – Generates SQL files and executes them on the database. If <code>toplink.ddl-generation</code> is set to <code>create-tables</code>, then <code>toplink.create-ddl-jdbc-file-name</code> is written to <code>toplink.application-location</code> and executed on the database. If <code>toplink.ddl-generation</code> is set to <code>drop-and-create-tables</code>, then both <code>toplink.create-ddl-jdbc-file-name</code> and <code>toplink.drop-ddl-jdbc-file-name</code> are written to <code>toplink.application-location</code> and both SQL files are executed on the database.</li> <li>■ <code>database</code> – Executes SQL on the database only (does not generate SQL files). If <code>toplink.ddl-generation</code> is set to <code>create-tables</code>, then <code>toplink.create-ddl-jdbc-file-name</code> is executed on the database. It is not written to <code>toplink.application-location</code>. If <code>toplink.ddl-generation</code> is set to <code>drop-and-create-tables</code>, then both <code>toplink.create-ddl-jdbc-file-name</code> and <code>toplink.drop-ddl-jdbc-file-name</code> are executed on the database. Neither is written to <code>toplink.application-location</code>.</li> <li>■ <code>sql-script</code> – Generates SQL files only (does not execute them on the database). If <code>toplink.ddl-generation</code> is set to <code>create-tables</code>, then <code>toplink.create-ddl-jdbc-file-name</code> is written to <code>toplink.application-location</code>. It is not executed on the database. If <code>toplink.ddl-generation</code> is set to <code>drop-and-create-tables</code>, then both <code>toplink.create-ddl-jdbc-file-name</code> and <code>toplink.drop-ddl-jdbc-file-name</code> are written to <code>toplink.application-location</code>. Neither is executed on the database.</li> </ul>

The following options of the `asadmin deploy` or `asadmin deploydir` command control the automatic creation of database tables at deployment.

TABLE 7-3 The `asadmin deploy` and `asadmin deploydir` Generation Options

Option	Default	Description
<code>--createtables</code>	none	If <code>true</code> , causes database tables to be created for entities that need them. If <code>false</code> , does not create tables. If not specified, the value of the <code>toplink.ddl-generation</code> property in <code>persistence.xml</code> is used.

TABLE 7-3 The `asadmin deploy` and `asadmin deploydir` Generation Options (Continued)

Option	Default	Description
<code>--dropandcreatetables</code>	none	<p>If <code>true</code>, and if tables were automatically created when this application was last deployed, tables from the earlier deployment are dropped and fresh ones are created.</p> <p>If <code>true</code>, and if tables were <i>not</i> automatically created when this application was last deployed, no attempt is made to drop any tables. If tables with the same names as those that would have been automatically created are found, the deployment proceeds, but a warning is thrown to indicate that tables could not be created.</p> <p>If <code>false</code>, the <code>toplink.ddl-generation</code> property setting in <code>persistence.xml</code> is overridden.</p>

The following options of the `asadmin undeploy` command control the automatic removal of database tables at undeployment.

TABLE 7-4 The `asadmin undeploy` Generation Options

Option	Default	Description
<code>--droptables</code>	none	<p>If <code>true</code>, causes database tables that were automatically created when the entities were last deployed to be dropped when the entities are undeployed. If <code>false</code>, does not drop tables.</p> <p>If not specified, tables are dropped only if the <code>toplink.ddl-generation</code> property setting in <code>persistence.xml</code> is <code>drop-and-create-tables</code>.</p>

For more information about the `asadmin deploy`, `asadmin deploydir`, and `asadmin undeploy` commands, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

When `asadmin` deployment options and `persistence.xml` options are both specified, the `asadmin` deployment options take precedence.

The `asant` tasks `sun-appserv-deploy` and `sun-appserv-undeploy` are equivalent to `asadmin deploy` and `asadmin undeploy`, respectively. These `asant` tasks also override the `persistence.xml` options. For details, see [Chapter 3, “The `asant` Utility.”](#)

## Query Hints

Query hints are additional, implementation-specific configuration settings. You can use hints in your queries in the following format:

```
setHint("hint-name", hint-value)
```

For example:

```
Customer customer = (Customer)entityMgr.  
    createNamedQuery("findCustomerBySSN").  
    setParameter("SSN", "123-12-1234").  
    setHint("toplink.refresh", true).  
    getSingleResult();
```

For more information about the query hints available with the default provider, see *Query Hints in TopLink JPA Extensions Reference* (<http://www.oracle.com/technology/products/ias/toplink/jpa/essentials/toplink-jpa-extensions.html>).

## Changing the Persistence Provider

---

**Note** – The previous sections in this chapter apply only to the default persistence provider. If you change the provider for a module or application, the provider-specific database properties, query hints, and schema generation features described in this chapter do not apply.

The `verifier` utility always uses the default provider to verify persistence settings. For information about the `verifier` utility, see “[The verifier Utility](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

---

You can change the persistence provider for an application in the manner described in the Java Persistence API Specification.

First, install the provider. Copy the provider JAR files to the `domain-dir/lib` directory, and restart the Enterprise Server. For more information about the `domain-dir/lib` directory, see “[Using the Common Class Loader](#)” on page 40. The new persistence provider is now available to all modules and applications deployed on servers that share the same configuration. However, the *default* provider remains the same.

In your persistence unit, specify the provider and any properties the provider requires in the `persistence.xml` file. For example:

```
<?xml version="1.0" encoding="UTF-8"?>  
  <persistence xmlns="http://java.sun.com/xml/ns/persistence">  
    <persistence-unit name="em3">  
      <provider>com.company22.persistence.PersistenceProviderImpl</provider>  
      <properties>  
        <property name="company22.database.name" value="MyDB"/>  
      </properties>  
    </persistence-unit>  
  </persistence>
```

# Restrictions and Optimizations

This section discusses restrictions and performance optimizations that affect using the Java Persistence API.

- “Extended Persistence Context Failover” on page 127
- “Using `@OrderBy` with a Shared Session Cache” on page 127
- “Using BLOB or CLOB Types with the Inet Oraxo JDBC Driver” on page 128
- “Database Case Sensitivity” on page 128
- “Sybase Finder Limitation” on page 129
- “MySQL Database Restrictions” on page 129

## Extended Persistence Context Failover

A reference to an extended persistence context in a stateful session bean or an `HttpSession` may not fail over successfully.

The Java Persistence API specification is not clear how the container and persistence provider should work together to passivate a stateful session bean with an extended persistence context in a stand-alone server instance. This also prevents successful serialization and storage of a reference to an extended persistence context in an `HttpSession`.

Even in a single-instance environment, if a stateful session bean is passivated, its extended persistence context could be lost when the stateful session bean is activated. In this environment, it is safe to store an extended persistence context in a stateful session bean only if you can safely disable stateful session bean passivation altogether. This is possible, but trade-offs in memory utilization must be carefully examined before choosing this option.

In a single-instance environment, it is safe to store a reference to an extended persistence context in an `HttpSession`.

## Using `@OrderBy` with a Shared Session Cache

Setting `@OrderBy` on a `ManyToMany` or `OneToMany` relationship field in which a `List` represents the `Many` side doesn't work if the session cache is shared. Use one of the following workarounds:

- Have the application maintain the order so the `List` is cached properly.
- Refresh the session cache using `EntityManager.refresh()` if you don't want to maintain the order during creation or modification of the `List`.
- Disable session cache sharing in `persistence.xml` as follows:

```
<property name="toplink.cache.shared.default" value="false"/>
```

## Using BLOB or CLOB Types with the Inet Oraxo JDBC Driver

To use BLOB or CLOB data types larger than 4 KB for persistence using the Inet Oraxo JDBC Driver for Oracle Databases, you must set the database's `streamsToLob` property value to `true`.

## Database Case Sensitivity

Mapping references to column or table names must be in accordance with the expected column or table name case, and ensuring this is the programmer's responsibility. If column or table names are not explicitly specified for a field or entity, the Enterprise Server uses upper case column names by default, so any mapping references to the column or table names must be in upper case. If column or table names are explicitly specified, the case of all mapping references to the column or table names must be in accordance with the case used in the specified names.

The following are examples of how case sensitivity affects mapping elements that refer to columns or tables. Programmers must keep case sensitivity in mind when writing these mappings.

### Unique Constraints

If column names are not explicitly specified on a field, unique constraints and foreign key mappings must be specified using uppercase references. For example:

```
@Table(name="Department", uniqueConstraints={ @UniqueConstraint ( columnNames= { "DEPTNAME" } ) } )
```

The other way to handle this is by specifying explicit column names for each field with the required case. For example:

```
@Table(name="Department", uniqueConstraints={ @UniqueConstraint ( columnNames= { "deptName" } ) } )  
public class Department{ @Column(name="deptName") private String deptName; }
```

Otherwise, the ALTER TABLE statement generated by the Enterprise Server uses the incorrect case, and the creation of the unique constraint fails.

### Foreign Key Mapping

Use `@OneToMany(mappedBy="COMPANY")` or specify an explicit column name for the Company field on the Many side of the relationship.

### SQL Result Set Mapping

Use the following elements:



```
<sql-result-set-mapping name="SRSMName" >
  <entity-result entity-class="entities.someEntity" />
  <column-result name="UPPERCASECOLUMNNAME" />
</sql-result-set-mapping>
```

Or specify an explicit column name for the upperCaseColumnName field.

## Named Native Queries and JDBC Queries

Column or table names specified in SQL queries must be in accordance with the expected case. For example, MySQL requires column names in the SELECT clause of JDBC queries to be uppercase, while PostgreSQL and Sybase require table names to be uppercase in all JDBC queries.

## PostgreSQL Case Sensitivity

PostgreSQL stores column and table names in lower case. JDBC queries on PostgreSQL retrieve column or table names in lowercase unless the names are quoted. For example:

```
use aliases Select m.ID AS \"ID\" from Department m
```

Use the backslash as an escape character in the class file, but not in the persistence.xml file.

## Sybase Finder Limitation

If a finder method with an input greater than 255 characters is executed and the primary key column is mapped to a VARCHAR column, Sybase attempts to convert type VARCHAR to type TEXT and generates the following error:

```
com.sybase.jdbc2.jdbc.SybSQLException: Implicit conversion from datatype
'TEXT' to 'VARCHAR' is not allowed. Use the CONVERT function to run this
query.
```

To avoid this error, make sure the finder method input is less than 255 characters.

## MySQL Database Restrictions

The following restrictions apply when you use a MySQL database with the Enterprise Server for persistence.

- MySQL treats int1 and int2 as reserved words. If you want to define int1 and int2 as fields in your table, use 'int1' and 'int2' field names in your SQL file.
- When VARCHAR fields get truncated, a warning is displayed instead of an error. To get an error message, start the MySQL database in strict SQL mode.

- The order of fields in a foreign key index must match the order in the explicitly created index on the primary table.
- The CREATE TABLE syntax in the SQL file must end with the following line.

```
) Engine=InnoDB;
```

InnoDB provides MySQL with a transaction-safe (ACID compliant) storage engine having commit, rollback, and crash recovery capabilities.

- For a FLOAT type field, the correct precision must be defined. By default, MySQL uses four bytes to store a FLOAT type that does not have an explicit precision definition. For example, this causes a number such as 12345.67890123 to be rounded off to 12345.7 during an INSERT. To prevent this, specify `FLOAT(10, 2)` in the DDL file, which forces the database to use an eight-byte double-precision column. For more information, see <http://dev.mysql.com/doc/mysql/en/numeric-types.html>.
- To use `||` as the string concatenation symbol, start the MySQL server with the `--sql-mode="PIPES_AS_CONCAT"` option. For more information, see <http://dev.mysql.com/doc/refman/5.0/en/server-sql-mode.html> and <http://dev.mysql.com/doc/mysql/en/ansi-mode.html>.
- MySQL always starts a new connection when `autoCommit=true` is set. This ensures that each SQL statement forms a single transaction on its own. If you try to rollback or commit an SQL statement, you get an error message.

```
javax.transaction.SystemException: java.sql.SQLException:
Can't call rollback when autocommit=true
```

```
javax.transaction.SystemException: java.sql.SQLException:
Error open transaction is not closed
```

To resolve this issue, add `relaxAutoCommit=true` to the JDBC URL. For more information, see <http://forums.mysql.com/read.php?39,31326,31404>.

- MySQL does not allow a DELETE on a row that contains a reference to itself. Here is an example that illustrates the issue.

```
create table EMPLOYEE (
    empId int NOT NULL,
    salary float(25,2) NULL,
    mgrId int NULL,
    PRIMARY KEY (empId),
    FOREIGN KEY (mgrId) REFERENCES EMPLOYEE (empId)
) ENGINE=InnoDB;
```

```
insert into Employee values (1, 1234.34, 1);
delete from Employee where empId = 1;
```

This example fails with the following error message.

ERROR 1217 (23000): Cannot delete or update a parent row:  
a foreign key constraint fails

To resolve this issue, change the table creation script to the following:

```
create table EMPLOYEE (  
    empId int NOT NULL,  
    salary float(25,2) NULL,  
    mgrId int NULL,  
    PRIMARY KEY (empId),  
    FOREIGN KEY (mgrId) REFERENCES EMPLOYEE (empId)  
    ON DELETE SET NULL  
    ) ENGINE=InnoDB;  
  
insert into Employee values (1, 1234.34, 1);  
delete from Employee where empId = 1;
```

This can be done only if the foreign key field is allowed to be null. For more information, see <http://bugs.mysql.com/bug.php?id=12449> and <http://dev.mysql.com/doc/mysql/en/innodb-foreign-key-constraints.html>.



# Developing Web Applications

---

This chapter describes how web applications are supported in the Sun GlassFish Enterprise Server and includes the following sections:

- “Using Servlets” on page 133
- “Using JavaServer Pages” on page 139
- “Creating and Managing Sessions” on page 144
- “Advanced Web Application Features” on page 152

For general information about web applications, see “Part One: The Web Tier” in the *Java EE 5 Tutorial* (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

## Using Servlets

Enterprise Server supports the Java Servlet Specification version 2.5.

---

**Note** – Servlet API version 2.5 is fully backward compatible with versions 2.3 and 2.4, so all existing servlets should work without modification or recompilation.

---

To develop servlets, use Sun Microsystems’ Java Servlet API. For information about using the Java Servlet API, see the documentation provided by Sun Microsystems at <http://java.sun.com/products/servlet/index.html>.

The Enterprise Server provides the `wscompile` and `wsdeploy` tools to help you implement a web service endpoint as a servlet. For more information about these tools, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

This section describes how to create effective servlets to control application interactions running on an Enterprise Server, including standard-based servlets. In addition, this section describes the Enterprise Server features to use to augment the standards.

This section contains the following topics:

- “Invoking a Servlet With a URL” on page 134
- “Servlet Output” on page 134
- “Caching Servlet Results” on page 135
- “About the Servlet Engine” on page 138

## Invoking a Servlet With a URL

You can call a servlet deployed to the Enterprise Server by using a URL in a browser or embedded as a link in an HTML or JSP file. The format of a servlet invocation URL is as follows:

```
http://server:port/context-root/servlet-mapping?name=value
```

The following table describes each URL section.

TABLE 8-1 URL Fields for Servlets Within an Application

URL element	Description
<i>server:port</i>	The IP address (or host name) and optional port number.  To access the default web module for a virtual server, specify only this URL section. You do not need to specify the <i>context-root</i> or <i>servlet-name</i> unless you also wish to specify name-value parameters.
<i>context-root</i>	For an application, the context root is defined in the <i>context-root</i> element of the <i>application.xml</i> , <i>sun-application.xml</i> , or <i>sun-web.xml</i> file. For an individually deployed web module, the context root is specified during deployment.  For both applications and individually deployed web modules, the default context root is the name of the WAR file minus the <i>.war</i> suffix.
<i>servlet-mapping</i>	The <i>servlet-mapping</i> as configured in the <i>web.xml</i> file.
<i>?name=value . . .</i>	Optional request parameters.

In this example, `localhost` is the host name, `MortPages` is the context root, and `calcMortgage` is the servlet mapping:

```
http://localhost:8080/MortPages/calcMortgage?rate=8.0&per=360&bal=180000
```

When invoking a servlet from within a JSP file, you can use a relative path. For example:

```
<jsp:forward page="TestServlet"/>
<jsp:include page="TestServlet"/>
```

## Servlet Output

`ServletContext.log` messages are sent to the server log.

By default, the `System.out` and `System.err` output of servlets are sent to the server log, and during startup, server log messages are echoed to the `System.err` output. Also by default, there is no Windows-only console for the `System.err` output.

You can change these defaults using the Admin Console. In the developer profile, select the Enterprise Server component and the Logging tab. In the cluster profile, select the Logger Settings component under the relevant configuration. Then check or uncheck Write to System Log. If this box is checked, `System.out` output is sent to the server log. If it is unchecked, `System.out` output is sent to the system default location only.

For more information, click the Help button in the Admin Console from the Logging page.

## Caching Servlet Results

The Enterprise Server can cache the results of invoking a servlet, a JSP, or any URL pattern to make subsequent invocations of the same servlet, JSP, or URL pattern faster. The Enterprise Server caches the request results for a specific amount of time. In this way, if another data call occurs, the Enterprise Server can return the cached data instead of performing the operation again. For example, if your servlet returns a stock quote that updates every 5 minutes, you set the cache to expire after 300 seconds.

Whether to cache results and how to cache them depends on the data involved. For example, it makes no sense to cache the results of a quiz submission, because the input to the servlet is different each time. However, it makes sense to cache a high level report showing demographic data taken from quiz results that is updated once an hour.

To define how an Enterprise Server web application handles response caching, you edit specific fields in the `sun-web.xml` file.

---

**Note** – A servlet that uses caching is not portable.

---

For Javadoc tool pages relevant to caching servlet results, go to <http://glassfish.dev.java.net/nonav/javaee5/api/index.html> and click on the `com.sun.appserv.web.cache` package.

For information about JSP caching, see “JSP Caching” on page 140.

The rest of this section covers the following topics:

- “Caching Features” on page 136
- “Default Cache Configuration” on page 136
- “Caching Example” on page 137
- “The CacheKeyGenerator Interface” on page 138

## Caching Features

The Enterprise Server has the following web application response caching capabilities:

- Caching is configurable based on the servlet name or the URI.
- When caching is based on the URI, this includes user specified parameters in the query string. For example, a response from `/garden/catalog?category=roses` is different from a response from `/garden/catalog?category=lilies`. These responses are stored under different keys in the cache.
- Cache size, entry timeout, and other caching behaviors are configurable.
- Entry timeout is measured from the time an entry is created or refreshed. To override this timeout for an individual cache mapping, specify the `cache-mapping subelement timeout`.
- To determine caching criteria programmatically, write a class that implements the `com.sun.appserv.web.cache.CacheHelper` interface. For example, if only a servlet knows when a back end data source was last modified, you can write a helper class to retrieve the last modified timestamp from the data source and decide whether to cache the response based on that timestamp.
- To determine cache key generation programmatically, write a class that implements the `com.sun.appserv.web.cache.CacheKeyGenerator` interface. See [“The CacheKeyGenerator Interface” on page 138](#).
- All non-ASCII request parameter values specified in cache key elements must be URL encoded. The caching subsystem attempts to match the raw parameter values in the request query string.
- Since newly updated classes impact what gets cached, the web container clears the cache during dynamic deployment or reloading of classes.
- The following `HttpServletRequest` request attributes are exposed.
  - `com.sun.appserv.web.cachedServletName`, the cached servlet target
  - `com.sun.appserv.web.cachedURLPattern`, the URL pattern being cached
- Results produced by resources that are the target of a `RequestDispatcher.include()` or `RequestDispatcher.forward()` call are cached if caching has been enabled for those resources. For details, see [“cache-mapping” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#) and [“dispatcher” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#). These are elements in the `sun-web.xml` file.

## Default Cache Configuration

If you enable caching but do not provide any special configuration for a servlet or JSP, the default cache configuration is as follows:

- The default cache timeout is 30 seconds.
- Only the HTTP GET method is eligible for caching.
- HTTP requests with cookies or sessions automatically disable caching.



- No special consideration is given to `Pragma:`, `Cache-control:`, or `Vary:` headers.
- The default key consists of the Servlet Path (minus `pathInfo` and the query string).
- A “least recently used” list is maintained to evict cache entries if the maximum cache size is exceeded.
- Key generation concatenates the servlet path with key field values, if any are specified.
- Results produced by resources that are the target of a `RequestDispatcher.include()` or `RequestDispatcher.forward()` call are never cached.

## Caching Example

Here is an example cache element in the `sun-web.xml` file:

```
<cache max-capacity="8192" timeout="60">
<cache-helper name="myHelper" class-name="MyCacheHelper"/>
<cache-mapping>
  <servlet-name>myservlet</servlet-name>
  <timeout name="timefield">120</timeout>
  <http-method>GET</http-method>
  <http-method>POST</http-method>
</cache-mapping>
<cache-mapping>
  <url-pattern> /catalog/* </url-pattern>
  <!-- cache the best selling category; cache the responses to
  -- this resource only when the given parameters exist. Cache
  -- only when the catalog parameter has 'lilies' or 'roses'
  -- but no other catalog varieties:
  -- /orchard/catalog?best&category='lilies'
  -- /orchard/catalog?best&category='roses'
  -- but not the result of
  -- /orchard/catalog?best&category='wild'
  -->
  <constraint-field name='best' scope='request.parameter' />
  <constraint-field name='category' scope='request.parameter'>
    <value> roses </value>
    <value> lilies </value>
  </constraint-field>
  <!-- Specify that a particular field is of given range but the
  -- field doesn't need to be present in all the requests -->
  <constraint-field name='SKUnum' scope='request.parameter'>
    <value match-expr='in-range'> 1000 - 2000 </value>
  </constraint-field>
  <!-- cache when the category matches with any value other than
  -- a specific value -->
  <constraint-field name="category" scope="request.parameter">
    <value match-expr="equals" cache-on-match-failure="true">
      bogus
```

```
        </value>
    </constraint-field>
</cache-mapping>
<cache-mapping>
    <servlet-name> InfoServlet </servlet-name>
    <cache-helper-ref>myHelper</cache-helper-ref>
</cache-mapping>
</cache>
```

For more information about the `sun-web.xml` caching settings, see “[cache](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## The CacheKeyGenerator Interface

The built-in default `CacheHelper` implementation allows web applications to customize the key generation. An application component (in a servlet or JSP) can set up a custom `CacheKeyGenerator` implementation as an attribute in the `ServletContext`.

The name of the context attribute is configurable as the value of the `cacheKeyGeneratorAttrName` property in the `default-helper` element of the `sun-web.xml` deployment descriptor. For more information, see “[default-helper](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## About the Servlet Engine

Servlets exist in and are managed by the servlet engine in the Enterprise Server. The servlet engine is an internal object that handles all servlet meta functions. These functions include instantiation, initialization, destruction, access from other components, and configuration management. This section covers the following topics:

- “[Instantiating and Removing Servlets](#)” on page 138
- “[Request Handling](#)” on page 139

## Instantiating and Removing Servlets

After the servlet engine instantiates the servlet, the servlet engine calls the servlet’s `init()` method to perform any necessary initialization. You can override this method to perform an initialization function for the servlet’s life, such as initializing a counter.

When a servlet is removed from service, the servlet engine calls the `destroy()` method in the servlet so that the servlet can perform any final tasks and deallocate resources. You can override this method to write log messages or clean up any lingering connections that won’t be caught in garbage collection.

## Request Handling

When a request is made, the Enterprise Server hands the incoming data to the servlet engine. The servlet engine processes the request's input data, such as form data, cookies, session information, and URL name-value pairs, into an `HttpServletRequest` request object type.

The servlet engine also creates an `HttpServletResponse` response object type. The engine then passes both as parameters to the servlet's `service()` method.

In an HTTP servlet, the default `service()` method routes requests to another method based on the HTTP transfer method: POST, GET, DELETE, HEAD, OPTIONS, PUT, or TRACE. For example, HTTP POST requests are sent to the `doPost()` method, HTTP GET requests are sent to the `doGet()` method, and so on. This enables the servlet to process request data differently, depending on which transfer method is used. Since the routing takes place in the `service()` method, you generally do not override `service()` in an HTTP servlet. Instead, override `doGet()`, `doPost()`, and so on, depending on the request type you expect.

To perform the tasks to answer a request, override the `service()` method for generic servlets, and the `doGet()` or `doPost()` methods for HTTP servlets. Very often, this means accessing EJB components to perform business transactions, then collating the information in the request object or in a `JDBC ResultSet` object.

## Using JavaServer Pages

The Enterprise Server supports the following JSP features:

- JavaServer Pages (JSP) Specification version 2.1
- Precompilation of JSP files, which is especially useful for production servers
- JSP tag libraries and standard portable tags

For information about creating JSP files, see Sun Microsystem's JavaServer Pages web site at <http://java.sun.com/products/jsp/index.html>.

For information about Java Beans, see Sun Microsystem's JavaBeans web page at <http://java.sun.com/beans/index.html>.

This section describes how to use JavaServer Pages (JSP files) as page templates in an Enterprise Server web application. This section contains the following topics:

- "JSP Tag Libraries and Standard Portable Tags" on page 140
- "JSP Caching" on page 140
- "Options for Compiling JSP Files" on page 144

## JSP Tag Libraries and Standard Portable Tags

Enterprise Server supports tag libraries and standard portable tags. For more information, see the JavaServer Pages Standard Tag Library (JSTL) page at <http://java.sun.com/products/jsp/jstl/index.jsp>.

Web applications don't need to bundle copies of the `jsf-impl.jar` or `appserv-jstl.jar` JSP tag libraries (in `as-install/lib`) to use JavaServer™ Faces technology or JSTL, respectively. These tag libraries are automatically available to all web applications.

However, the `as-install/lib/appserv-tags.jar` tag library for JSP caching is not automatically available to web applications. See “JSP Caching” on page 140, next.

## JSP Caching

JSP caching lets you cache tag invocation results within the Java engine. Each can be cached using different cache criteria. For example, suppose you have invocations to view stock quotes, weather information, and so on. The stock quote result can be cached for 10 minutes, the weather report result for 30 minutes, and so on. JSP caching is described in the following topics:

- “The `appserv-tags.jar` File” on page 140
- “Caching Scope” on page 141
- “The cache Tag” on page 142
- “The flush Tag” on page 143

For more information about response caching as it pertains to servlets, see “Caching Servlet Results” on page 135.

### The `appserv-tags.jar` File

JSP caching is implemented by a tag library packaged into the `as-install/lib/appserv-tags.jar` file, which you can copy into the `WEB-INF/lib` directory of your web application. The `appserv-tags.tld` tag library descriptor file is in the `META-INF` directory of this JAR file.

---

**Note** – Web applications that use this tag library without bundling it are not portable.

---

To allow all web applications to share this tag library, change the following elements in the `domain.xml` file. Change this:

```
<jvm-options>
-Dcom.sun.enterprise.taglibs=appserv-jstl.jar,jsf-impl.jar
</jvm-options>
```

to this:

```
<jvm-options>
-Dcom.sun.enterprise.taglibs=appserv-jstl.jar,jsf-impl.jar,appserv-tags.jar
</jvm-options>
```

and this:

```
<jvm-options>
-Dcom.sun.enterprise.taglisteners=jsf-impl.jar
</jvm-options>
```

to this:

```
<jvm-options>
-Dcom.sun.enterprise.taglisteners=jsf-impl.jar,appserv-tags.jar
</jvm-options>
```

For more information about the `domain.xml` file, see the [Sun GlassFish Enterprise Server 2.1 Administration Reference](#).

Refer to these tags in JSP files as follows:

```
<%@ taglib prefix="prefix" uri="Sun ONE Application Server Tags" %>
```

Subsequently, the cache tags are available as `<prefix:cache>` and `<prefix:flush>`. For example, if your *prefix* is *mypfx*, the cache tags are available as `<mypfx:cache>` and `<mypfx:flush>`.

## Caching Scope

JSP caching is available in three different scopes: request, session, and application. The default is application. To use a cache in request scope, a web application must specify the `com.sun.appserv.web.taglibs.cache.CacheRequestListener` in its `web.xml` deployment descriptor, as follows:

```
<listener>
  <listener-class>
    com.sun.appserv.web.taglibs.cache.CacheRequestListener
  </listener-class>
</listener>
```

Likewise, for a web application to utilize a cache in session scope, it must specify the `com.sun.appserv.web.taglibs.cache.CacheSessionListener` in its `web.xml` deployment descriptor, as follows:

```
<listener>
  <listener-class>
    com.sun.appserv.web.taglibs.cache.CacheSessionListener
  </listener-class>
</listener>
```

To utilize a cache in application scope, a web application need not specify any listener. The `com.sun.appserv.web.taglibs.cache.CacheContextListener` is already specified in the `appserv-tags.tld` file.

## The cache Tag

The cache tag caches the body between the beginning and ending tags according to the attributes specified. The first time the tag is encountered, the body content is executed and cached. Each subsequent time it is run, the cached content is checked to see if it needs to be refreshed and if so, it is executed again, and the cached data is refreshed. Otherwise, the cached data is served.

### Attributes of cache

The following table describes attributes for the cache tag.

TABLE 8-2 The cache Attributes

Attribute	Default	Description
key	<i>ServletPath_Suffix</i>	(optional) The name used by the container to access the cached entry. The cache key is suffixed to the servlet path to generate a key to access the cached entry. If no key is specified, a number is generated according to the position of the tag in the page.
timeout	60s	(optional) The time in seconds after which the body of the tag is executed and the cache is refreshed. By default, this value is interpreted in seconds. To specify a different unit of time, add a suffix to the timeout value as follows: s for seconds, m for minutes, h for hours, d for days. For example, 2h specifies two hours.
nocache	false	(optional) If set to true, the body content is executed and served as if there were no cache tag. This offers a way to programmatically decide whether the cached response is sent or whether the body has to be executed, though the response is not cached.
refresh	false	(optional) If set to true, the body content is executed and the response is cached again. This lets you programmatically refresh the cache immediately regardless of the timeout setting.
scope	application	(optional) The scope of the cache. Can be request, session, or application. See <a href="#">“Caching Scope” on page 141</a> .

### Example of cache

The following example represents a cached JSP file:

```
<%@ taglib prefix="mypfx" uri="Sun ONE Application Server Tags" %>
<%@ taglib prefix="c" uri="http://java.sun.com/jsp/jstl/core" %>
<mypfx:cache
    key="${sessionScope.loginId}"
```

```

        nocache="${param.nocache}"
        refresh="${param.refresh}"
        timeout="10m">
<c:choose>
  <c:when test="${param.page == 'frontPage'}">
    <!-- get headlines from database -->
  </c:when>
  <c:otherwise>
    ...
  </c:otherwise>
</c:choose>
</myafx:cache>
<myafx:cache timeout="1h">
<h2> Local News </h2>
  <!-- get the headline news and cache them -->
</myafx:cache>

```

## The flush Tag

Forces the cache to be flushed. If a key is specified, only the entry with that key is flushed. If no key is specified, the entire cache is flushed.

### Attributes of flush

The following table describes attributes for the flush tag.

TABLE 8-3 The flush Attributes

Attribute	Default	Description
key	<i>ServletPath_Suffix</i>	(optional) The name used by the container to access the cached entry. The cache key is suffixed to the servlet path to generate a key to access the cached entry. If no key is specified, a number is generated according to the position of the tag in the page.
scope	application	(optional) The scope of the cache. Can be request, session, or application. See “ <a href="#">Caching Scope</a> ” on page 141.

### Examples of flush

To flush the entry with key="foobar":

```
<myafx:flush key="foobar"/>
```

To flush the entire cache:

```
<c:if test="${empty sessionScope.clearCache}">
  <myafx:flush />
</c:if>
```

## Options for Compiling JSP Files

Enterprise Server provides the following ways of compiling JSP 2.1 compliant source files into servlets:

- JSP files are automatically compiled at runtime.
- The `asadmin deploy` command has a `precompilejsp` option. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.
- The `sun-appserv-jspc` Ant task allows you to precompile JSP files; see “The `sun-appserv-jspc` Task” on page 58.
- The `jspc` command line tool allows you to precompile JSP files at the command line. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Creating and Managing Sessions

This chapter describes how to create and manage HTTP sessions that allows users and transaction information to persist between interactions.

This chapter contains the following sections:

- “Configuring Sessions” on page 144
- “Session Managers” on page 147

## Configuring Sessions

This section covers the following topics:

- “HTTP Sessions, Cookies, and URL Rewriting” on page 144
- “Coordinating Session Access” on page 145
- “Distributed Sessions and Persistence” on page 145

### HTTP Sessions, Cookies, and URL Rewriting

To configure whether and how HTTP sessions use cookies and URL rewriting, edit the `session-properties` and `cookie-properties` elements in the `sun-web.xml` file for an individual web application. For more about the properties you can configure, see “`session-properties`” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide* and “`cookie-properties`” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

For information about configuring default session properties for the entire web container, see Chapter 8, “SIP, Web, and EJB Containers,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide* and the *Sun GlassFish Enterprise Server 2.1 High Availability Administration Guide*.



## Coordinating Session Access

Make sure that multiple threads don't simultaneously modify the same session object in conflicting ways. If the persistence type is `ha` (see “[The ha Persistence Type](#)” on page 150), the following message in the log file indicates that this might be happening:

```
Primary Key Constraint violation while saving session session_id
```

This is especially likely to occur in web applications that use HTML frames where multiple servlets are executing simultaneously on behalf of the same client. A good solution is to ensure that one of the servlets modifies the session and the others have read-only access.

## Distributed Sessions and Persistence

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “[Usage Profiles](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

A distributed HTTP session can run in multiple Enterprise Server instances, provided the following criteria are met:

- Each server instance has the same distributable web application deployed to it. The `web-app` element of the `web.xml` deployment descriptor file must have the `distributable` subelement specified.
- The web application uses high-availability session persistence. If a non-distributable web application is configured to use high-availability session persistence, a warning is written to the server log, and the session persistence type reverts to memory. See “[The replicated Persistence Type](#)” on page 149 or “[The ha Persistence Type](#)” on page 150.
- All objects bound into a distributed session must be of the types listed in [Table 8–4](#).
- The web application must be deployed using the `deploy` or `deploydir` command with the `--availabilityenabled` option set to `true`. See the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.
- Optionally, each server instance has access to the same high-availability database (HADB). HADB must be installed and the enterprise profile must be selected. For information about profiles, see “[Usage Profiles](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*. For information about HADB, see the description of the `configure-ha-cluster` command in the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

---

**Note** – Contrary to the Servlet 2.5 specification, Enterprise Server does not throw an `IllegalArgumentException` if an object type not supported for failover is bound into a distributed session.

Keep the distributed session size as small as possible. Session size has a direct impact on overall system throughput.

---

In the event of an instance or hardware failure, another server instance can take over a distributed session, with the following limitations:

- If a distributable web application references a Java EE component or resource, the reference might be lost. See [Table 8-4](#) for a list of the types of references that `HTTPSession` failover supports.
- References to open files or network connections are lost.

In the following table, *No* indicates that failover for the object type might not work in all cases and that no failover support is provided. However, failover might work in some cases for that object type. For example, failover might work because the class implementing that type is serializable.

For more information about the `InitialContext`, see [“Accessing the Naming Context” on page 269](#). For more information about transaction recovery, see [Chapter 16, “Using the Transaction Service.”](#) For more information about Administered Objects, see [“Creating Physical Destinations” on page 279](#).

**TABLE 8-4** Object Types Supported for Java EE Web Application Session State Failover

Java Object Type	Failover Support
Colocated or distributed stateless session, stateful session, or entity bean reference	Yes
JNDI context	Yes, <code>InitialContext</code> and <code>java:comp/env</code>
UserTransaction	Yes, but if the instance that fails is never restarted, any prepared global transactions are lost and might not be correctly rolled back or committed.
JDBC DataSource	No
Java Message Service (JMS) ConnectionFactory, Destination	No
JavaMail Session	No
Connection Factory	No

TABLE 8-4 Object Types Supported for Java EE Web Application Session State Failover (Continued)

Java Object Type	Failover Support
Administered Object	No
Web service reference	No
Serializable Java types	Yes
Extended persistence context	No

## Session Managers

A session manager automatically creates new session objects whenever a new session starts. In some circumstances, clients do not join the session, for example, if the session manager uses cookies and the client does not accept cookies.

Enterprise Server offers these session management options, determined by the session-manager element's persistence-type attribute in the sun-web.xml file:

- “The memory Persistence Type” on page 147, the default
- “The file Persistence Type” on page 148, which uses a file to store session data
- “The replicated Persistence Type” on page 149, which uses other servers in the cluster for session persistence
- “The ha Persistence Type” on page 150, which uses the high-availability database for session persistence

---

**Note** – If the session manager configuration contains an error, the error is written to the server log and the default (memory) configuration is used.

---

For more information, see “session-manager” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

### The memory Persistence Type

This persistence type is not designed for a production environment that requires session persistence. It provides no session persistence. However, you can configure it so that the session state in memory is written to the file system prior to server shutdown.

To specify the memory persistence type for the entire web container, use the configure-ha-persistence command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

To specify the memory persistence type for a specific web application, edit the sun-web.xml file as in the following example. The persistence-type property is optional, but must be set to memory if included. This overrides the web container availability settings for the web application.

```
<sun-web-app>
...
<session-config>
  <session-manager persistence-type="memory" />
  <manager-properties>
    <property name="sessionFilename" value="sessionstate" />
  </manager-properties>
</session-manager>
...
</session-config>
...
</sun-web-app>
```

The only manager property that the memory persistence type supports is `sessionFilename`, which is listed under “[manager-properties](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

For more information about the `sun-web.xml` file, see *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## The file Persistence Type

This persistence type provides session persistence to the local file system, and allows a single server domain to recover the session state after a failure and restart. The session state is persisted in the background, and the rate at which this occurs is configurable. The store also provides passivation and activation of the session state to help control the amount of memory used. This option is not supported in a production environment. However, it is useful for a development system with a single server instance.

---

**Note** – Make sure the `delete` option is set in the `server.policy` file, or expired file-based sessions might not be deleted properly. For more information about `server.policy`, see “[The server.policy File](#)” on page 89.

---

To specify the file persistence type for the entire web container, use the `configure-ha-persistence` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

To specify the file persistence type for a specific web application, edit the `sun-web.xml` file as in the following example. Note that `persistence-type` must be set to `file`. This overrides the web container availability settings for the web application.

```
<sun-web-app>
...
<session-config>
  <session-manager persistence-type="file">
```

```

        <store-properties>
            <property name="directory" value="sessiondir" />
        </store-properties>
    </session-manager>
    ...
</session-config>
...
</sun-web-app>

```

The file persistence type supports all the manager properties listed under “manager-properties” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide* except `sessionFilename`, and supports the `directory` store property listed under “store-properties” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

For more information about the `sun-web.xml` file, see *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## The replicated Persistence Type

The replicated persistence type uses other servers in the cluster for session persistence. Clustered server instances replicate session state in a ring topology. Each backup instance stores the replicated data in memory. This allows sessions to be distributed. For details, see “Distributed Sessions and Persistence” on page 145. In addition, you can configure the frequency and scope of session persistence. The other servers are also used as the passivation and activation store. Use this option in a production environment that requires session persistence.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

To use the replicated persistence type, you must enable availability. Select the Availability Service component under the relevant configuration in the Admin Console. Check the Availability Service box. To enable availability for the web container, select the Web Container Availability tab, then check the Availability Service box. All instances in an Enterprise Server cluster should have the same availability settings to ensure consistent behavior. For details, see the *Sun GlassFish Enterprise Server 2.1 High Availability Administration Guide*.

To change settings such as persistence frequency and persistence scope for the entire web container, use the Persistence Frequency and Persistence Scope drop-down lists on the Web Container Availability tab in the Admin Console, or use the `asadmin set` command. For example:

```

asadmin set
server-config.availability-service.web-container-availability.persistence-frequency=time-based

```

For more information, see the description of the `asadmin set` command in the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

To specify the replicated persistence type for a specific web application, edit the `sun-web.xml` file as in the following example. Note that `persistence-type` must be set to `replicated`. This overrides the web container availability settings for the web application.

```
<sun-web-app>
...
<session-config>
  <session-manager persistence-type="replicated">
    <manager-properties>
      <property name="persistenceFrequency" value="web-method" />
    </manager-properties>
    <store-properties>
      <property name="persistenceScope" value="session" />
    </store-properties>
  </session-manager>
  ...
</session-config>
...
</sun-web-app>
```

The replicated persistence type supports all the manager properties listed under “[manager-properties](#)” in [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#) except `sessionFilename`, and supports the `persistenceScope` store property listed under “[store-properties](#)” in [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

For more information about the `sun-web.xml` file, see [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

## The ha Persistence Type

The ha persistence type uses the high-availability database (HADB) for session persistence. The HADB must be installed and the enterprise profile must be selected. For information about profiles, see “[Usage Profiles](#)” in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

The HADB allows sessions to be distributed. For details, see “[Distributed Sessions and Persistence](#)” on page 145. In addition, you can configure the frequency and scope of session persistence. The HADB is also used as the passivation and activation store. Use this option in a production environment that requires session persistence.

The HADB must be configured and enabled before you can use distributed sessions. For configuration details, see the description of the `configure-ha-cluster` command in the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

To enable the HADB, select the Availability Service component under the relevant configuration in the Admin Console. Check the Instance Level Availability box. To enable availability for the web container, select the Web Container Availability tab, then check the Availability Service box. All instances in an Enterprise Server cluster should have the same availability settings to ensure consistent behavior. For details, see the *Sun GlassFish Enterprise Server 2.1 High Availability Administration Guide*.

To change settings such as persistence frequency and persistence scope for the entire web container, see the description of the `configure-ha-persistence` command in the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

To specify the `ha` persistence type for a specific web application, edit the `sun-web.xml` file as in the following example. Note that `persistence-type` must be set to `ha`. This overrides the web container availability settings for the web application.

```
<sun-web-app>
...
<session-config>
  <session-manager persistence-type="ha">
    <manager-properties>
      <property name="persistenceFrequency" value="web-method" />
    </manager-properties>
    <store-properties>
      <property name="persistenceScope" value="session" />
    </store-properties>
  </session-manager>
  ...
</session-config>
...
</sun-web-app>
```

The `ha` persistence type supports all the manager properties listed under “`manager-properties`” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide* except `sessionFilename`, and supports the `persistenceScope` store property listed under “`store-properties`” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

For more information about the `sun-web.xml` file, see *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

# Advanced Web Application Features

This section includes summaries of the following topics:

- “Internationalization Issues” on page 152
- “Virtual Servers” on page 153
- “Default Web Modules” on page 154
- “Class Loader Delegation” on page 155
- “Using the default-web.xml File” on page 155
- “Configuring Logging and Monitoring in the Web Container” on page 156
- “Configuring Idempotent URL Requests” on page 156
- “Header Management” on page 157
- “Configuring Valves and Catalina Listeners” on page 157
- “Alternate Document Roots” on page 158
- “Redirecting URLs” on page 159
- “Enabling Comet Support” on page 160
- “Using a context.xml File” on page 160
- “Enabling WebDav” on page 160
- “Using mod\_jk” on page 162

## Internationalization Issues

This section covers internationalization as it applies to the following:

- “The Server's Default Locale” on page 152
- “Servlet Character Encoding” on page 152

### The Server's Default Locale

To set the default locale of the entire Enterprise Server, which determines the locale of the Admin Console, the logs, and so on, use the Admin Console. In the developer profile, select the Enterprise Server component, the Advanced tab, and the Domain Attributes tab. In the cluster profile, select the domain component. Then type a value in the Locale field. For details, click the Help button in the Admin Console.

### Servlet Character Encoding

This section explains how the Enterprise Server determines the character encoding for the servlet request and the servlet response. For encodings you can use, see <http://java.sun.com/javase/6/docs/technotes/guides/intl/encoding.doc.html>.

### Servlet Request

When processing a servlet request, the server uses the following order of precedence, first to last, to determine the request character encoding:



- The `getCharacterEncoding()` method
- A hidden field in the form, specified by the `form-hint-field` attribute of the `parameter-encoding` element in the `sun-web.xml` file
- The `default-charset` attribute of the `parameter-encoding` element in the `sun-web.xml` file
- The default, which is ISO-8859-1

For details about the `parameter-encoding` element, see “[parameter-encoding](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## Servlet Response

When processing a servlet response, the server uses the following order of precedence, first to last, to determine the response character encoding:

- The `setCharacterEncoding()` or `setContentType()` method
- The `setLocale()` method
- The default, which is ISO-8859-1

## Virtual Servers

A virtual server, also called a virtual host, is a virtual web server that serves content targeted for a specific URL. Multiple virtual servers can serve content using the same or different host names, port numbers, or IP addresses. The HTTP service directs incoming web requests to different virtual servers based on the URL.

When you first install the Enterprise Server, a default virtual server is created. You can also assign a default virtual server to each new HTTP listener you create.

Web applications and Java EE applications containing web components (web modules) can be assigned to virtual servers during deployment. A web module can be assigned to more than one virtual server, and a virtual server can have more than one web module assigned to it.

For more information about virtual servers, see “[virtual-server](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Reference*.

### ▼ To Assign a Default Virtual Server

- 1 In the Admin Console, open the HTTP Service component under the relevant configuration.
- 2 Open the HTTP Listeners component under the HTTP Service component.
- 3 Select or create a new HTTP listener.

**4 Select from the Default Virtual Server drop-down list.**

For more information, see [“Default Web Modules” on page 154](#).

**See Also** For details, click the Help button in the Admin Console from the HTTP Listeners page.

**▼ To Assign Virtual Servers****1 Deploy the application or web module and assign the desired virtual servers to it.**

For more information, see *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

**2 In the Admin Console, open the HTTP Service component under the relevant configuration.****3 Open the Virtual Servers component under the HTTP Service component.****4 Select the virtual server to which you want to assign a default web module.****5 Select the application or web module from the Default Web Module drop-down list.**

For more information, see [“Default Web Modules” on page 154](#).

**See Also** For details, click the Help button in the Admin Console from the Virtual Servers page.

## Default Web Modules

A default web module can be assigned to the default virtual server and to each new virtual server. For details, see [“Virtual Servers” on page 153](#). To access the default web module for a virtual server, point the browser to the URL for the virtual server, but do not supply a context root. For example:

```
http://myvserver:3184/
```

A virtual server with no default web module assigned serves HTML or JavaServer Pages™ (JSP™) content from its document root, which is usually *domain-dir/doc root*. To access this HTML or JSP content, point your browser to the URL for the virtual server, do not supply a context root, but specify the target file.

For example:

```
http://myvserver:3184/hellothere.jsp
```

## Class Loader Delegation

The Servlet specification recommends that the Web class loader look in the local class loader before delegating to its parent. To make the Web class loader follow the delegation model in the Servlet specification, set `delegate="false"` in the `class-loader` element of the `sun-web.xml` file. It's safe to do this only for a web module that does not interact with any other modules.

The default value is `delegate="true"`, which causes the Web class loader to delegate in the same manner as the other class loaders. Use `delegate="true"` for a web application that accesses EJB components or that acts as a web service client or endpoint. For details about `sun-web.xml`, see [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

For general information about class loaders, see [Chapter 2, "Class Loaders."](#)

## Using the default-web.xml File

You can use the `default-web.xml` file to define features such as filters and security constraints that apply to all web applications.

For example, you can disable directory listings for added security. In your domain's `default-web.xml` file, search for the definition of the servlet whose `servlet-name` is equal to `default`, and set the value of the `init-param` named `listings` to `false`. Then redeploy your web application if it has already been deployed.

```
<init-param>
  <param-name>listings</param-name>
  <param-value>>false</param-value>
</init-param>
```

The `mime-mapping` elements in `default-web.xml` are global and inherited by all web applications. You can override these mappings or define your own using `mime-mapping` elements in your web application's `web.xml` file. For more information about `mime-mapping` elements, see the Servlet specification.

You can use the Admin Console to edit the `default-web.xml` file. For details, click the Help button in the Admin Console. As an alternative, you can edit the file directly using the following steps.

### ▼ To Use the default-web.xml File

- 1 Place the JAR file for the filter, security constraint, or other feature in the `domain-dir/lib` directory.
- 2 Edit the `domain-dir/config/default-web.xml` file to refer to the JAR file.

### 3 Restart the server.

## Configuring Logging and Monitoring in the Web Container

For information about configuring logging and monitoring in the web container using the Admin Console, click the Help button in the Admin Console. In the developer profile, Logging and Monitor tabs are accessible from the Application Server page. In the Cluster profile, select Logger Settings under the relevant configuration, or select the Stand-Alone Instances component, select the instance from the table, and select the Monitor tab.

## Configuring Idempotent URL Requests

An *idempotent* request is one that does not cause any change or inconsistency in an application when retried. To enhance the availability of your applications deployed on an Enterprise Server cluster, configure the load balancer to retry failed idempotent HTTP requests on all the Enterprise Server instances in a cluster. This option can be used for read-only requests, for example, to retry a search request.

This section describes the following topics:

- “Specifying an Idempotent URL” on page 156
- “Characteristics of an Idempotent URL” on page 156

### Specifying an Idempotent URL

To configure idempotent URL response, specify the URLs that can be safely retried in `idempotent-url-pattern` elements in the `sun-web.xml` file. For example:

```
<idempotent-url-pattern url-pattern="sun_java/*" no-of-retries="10"/>
```

For details, see “`idempotent-url-pattern`” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

If none of the server instances can successfully serve the request, an error page is returned.

### Characteristics of an Idempotent URL

Since all requests for a given session are sent to the same application server instance, and if that Enterprise Server instance is unreachable, the load balancer returns an error message. Normally, the request is not retried on another Enterprise Server instance. However, if the URL pattern matches that specified in the `sun-web.xml` file, the request is implicitly retried on another Enterprise Server instance in the cluster.

In HTTP, some methods (such as GET) are idempotent, while other methods (such as POST) are not. In effect, retrying an idempotent URL should not cause values to change on the server or in the database. The only difference should be a change in the response received by the user.

Examples of idempotent requests include search engine queries and database queries. The underlying principle is that the retry does not cause an update or modification of data.

A search engine, for example, sends HTTP requests with the same URL pattern to the load balancer. Specifying the URL pattern of the search request to the load balancer ensures that HTTP requests with the specified URL pattern are implicitly retried on another Enterprise Server instance.

For example, if the request URL sent to the Enterprise Server is of the type `/search/something.html`, then the URL pattern can be specified as `/search/*`.

Examples of non-idempotent requests include banking transactions and online shopping. If you retry such requests, money might be transferred twice from your account.

## Header Management

In all Editions of the Enterprise Server, the Enumeration from `request.getHeaders()` contains multiple elements (one element per request header) instead of a single, aggregated value.

The header names used in `HttpServletResponse.addXXXHeader()` and `HttpServletResponse.setXXXHeader()` are returned as they were created.

## Configuring Valves and Catalina Listeners

You can configure custom valves and Catalina listeners for web modules or virtual servers by defining properties. In the `domain.xml` file, valve and listener properties look like this:

```
<web-module ...>
  <property name="valve_1" value="org.glassfish.extension.Valve"/>
  <property name="listener_1" value="org.glassfish.extension.MyLifecycleListener"/>
</web-module>
```

You can define these properties in one of the following ways, then restart the server:

- You can define properties using the `asadmin set` command. For example:

```
asadmin set domain1.applications.web-module.MyWebMod.property.valve_1="org.glassfish.extension.Valve"
```

```
asadmin set config1.http-service.virtual-server.MyVS.property.valve_1="org.glassfish.extension.Valve"
```

- You can define virtual server properties using the Admin Console. Select the HTTP Service component under the relevant configuration, select Virtual Servers, and select the desired virtual server. Select Add Property, enter the property name and value, check the enable box, and select Save. For details, click the Help button in the Admin Console.

## Alternate Document Roots

An alternate document root (docroot) allows a web application to serve requests for certain resources from outside its own docroot, based on whether those requests match one (or more) of the URI patterns of the web application's alternate docroots.

To specify an alternate docroot for a web application or a virtual server, use the `alternatedocroot_n` property, where *n* is a positive integer that allows specification of more than one. This property can be a subelement of a `sun-web-app` element in the `sun-web.xml` file or a `virtual-server` element in the `domain.xml` file. For more information about these elements, see “[sun-web-app](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide* and “[virtual-server](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Reference*.

A virtual server's alternate docroots are considered only if a request does not map to any of the web modules deployed on that virtual server. A web module's alternate docroots are considered only once a request has been mapped to that web module.

If a request matches an alternate docroot's URI pattern, it is mapped to the alternate docroot by appending the request URI (minus the web application's context root) to the alternate docroot's physical location (directory). If a request matches multiple URI patterns, the alternate docroot is determined according to the following precedence order:

- Exact match
- Longest path match
- Extension match

For example, the following properties specify three docroots. The URI pattern of the first alternate docroot uses an exact match, whereas the URI patterns of the second and third alternate docroots use extension and longest path prefix matches, respectively.

```
<property name="alternatedocroot_1" value="from=/my.jpg dir=/srv/images/jpg"/>
<property name="alternatedocroot_2" value="from=*.jpg dir=/srv/images/jpg"/>
<property name="alternatedocroot_3" value="from=/jpg/* dir=/src/images"/>
```

The value of each alternate docroot has two components: The first component, `from`, specifies the alternate docroot's URI pattern, and the second component, `dir`, specifies the alternate docroot's physical location (directory).

Suppose the above examples belong to a web application deployed at `http://company22.com/myapp`. The first alternate docroot maps any requests with this URL:

```
http://company22.com/myapp/my.jpg
```

To this resource:

```
/svr/images/jpg/my.jpg
```

The second alternate docroot maps any requests with a \*.jpg suffix, such as:

```
http://company22.com/myapp/*.jpg
```

To this physical location:

```
/svr/images/jpg
```

The third alternate docroot maps any requests whose URI starts with /myapp/jpg/, such as:

```
http://company22.com/myapp/jpg/*
```

To the same directory as the second alternate docroot.

For example, the second alternate docroot maps this request:

```
http://company22.com/myapp/abc/def/my.jpg
```

To:

```
/srv/images/jpg/abc/def/my.jpg
```

The third alternate docroot maps:

```
http://company22.com/myapp/jpg/abc/resource
```

To:

```
/srv/images/jpg/abc/resource
```

If a request does not match any of the target web application's alternate docroots, or if the target web application does not specify any alternate docroots, the request is served from the web application's standard docroot, as usual.

## Redirecting URLs

You can specify that a request for an old URL is treated as a request for a new URL. This is called *redirecting* a URL.

To specify a redirected URL for a virtual server, use the `redirect_n` property, where *n* is a positive integer that allows specification of more than one. This property is a subelement of a `virtual-server` element in the `domain.xml` file. For more information about this element, see “[virtual-server](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Reference*. Each of these `redirect_n` properties is inherited by all web applications deployed on the virtual server.

The value of each `redirect_n` property has two components, which may be specified in any order:

The first component, `from`, specifies the prefix of the requested URI to match.

The second component, `url-prefix`, specifies the new URL prefix to return to the client. The from prefix is simply replaced by this URL prefix.

For example:

```
<property name="redirect_1" value="from=/dummy url-prefix=http://etude"/>
```

## Enabling Comet Support

To enable Comet support for an HTTP listener and all its associated web applications, set the `cometSupport` property to `true`. For more information, see [“http-listener” in Sun GlassFish Enterprise Server 2.1 Administration Reference](#).

If your servlet or JSP page uses Comet technology, make sure it is initialized when the Enterprise Server starts up by adding the `load-on-startup` element to your `web.xml` file. For example:

```
<servlet>
  <servlet-name>CheckIn</servlet-name>
  <servlet-class>CheckInServlet</servlet-class>
  <load-on-startup>0</load-on-startup>
</servlet>
```

## Using a context.xml File

Use the `contextXmlDefault` property to specify the location, relative to *domain-dir*, of the `context.xml` file for a virtual server. For more information about virtual servers, see [“Virtual Servers” on page 153](#). For more information about the `context.xml` file, see [The Context Container \(http://tomcat.apache.org/tomcat-5.5-doc/config/context.html\)](http://tomcat.apache.org/tomcat-5.5-doc/config/context.html).

## Enabling WebDav

To enable WebDav in the Enterprise Server, you edit the `web.xml` and `sun-web.xml` files as follows.

First, enable the WebDav servlet in your `web.xml` file:

```
<servlet>
  <servlet-name>webdav</servlet-name>
  <servlet-class>org.apache.catalina.servlets.WebdavServlet</servlet-class>
```



```

<init-param>
  <param-name>debug</param-name>
  <param-value>0</param-value>
</init-param>
<init-param>
  <param-name>listings</param-name>
  <param-value>>true</param-value>
</init-param>
<init-param>
  <param-name>readonly</param-name>
  <param-value>>false</param-value>
</init-param>
</servlet>

```

Then define the servlet mapping associated with your WebDav servlet in your `web.xml` file:

```

<servlet-mapping>
  <servlet-name>webdav</servlet-name>
  <url-pattern>/webdav/*</url-pattern>
</servlet-mapping>

```

To protect the WebDav servlet so other users can't modify it, add a security constraint in your `web.xml` file:

```

<security-constraint>
  <web-resource-collection>
    <web-resource-name>Login Resources</web-resource-name>
    <url-pattern>/webdav/*</url-pattern>
  </web-resource-collection>
  <auth-constraint>
    <role-name>Admin</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>NONE</transport-guarantee>
  </user-data-constraint>
  <login-config>
    <auth-method>BASIC</auth-method>
    <realm-name>default</realm-name>
  </login-config>
  <security-role>
    <role-name>Admin</role-name>
  </security-role>
</security-constraint>

```

Then define a security role mapping in your `sun-web.xml` file:

```

<security-role-mapping>
  <role-name>Admin</role-name>

```

```
<group-name>Admin</group-name>
</security-role-mapping>
```

If you are using the file realm, create a user and password. For example:

```
asadmin create-file-user --user admin --host localhost --port 4848 --terse=true
--groups Admin --authrealmname default admin
```

You can now use any WebDav client by connecting to the WebDav servlet URL, which has this format:

```
http://host:port/context-root/webdav/file
```

For example:

```
http://localhost:80/glassfish-webdav/webdav/index.html
```

You can add the WebDav servlet to your default-web.xml file to enable it for all applications, but you can't set up a security role mapping to protect it.

## Using mod\_jk

To set up mod\_jk, follow these steps:

1. Obtain the following software:
  - Apache 2.0.x
  - Apache Tomcat Connectors (<http://www.apache.org/dist/tomcat/tomcat-connectors/jk/binaries/>)
  - Apache Tomcat 5.5.16, needed for just one JAR file (<http://archive.apache.org/dist/tomcat/tomcat-5/v5.5.16/bin/apache-tomcat-5.5.16.tar.gz>)
  - Apache Commons Logging 1.0.4 (<http://archive.apache.org/dist/jakarta/commons/logging/binaries/commons-logging-1.0.4.tar.gz>)
  - Apache Commons Modeler 1.1 (<http://archive.apache.org/dist/jakarta/commons/modeler/binaries/modeler-1.1.tar.gz>)
2. Install mod\_jk as described at [http://tomcat.apache.org/connectors-doc/webserver\\_howto/apache.html](http://tomcat.apache.org/connectors-doc/webserver_howto/apache.html).
3. Copy the following Tomcat and Jakarta Commons files to *as-install/lib*:
  - tomcat-ajp.jar
  - commons-logging.jar
  - commons-modeler.jar
4. Create and configure the following files:
  - /etc/httpd/conf/httpd.conf

- `/etc/httpd/conf/worker.properties` or `domain-dir/config/glassfish-jk.properties` (to use non-default values of attributes described at <http://tomcat.apache.org/tomcat-5.5-doc/config/ajp.html>)

Examples of these files are shown after these steps. If you use both `worker.properties` and `glassfish-jk.properties` files, the file referenced by `httpd.conf`, or referenced by `httpd.conf` first, takes precedence.

5. Start `httpd`.
6. Enable `mod_jk` using the following command:

```
asadmin create-jvm-options -Dcom.sun.enterprise.web.connector.enableJK=8009
```

7. If you are using the `glassfish-jk.properties` file and not referencing it in `httpd.conf`, point to it using the following command:

```
asadmin create-jvm-options
-Dcom.sun.enterprise.web.connector.enableJK.propertyFile=domain-dir/config/glassfish-jk.properties
```

8. Restart the Enterprise Server.

Here is an example `httpd.conf` file:

```
LoadModule jk_module /usr/lib/httpd/modules/mod_jk.so
JkWorkersFile /etc/httpd/conf/worker.properties
# Where to put jk logs
JkLogFile /var/log/httpd/mod_jk.log
# Set the jk log level [debug/error/info]
JkLogLevel debug
# Select the log format
JkLogStampFormat "[%a %b %d %H:%M:%S %Y] "
# JkOptions indicate to send SSL KEY SIZE,
JkOptions +ForwardKeySize +ForwardURICompat -ForwardDirectories
# JkRequestLogFormat set the request format
JkRequestLogFormat "%w %V %T"
# Send all jsp requests to GlassFish
JkMount /*.jsp worker1
# Send all glassfish-test requests to GlassFish
JkMount /glassfish-test/* worker1
```

Here is an example `worker.properties` or `glassfish-jk.properties` file:

```
# Define 1 real worker using ajp13
worker.list=worker1
# Set properties for worker1 (ajp13)
worker.worker1.type=ajp13
worker.worker1.host=localhost.localdomain
worker.worker1.port=8009
worker.worker1.lbfactor=50
```

```
worker.worker1.cachesize=10  
worker.worker1.cache_timeout=600  
worker.worker1.socket_keepalive=1  
worker.worker1.socket_timeout=300
```

# Using Enterprise JavaBeans Technology

---

This chapter describes how Enterprise JavaBeans™ (EJB™) technology is supported in the Sun GlassFish Enterprise Server. This chapter addresses the following topics:

- “Summary of EJB 3.0 Changes” on page 165
- “Value Added Features” on page 166
- “EJB Timer Service” on page 169
- “Using Session Beans” on page 170
- “Using Read-Only Beans” on page 178
- “Using Message-Driven Beans” on page 181
- “Handling Transactions With Enterprise Beans” on page 184

For general information about enterprise beans, see “Part Three: Enterprise Beans” in the *Java EE 5 Tutorial* (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

## Summary of EJB 3.0 Changes

The Enterprise Server supports and is compliant with the Sun Microsystems Enterprise JavaBeans (EJB) architecture as defined by the Enterprise JavaBeans Specification, v3.0, also known as JSR 220 (<http://jcp.org/en/jsr/detail?id=220>).

---

**Note** – The Enterprise Server is backward compatible with 1.1, 2.0, and 2.1 enterprise beans. However, to take advantage of version 3.0 features, you should develop new beans as 3.0 enterprise beans.

---

The main changes in the Enterprise JavaBeans Specification, v3.0 that impact enterprise beans in the Enterprise Server environment are as follows:

- Definition of the Java language metadata *annotations* that can be used to annotate EJB applications. These metadata annotations are targeted at simplifying the developer's task, at reducing the number of program classes and interfaces the developer is required to

implement, at encapsulation of environmental dependencies and JNDI access, and at eliminating the need for the developer to provide an EJB deployment descriptor.

- Elimination of the requirement for home or EJB component interfaces for session beans. The required business interface for a session bean can be a plain Java interface rather than an `EJBObject`, `EJBLocalObject`, or `java.rmi.Remote` interface.
- Elimination of all required interfaces for persistent entities. Specification of Java language metadata annotations and XML deployment descriptor elements for the object/relational mapping of persistent entities. For details about Java Persistence in the Enterprise Server, see [Chapter 7, “Using the Java Persistence API.”](#)

Container-managed persistence (CMP) is still supported for EJB 2.1 beans, for backward compatibility. For details, see [Chapter 10, “Using Container-Managed Persistence.”](#)

## Value Added Features

The Enterprise Server provides a number of value additions that relate to EJB development. These capabilities are discussed in the following sections. References to more in-depth material are included.

- “Read-Only Beans” on page 166
- “The pass-by-reference Element” on page 167
- “Pooling and Caching” on page 167
- “Bean-Level Container-Managed Transaction Timeouts” on page 168
- “Priority Based Scheduling of Remote Bean Invocations” on page 169
- “Immediate Flushing” on page 169

## Read-Only Beans

Another feature that the Enterprise Server provides is the *read-only bean*, an EJB 2.1 entity bean that is never modified by an EJB client. Read-only beans avoid database updates completely.

---

**Note** – Read-only beans are specific to the Enterprise Server and are not part of the Enterprise JavaBeans Specification, v2.1. Use of this feature for an EJB 2.1 bean results in a non-portable application.

To make an EJB 3.0 entity read-only, use `@Column` annotations to mark its columns `insertable=false` and `updatable=false`.

---

A read-only bean can be used to cache a database entry that is frequently accessed but rarely updated (externally by other beans). When the data that is cached by a read-only bean is updated by another bean, the read-only bean can be notified to refresh its cached data.

The Enterprise Server provides a number of ways by which a read-only bean's state can be refreshed. By setting the `refresh-period-in-seconds` element in the `sun-ejb-jar.xml` file and the `trans-attribute` element (or `@TransactionAttribute` annotation) in the `ejb-jar.xml` file, it is easy to configure a read-only bean that is one of the following:

- Always refreshed
- Periodically refreshed
- Never refreshed
- Programmatically refreshed

Read-only beans are best suited for situations where the underlying data never changes, or changes infrequently. For further information and usage guidelines, see [“Using Read-Only Beans” on page 178](#).

## The pass-by-reference Element

The `pass-by-reference` element in the `sun-ejb-jar.xml` file allows you to specify the parameter passing semantics for colocated remote EJB invocations. This is an opportunity to improve performance. However, use of this feature results in non-portable applications. See [“pass-by-reference” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#).

## Pooling and Caching

The EJB container of the Enterprise Server pools anonymous instances (message-driven beans, stateless session beans, and entity beans) to reduce the overhead of creating and destroying objects. The EJB container maintains the free pool for each bean that is deployed. Bean instances in the free pool have no identity (that is, no primary key associated) and are used to serve method calls. The free beans are also used to serve all methods for stateless session beans.

Bean instances in the free pool transition from a Pooled state to a Cached state after `ejbCreate` and the business methods run. The size and behavior of each pool is controlled using pool-related properties in the EJB container or the `sun-ejb-jar.xml` file.

In addition, the Enterprise Server supports a number of tunable parameters that can control the number of “stateful” instances (stateful session beans and entity beans) cached as well as the duration they are cached. Multiple bean instances that refer to the same database row in a table can be cached. The EJB container maintains a cache for each bean that is deployed.

To achieve scalability, the container selectively evicts some bean instances from the cache, usually when cache overflows. These evicted bean instances return to the free bean pool. The size and behavior of each cache can be controlled using the cache-related properties in the EJB container or the `sun-ejb-jar.xml` file.

Pooling and caching parameters for the `sun-ejb-jar.xml` file are described in [“bean-cache” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#).

## Pooling Parameters

One of the most important parameters of Enterprise Server pooling is `steady-pool-size`. When `steady-pool-size` is set to greater than 0, the container not only pre-populates the bean pool with the specified number of beans, but also attempts to ensure that this number of beans is always available in the free pool. This ensures that there are enough beans in the ready to serve state to process user requests.

This parameter does not necessarily guarantee that no more than `steady-pool-size` instances exist at a given time. It only governs the number of instances that are pooled over a long period of time. For example, suppose an idle stateless session container has a fully-populated pool with a `steady-pool-size` of 10. If 20 concurrent requests arrive for the EJB component, the container creates 10 additional instances to satisfy the burst of requests. The advantage of this is that it prevents the container from blocking any of the incoming requests. However, if the activity dies down to 10 or fewer concurrent requests, the additional 10 instances are discarded.

Another parameter, `pool-idle-timeout-in-seconds`, allows the administrator to specify the amount of time a bean instance can be idle in the pool. When `pool-idle-timeout-in-seconds` is set to greater than 0, the container removes or destroys any bean instance that is idle for this specified duration.

## Caching Parameters

Enterprise Server provides a way that completely avoids caching of entity beans, using commit option C. Commit option C is particularly useful if beans are accessed in large number but very rarely reused. For additional information, refer to [“Commit Options” on page 185](#).

The Enterprise Server caches can be either bounded or unbounded. *Bounded caches* have limits on the number of beans that they can hold beyond which beans are passivated. For stateful session beans, there are three ways (LRU, NRU and FIFO) of picking victim beans when cache overflow occurs. Caches can also passivate beans that are idle (not accessed for a specified duration).

## Bean-Level Container-Managed Transaction Timeouts

The default transaction timeout for the domain is specified using the Transaction Timeout setting of the Transaction Service. A transaction started by the container must commit (or rollback) within this time, regardless of whether the transaction is suspended (and resumed), or the transaction is marked for rollback.

To override this timeout for an individual bean, use the optional `cmt-timeout-in-seconds` element in `sun-ejb-jar.xml`. The default value, 0, specifies that the default Transaction Service timeout is used. The value of `cmt-timeout-in-seconds` is used for all methods in the bean that start a new container-managed transaction. This value is *not* used if the bean joins a client transaction.



## Priority Based Scheduling of Remote Bean Invocations

You can create multiple thread pools, each having its own work queues. An optional element in the `sun-ejb-jar.xml` file, `use-thread-pool-id`, specifies the thread pool that processes the requests for the bean. The bean must have a remote interface, or `use-thread-pool-id` is ignored. You can create different thread pools and specify the appropriate thread pool ID for a bean that requires a quick response time. If there is no such thread pool configured or if the element is absent, the default thread pool is used.

## Immediate Flushing

Normally, all entity bean updates within a transaction are batched and executed at the end of the transaction. The only exception is the database flush that precedes execution of a finder or select query.

Since a transaction often spans many method calls, you might want to find out if the updates made by a method succeeded or failed immediately after method execution. To force a flush at the end of a method's execution, use the `flush-at-end-of-method` element in the `sun-ejb-jar.xml` file. Only non-finder methods in an entity bean can be flush-enabled. (For an EJB 2.1 bean, these methods must be in the Local, Local Home, Remote, or Remote Home interface.) See “flush-at-end-of-method” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

Upon completion of the method, the EJB container updates the database. Any exception thrown by the underlying data store is wrapped as follows:

- If the method that triggered the flush is a create method, the exception is wrapped with `CreateException`.
- If the method that triggered the flush is a remove method, the exception is wrapped with `RemoveException`.
- For all other methods, the exception is wrapped with `EJBException`.

All normal end-of-transaction database synchronization steps occur regardless of whether the database has been flushed during the transaction.

## EJB Timer Service

The EJB Timer Service uses a database to store persistent information about EJB timers. In the developer profile, the EJB Timer Service in Enterprise Server is preconfigured to use an embedded version of the Java DB database. The EJB Timer Service configuration can store persistent timer information in any database supported by the Enterprise Server for persistence.

For a list of the JDBC drivers currently supported by the Enterprise Server, see the [Sun GlassFish Enterprise Server 2.1 Release Notes](#). For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

To change the database used by the EJB Timer Service, set the EJB Timer Service’s Timer DataSource setting to a valid JDBC resource. You must also create the timer database table. DDL files are located in `as-install/lib/install/databases`. Ideally, each cluster should have its own timer table.

Using the EJB Timer Service is equivalent to interacting with a single JDBC resource manager. If an EJB component or application accesses a database either directly through JDBC or indirectly (for example, through an entity bean’s persistence mechanism), and also interacts with the EJB Timer Service, its data source must be configured with an XA JDBC driver.

You can change the following EJB Timer Service settings. You must restart the server for the changes to take effect.

- Minimum Delivery Interval - Specifies the minimum time in milliseconds before an expiration for a particular timer can occur. This guards against extremely small timer increments that can overload the server. The default is 7000.
- Maximum Redeliveries - Specifies the maximum number of times the EJB timer service attempts to redeliver a timer expiration due for exception or rollback. The default is 1.
- Redelivery Interval - Specifies how long in milliseconds the EJB timer service waits after a failed `ejbTimeout` delivery before attempting a redelivery. The default is 5000.
- Timer DataSource - Specifies the database used by the EJB Timer Service. In the developer profile, the default is `jdbc/__TimerPool`.

For information about configuring EJB Timer Service settings, see [Chapter 8, “SIP, Web, and EJB Containers,”](#) in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#). For information about the `asadmin list-timers` and `asadmin migrate-timers` commands, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## Using Session Beans

This section provides guidelines for creating session beans in the Enterprise Server environment. This section addresses the following topics:

- “About the Session Bean Containers” on page 171
- “Stateful Session Bean Failover” on page 172
- “Session Bean Restrictions and Optimizations” on page 177

Extensive information on session beans is contained in Chapters 3 and 4 of the Enterprise JavaBeans Specification, v3.0, EJB Core Contracts and Requirements.

## About the Session Bean Containers

Like an entity bean, a session bean can access a database through Java Database Connectivity (JDBC) calls. A session bean can also provide transaction settings. These transaction settings and JDBC calls are referenced by the session bean's container, allowing it to participate in transactions managed by the container.

A container managing stateless session beans has a different charter from a container managing stateful session beans. This section addresses the following topics:

- “Stateless Container” on page 171
- “Stateful Container” on page 171

### Stateless Container

The *stateless container* manages stateless session beans, which, by definition, do not carry client-specific states. All session beans (of a particular type) are considered equal.

A stateless session bean container uses a bean pool to service requests. The Enterprise Server specific deployment descriptor file, `sun-ejb-jar.xml`, contains the properties that define the pool:

- `steady-pool-size`
- `resize-quantity`
- `max-pool-size`
- `pool-idle-timeout-in-seconds`

For more information about `sun-ejb-jar.xml`, see “The `sun-ejb-jar.xml` File” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

The Enterprise Server provides the `wscmpile` and `wdeploy` tools to help you implement a web service endpoint as a stateless session bean. For more information about these tools, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

### Stateful Container

The *stateful container* manages the stateful session beans, which, by definition, carry the client-specific state. There is a one-to-one relationship between the client and the stateful session beans. At creation, each stateful session bean (SFSB) is given a unique session ID that is used to access the session bean so that an instance of a stateful session bean is accessed by a single client only.

Stateful session beans are managed using cache. The size and behavior of stateful session beans cache are controlled by specifying the following `sun-ejb-jar.xml` parameters:

- `max-cache-size`
- `resize-quantity`

- `cache-idle-timeout-in-seconds`
- `removal-timeout-in-seconds`
- `victim-selection-policy`

The `max-cache-size` element specifies the maximum number of session beans that are held in cache. If the cache overflows (when the number of beans exceeds `max-cache-size`), the container then passivates some beans or writes out the serialized state of the bean into a file. The directory in which the file is created is obtained from the EJB container using the configuration APIs.

For more information about `sun-ejb-jar.xml`, see [“The sun-ejb-jar.xml File” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#).

The passivated beans are stored on the file system. The Session Store Location setting in the EJB container allows the administrator to specify the directory where passivated beans are stored. By default, passivated stateful session beans are stored in application-specific subdirectories created under `domain-dir/session-store`.

---

**Note** – Make sure the `delete` option is set in the `server.policy` file, or expired file-based sessions might not be deleted properly. For more information about `server.policy`, see [“The server.policy File” on page 89](#).

---

The Session Store Location setting also determines where the session state is persisted if it is not highly available; see [“Choosing a Persistence Store” on page 174](#).

## Stateful Session Bean Failover

An SFSB's state can be saved in a persistent store in case a server instance fails. The state of an SFSB is saved to the persistent store at predefined points in its life cycle. This is called *checkpointing*. If SFSB checkpointing is enabled, checkpointing generally occurs after any transaction involving the SFSB is completed, even if the transaction rolls back.

However, if an SFSB participates in a bean-managed transaction, the transaction might be committed in the middle of the execution of a bean method. Since the bean's state might be undergoing transition as a result of the method invocation, this is not an appropriate instant to checkpoint the bean's state. In this case, the EJB container checkpoints the bean's state at the end of the corresponding method, provided the bean is not in the scope of another transaction when that method ends. If a bean-managed transaction spans across multiple methods, checkpointing is delayed until there is no active transaction at the end of a subsequent method.

The state of an SFSB is not necessarily transactional and might be significantly modified as a result of non-transactional business methods. If this is the case for an SFSB, you can specify a list of checkpointed methods. If SFSB checkpointing is enabled, checkpointing occurs after any checkpointed methods are completed.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see [“Usage Profiles” in Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

---

The following table lists the types of references that SFSB failover supports. All objects bound into an SFSB must be one of the supported types. In the table, *No* indicates that failover for the object type might not work in all cases and that no failover support is provided. However, failover might work in some cases for that object type. For example, failover might work because the class implementing that type is serializable.

**TABLE 9-1** Object Types Supported for Java EE Stateful Session Bean State Failover

Java Object Type	Failover Support
Colocated or distributed stateless session, stateful session, or entity bean reference	Yes
JNDI context	Yes, <code>InitialContext</code> and <code>java:comp/env</code>
UserTransaction	Yes, but if the instance that fails is never restarted, any prepared global transactions are lost and might not be correctly rolled back or committed.
JDBC DataSource	No
Java Message Service (JMS) ConnectionFactory, Destination	No
JavaMail Session	No
Connection Factory	No
Administered Object	No
Web service reference	No
Serializable Java types	Yes
Extended persistence context	No

For more information about the `InitialContext`, see [“Accessing the Naming Context” on page 269](#). For more information about transaction recovery, see [Chapter 16, “Using the Transaction Service.”](#) For more information about Administered Objects, see [“Creating Physical Destinations” on page 279](#).

**Note** – Idempotent URLs are supported along the HTTP path, but not the RMI-IIOP path. For more information, see [“Configuring Idempotent URL Requests” on page 156](#).

If a server instance to which an RMI-IIOP client request is sent crashes during the request processing (before the response is prepared and sent back to the client), an error is sent to the client. The client must retry the request explicitly. When the client retries the request, the request is sent to another server instance in the cluster, which retrieves session state information for this client.

HTTP sessions can also be saved in a persistent store in case a server instance fails. In addition, if a distributable web application references an SFSB, and the web application’s session fails over, the EJB reference is also failed over. For more information, see [“Distributed Sessions and Persistence” on page 145](#).

If an SFSB that uses session persistence is undeployed while the Enterprise Server instance is stopped, the session data in the persistence store might not be cleared. To prevent this, undeploy the SFSB while the Enterprise Server instance is running.

---

Configure SFSB failover by:

- [“Choosing a Persistence Store” on page 174](#)
- [“Enabling Checkpointing” on page 175](#)
- [“Specifying Methods to Be Checkpointed” on page 176](#) (optional)

## Choosing a Persistence Store

The following types of persistent storage are supported for passivation and checkpointing of the SFSB state:

- The local file system - Allows a single server instance to recover the SFSB state after a failure and restart. This store also provides passivation and activation of the state to help control the amount of memory used. This option is not supported in a production environment that requires SFSB state persistence. This is the default storage mechanism.
- Other servers - Uses other server instances in the cluster for session persistence. Clustered server instances replicate session state in a ring topology. Each backup instance stores the replicated data in memory.
- The high-availability database (HADB) - Allows a cluster of server instances to recover the SFSB state if any server instance fails. The HADB is also used as the passivation and activation store. Use this option in a production environment that requires SFSB state persistence. For information about how to set up and configure this database, see the description of the `configure-ha-cluster` command in the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

Choose the persistence store in one of the following ways:

- To use the local file system, first disable availability. Select the Availability Service component under the relevant configuration in the Admin Console. Uncheck the Availability Service box. Then select the EJB Container component and edit the Session Store Location value. The default is `domain-dir/session-store`.
- To use other servers, select the Availability Service component under the relevant configuration in the Admin Console. Check the Availability Service box. To enable availability for the EJB container, select the EJB Container Availability tab, then check the Availability Service box. All instances in an Enterprise Server cluster should have the same availability settings to ensure consistent behavior.
- To use the HADB, you must first have the HADB installed and the enterprise profile selected. For information about profiles, see [“Usage Profiles” in Sun GlassFish Enterprise Server 2.1 Administration Guide](#). For information about HADB, see the description of the `configure-ha-cluster` command in the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#). Then you must enable availability as described in the previous item.

For more information, see the [Sun GlassFish Enterprise Server 2.1 High Availability Administration Guide](#).

## Enabling Checkpointing

The following sections describe how to enable SFSB checkpointing:

- [“Server Instance and EJB Container Levels” on page 175](#)
- [“Application and EJB Module Levels” on page 175](#)
- [“SFSB Level” on page 175](#)

### Server Instance and EJB Container Levels

To enable SFSB checkpointing at the server instance or EJB container level, see [“Choosing a Persistence Store” on page 174](#).

### Application and EJB Module Levels

To enable SFSB checkpointing at the application or EJB module level during deployment, use the `asadmin deploy` or `asadmin deploydir` command with the `--availabilityenabled` option set to `true`. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

### SFSB Level

To enable SFSB checkpointing at the SFSB level, set `availability-enabled="true"` in the `ejb` element of the SFSB's `sun-ejb-jar.xml` file as follows:

```
<sun-ejb-jar>
  ...
  <enterprise-beans>
```

```
...
<ejb availability-enabled="true">
  <ejb-name>MySFSB</ejb-name>
</ejb>
...
</enterprise-beans>
</sun-ejb-jar>
```

## Specifying Methods to Be Checkpointed

If SFSB checkpointing is enabled, checkpointing generally occurs after any transaction involving the SFSB is completed, even if the transaction rolls back.

To specify additional optional checkpointing of SFSBs at the end of non-transactional business methods that cause important modifications to the bean's state, use the `checkpoint-at-end-of-method` element within the `ejb` element in `sun-ejb-jar.xml`.

For example:

```
<sun-ejb-jar>
...
<enterprise-beans>
...
<ejb availability-enabled="true">
  <ejb-name>ShoppingCartEJB</ejb-name>
  <checkpoint-at-end-of-method>
    <method>
      <method-name>addToCart</method-name>
    </method>
  </checkpoint-at-end-of-method>
</ejb>
...
</enterprise-beans>
</sun-ejb-jar>
```

For details, see “[checkpoint-at-end-of-method](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

The non-transactional methods in the `checkpoint-at-end-of-method` element can be the following:

- `create()` methods defined in the home or business interface of the SFSB, if you want to checkpoint the initial state of the SFSB immediately after creation
- For SFSBs using container managed transactions only, methods in the remote interface of the bean marked with the transaction attribute `TX_NOT_SUPPORTED` or `TX_NEVER`
- For SFSBs using bean managed transactions only, methods in which a bean managed transaction is neither started nor committed



---

Any other methods mentioned in this list are ignored. At the end of invocation of each of these methods, the EJB container saves the state of the SFSB to persistent store.

---

**Note** – If an SFSB does not participate in any transaction, and if none of its methods are explicitly specified in the `checkpoint-at-end-of-method` element, the bean's state is not checkpointed at all even if `availability-enabled="true"` for this bean.

For better performance, specify a *small* subset of methods. The methods chosen should accomplish a significant amount of work in the context of the Java EE application or should result in some important modification to the bean's state.

---

## Session Bean Restrictions and Optimizations

This section discusses restrictions on developing session beans and provides some optimization guidelines:

- [“Optimizing Session Bean Performance” on page 177](#)
- [“Restricting Transactions” on page 177](#)

### Optimizing Session Bean Performance

For stateful session beans, colocating the stateful beans with their clients so that the client and bean are executing in the same process address space improves performance.

### Restricting Transactions

The following restrictions on transactions are enforced by the container and must be observed as session beans are developed:

- A session bean can participate in, at most, a single transaction at a time.
- If a session bean is participating in a transaction, a client cannot invoke a method on the bean such that the `trans-attribute` element (or `@TransactionAttribute` annotation) in the `ejb-jar.xml` file would cause the container to execute the method in a different or unspecified transaction context or an exception is thrown.
- If a session bean instance is participating in a transaction, a client cannot invoke the `remove` method on the session object's home or business interface object, or an exception is thrown.

## Using Read-Only Beans

A *read-only bean* is an EJB 2.1 entity bean that is never modified by an EJB client. The data that a read-only bean represents can be updated externally by other enterprise beans, or by other means, such as direct database updates.

---

**Note** – Read-only beans are specific to the Enterprise Server and are not part of the Enterprise JavaBeans Specification, v2.1. Use of this feature for an EJB 2.1 bean results in a non-portable application.

To make an EJB 3.0 entity bean read-only, use `@Column` annotations to mark its columns `insertable=false` and `updatable=false`.

---

Read-only beans are best suited for situations where the underlying data never changes, or changes infrequently. The following topics are addressed in this section:

- “Read-Only Bean Characteristics and Life Cycle” on page 178
- “Read-Only Bean Good Practices” on page 179
- “Refreshing Read-Only Beans” on page 179
- “Deploying Read-Only Beans” on page 181

## Read-Only Bean Characteristics and Life Cycle

Read-only beans are best suited for situations where the underlying data never changes, or changes infrequently. For example, a read-only bean can be used to represent a stock quote for a particular company, which is updated externally. In such a case, using a regular entity bean might incur the burden of calling `ejbStore`, which can be avoided by using a read-only bean.

Read-only beans have the following characteristics:

- Only entity beans can be read-only beans.
- Either bean-managed persistence (BMP) or container-managed persistence (CMP) is allowed. If CMP is used, do not create the database schema during deployment. Instead, work with your database administrator to populate the data into the tables. See [Chapter 10](#), “Using Container-Managed Persistence.”
- Only container-managed transactions are allowed; read-only beans cannot start their own transactions.
- Read-only beans don’t update any bean state.
- `ejbStore` is never called by the container.
- `ejbLoad` is called only when a transactional method is called or when the bean is initially created (in the cache), or at regular intervals controlled by the bean’s `refresh-period-in-seconds` element in the `sun-ejb-jar.xml` file.

- The home interface can have any number of find methods. The return type of the find methods must be the primary key for the same bean type (or a collection of primary keys).
- If the data that the bean represents can change, then `refresh-period-in-seconds` must be set to refresh the beans at regular intervals. `ejbLoad` is called at this regular interval.

A read-only bean comes into existence using the appropriate find methods.

Read-only beans are cached and have the same cache properties as entity beans. When a read-only bean is selected as a victim to make room in the cache, `ejbPassivate` is called and the bean is returned to the free pool. When in the free pool, the bean has no identity and is used only to serve any finder requests.

Read-only beans are bound to the naming service like regular read-write entity beans, and clients can look up read-only beans the same way read-write entity beans are looked up.

## Read-Only Bean Good Practices

For best results, follow these guidelines when developing read-only beans:

- Avoid having any `create` or `remove` methods in the home interface.
- Use any of the valid EJB 2.1 transaction attributes for the `trans-attribute` element.  
The reason for having `TX_SUPPORTED` is to allow reading uncommitted data in the same transaction. Also, the transaction attributes can be used to force `ejbLoad`.

## Refreshing Read-Only Beans

There are several ways of refreshing read-only beans as addressed in the following sections:

- [“Invoking a Transactional Method” on page 179](#)
- [“Refreshing Periodically” on page 179](#)
- [“Refreshing Programmatically” on page 180](#)

### Invoking a Transactional Method

Invoking any transactional method invokes `ejbLoad`.

### Refreshing Periodically

Use the `refresh-period-in-seconds` element in the `sun-ejb-jar.xml` file to refresh a read-only bean periodically.

- If the value specified in `refresh-period-in-seconds` is zero or not specified, which is the default, the bean is never refreshed (unless a transactional method is accessed).
- If the value is greater than zero, the bean is refreshed at the rate specified.

---

**Note** – This is the only way to refresh the bean state if the data can be modified external to the Enterprise Server.

---

By default, a single timer is used for all instances of a read-only bean. When that timer fires, all bean instances are marked as expired and are refreshed from the database the next time they are used.

Use the `-Dcom.sun.ejb.containers.readonly.relative.refresh.mode=true` flag to refresh each bean instance independently upon access if its refresh period has expired. The default is `false`. Note that each instance still has the same refresh period. This additional level of granularity can improve the performance of read-only beans that do not need to be refreshed at the same time.

To set this flag, use the `asadmin create-jvm-options` command. For example:

```
asadmin create-jvm-options --user adminuser -Dcom.sun.ejb.containers.readonly.relative.refresh.mode=true
```

## Refreshing Programmatically

Typically, beans that update any data that is cached by read-only beans need to notify the read-only beans to refresh their state. Use `ReadOnlyBeanNotifier` to force the refresh of read-only beans.

To do this, invoke the following methods on the `ReadOnlyBeanNotifier` bean:

```
public interface ReadOnlyBeanNotifier extends java.rmi.Remote {
    refresh(Object PrimaryKey) throws RemoteException;
}
```

The implementation of the `ReadOnlyBeanNotifier` interface is provided by the container. The bean looks up `ReadOnlyBeanNotifier` using a fragment of code such as the following example:

```
com.sun.appserv.ejb.ReadOnlyBeanHelper helper =
    new com.sun.appserv.ejb.ReadOnlyBeanHelper();
com.sun.appserv.ejb.ReadOnlyBeanNotifier notifier =
    helper.getReadOnlyBeanNotifier("java:comp/env/ejb/ReadOnlyCustomer");
notifier.refresh(PrimaryKey);
```

For a local read-only bean notifier, the lookup has this modification:

```
helper.getReadOnlyBeanLocalNotifier("java:comp/env/ejb/LocalReadOnlyCustomer");
```

Beans that update any data that is cached by read-only beans need to call the `refresh` methods. The next (non-transactional) call to the read-only bean invokes `ejbLoad`.

For Javadoc tool pages relevant to read-only beans, go to <http://glassfish.dev.java.net/nonav/javaee5/api/index.html> and click on the `com.sun.appserv.ejb` package.

## Deploying Read-Only Beans

Read-only beans are deployed in the same manner as other entity beans. However, in the entry for the bean in the `sun-ejb-jar.xml` file, the `is-read-only-bean` element must be set to `true`. That is:

```
<is-read-only-bean>true</is-read-only-bean>
```

Also, the `refresh-period-in-seconds` element in the `sun-ejb-jar.xml` file can be set to some value that specifies the rate at which the bean is refreshed. If this element is missing, no refresh occurs.

All requests in the same transaction context are routed to the same read-only bean instance. Set the `allow-concurrent-access` element to either `true` (to allow concurrent accesses) or `false` (to serialize concurrent access to the same read-only bean). The default is `false`.

For further information on these elements, refer to “The `sun-ejb-jar.xml` File” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## Using Message-Driven Beans

This section describes message-driven beans and explains the requirements for creating them in the Enterprise Server environment. This section contains the following topics:

- “Message-Driven Bean Configuration” on page 181
- “Message-Driven Bean Restrictions and Optimizations” on page 183

## Message-Driven Bean Configuration

This section addresses the following configuration topics:

- “Connection Factory and Destination” on page 182
- “Message-Driven Bean Pool” on page 182
- “Domain-Level Settings” on page 182

For information about setting up load balancing for message-driven beans, see “Load-Balanced Message Inflow” on page 282.

## Connection Factory and Destination

A message-driven bean is a client to a Connector inbound resource adapter. The message-driven bean container uses the JMS service integrated into the Enterprise Server for message-driven beans that are JMS clients. JMS clients use JMS Connection Factory- and Destination-administered objects. A JMS Connection Factory administered object is a resource manager Connection Factory object that is used to create connections to the JMS provider.

The `mdb-connection-factory` element in the `sun-ejb-jar.xml` file for a message-driven bean specifies the connection factory that creates the container connection to the JMS provider.

The `jndi-name` element of the `ejb` element in the `sun-ejb-jar.xml` file specifies the JNDI name of the administered object for the JMS Queue or Topic destination that is associated with the message-driven bean.

## Message-Driven Bean Pool

The container manages a pool of message-driven beans for the concurrent processing of a stream of messages. The `sun-ejb-jar.xml` file contains the elements that define the pool (that is, the `bean-pool` element):

- `steady-pool-size`
- `resize-quantity`
- `max-pool-size`
- `pool-idle-timeout-in-seconds`

For more information about `sun-ejb-jar.xml`, see [“The sun-ejb-jar.xml File”](#) in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## Domain-Level Settings

You can control the following domain-level message-driven bean settings in the EJB container:

- **Initial and Minimum Pool Size** - Specifies the initial and minimum number of beans maintained in the pool. The default is 0.
- **Maximum Pool Size** - Specifies the maximum number of beans that can be created to satisfy client requests. The default is 32.
- **Pool Resize Quantity** - Specifies the number of beans to be created if a request arrives when the pool is empty (subject to the Initial and Minimum Pool Size), or the number of beans to remove if idle for more than the Idle Timeout. The default is 8.
- **Idle Timeout** - Specifies the maximum time in seconds that a bean can remain idle in the pool. After this amount of time, the bean is destroyed. The default is 600 (10 minutes). A value of 0 means a bean can remain idle indefinitely.

For information on monitoring message-driven beans, click the Help button in the Admin Console. In the developer profile, the Monitor tab is accessible from the Application Server page. In the Cluster profile, select the Stand-Alone Instances component, select the instance from the table, and select the Monitor tab.

---

**Note** – Running monitoring when it is not needed might impact performance, so you might choose to turn monitoring off when it is not in use. For details, see [Chapter 18, “Monitoring Components and Services,”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

## Message-Driven Bean Restrictions and Optimizations

This section discusses the following restrictions and performance optimizations that pertain to developing message-driven beans:

- “Pool Tuning and Monitoring” on page 183
- “The `onMessage` Runtime Exception” on page 183

### Pool Tuning and Monitoring

The message-driven bean pool is also a pool of threads, with each message-driven bean instance in the pool associating with a server session, and each server session associating with a thread. Therefore, a large pool size also means a high number of threads, which impacts performance and server resources.

When configuring message-driven bean pool properties, make sure to consider factors such as message arrival rate and pattern, `onMessage` method processing time, overall server resources (threads, memory, and so on), and any concurrency requirements and limitations from other resources that the message-driven bean accesses.

When tuning performance and resource usage, make sure to consider potential JMS provider properties for the connection factory used by the container (the `mdb-connection-factory` element in the `sun-ejb-jar.xml` file). For example, you can tune the Sun GlassFish Message Queue flow control related properties for connection factory in situations where the message incoming rate is much higher than `max-pool-size` can handle.

Refer to [Chapter 18, “Monitoring Components and Services,”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide* for information on how to get message-driven bean pool statistics.

### The `onMessage` Runtime Exception

Message-driven beans, like other well-behaved `MessageListeners`, should not, in general, throw runtime exceptions. If a message-driven bean’s `onMessage` method encounters a system-level exception or error that does not allow the method to successfully complete, the Enterprise JavaBeans Specification, v3.0 provides the following guidelines:

- If the bean method encounters a runtime exception or error, it should simply propagate the error from the bean method to the container.
- If the bean method performs an operation that results in a checked exception that the bean method cannot recover, the bean method should throw the `javax.ejb.EJBException` that wraps the original exception.
- Any other unexpected error conditions should be reported using `javax.ejb.EJBException` (`javax.ejb.EJBException` is a subclass of `java.lang.RuntimeException`).

Under container-managed transaction demarcation, upon receiving a runtime exception from a message-driven bean's `onMessage` method, the container rolls back the container-started transaction and the message is redelivered. This is because the message delivery itself is part of the container-started transaction. By default, the Enterprise Server container closes the container's connection to the JMS provider when the first runtime exception is received from a message-driven bean instance's `onMessage` method. This avoids potential message redelivery looping and protects server resources if the message-driven bean's `onMessage` method continues misbehaving. To change this default container behavior, use the `cmt-max-runtime-exceptions` property of the `mdb-container` element in the `domain.xml` file.

The `cmt-max-runtime-exceptions` property specifies the maximum number of runtime exceptions allowed from a message-driven bean's `onMessage` method before the container starts to close the container's connection to the message source. By default this value is 1; -1 disables this container protection.

A message-driven bean's `onMessage` method can use the `javax.jms.Message` `getJMSRedelivered` method to check whether a received message is a redelivered message.

---

**Note** – The `cmt-max-runtime-exceptions` property might be deprecated in the future.

---

## Handling Transactions With Enterprise Beans

This section describes the transaction support built into the Enterprise JavaBeans programming model for the Enterprise Server.

As a developer, you can write an application that updates data in multiple databases distributed across multiple sites. The site might use EJB servers from different vendors. This section provides overview information on the following topics:

- [“Flat Transactions” on page 185](#)
- [“Global and Local Transactions” on page 185](#)
- [“Commit Options” on page 185](#)
- [“Administration and Monitoring” on page 186](#)



## Flat Transactions

The Enterprise JavaBeans Specification, v3.0 requires support for flat (as opposed to nested) transactions. In a flat transaction, each transaction is decoupled from and independent of other transactions in the system. Another transaction cannot start in the same thread until the current transaction ends.

Flat transactions are the most prevalent model and are supported by most commercial database systems. Although nested transactions offer a finer granularity of control over transactions, they are supported by far fewer commercial database systems.

## Global and Local Transactions

Understanding the distinction between global and local transactions is crucial in understanding the Enterprise Server support for transactions. See [“Transaction Scope” on page 262](#).

Both local and global transactions are demarcated using the `javax.transaction.UserTransaction` interface, which the client must use. Local transactions bypass the transaction manager and are faster. For more information, see [“The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction” on page 264](#).

## Commit Options

The EJB protocol is designed to give the container the flexibility to select the disposition of the instance state at the time a transaction is committed. This allows the container to best manage caching an entity object’s state and associating an entity object identity with the EJB instances.

There are three commit-time options:

- **Option A** – The container caches a ready instance between transactions. The container ensures that the instance has exclusive access to the state of the object in persistent storage. In this case, the container does *not* have to synchronize the instance’s state from the persistent storage at the beginning of the next transaction.

---

**Note** – Commit option A is not supported for this Enterprise Server release.

---

- **Option B** – The container caches a ready instance between transactions, but the container does *not* ensure that the instance has exclusive access to the state of the object in persistent storage. This is the default.

In this case, the container must synchronize the instance’s state by invoking `ejbLoad` from persistent storage at the beginning of the next transaction.

- **Option C** – The container does *not* cache a ready instance between transactions, but instead returns the instance to the pool of available instances after a transaction has completed.

The life cycle for every business method invocation under commit option C looks like this.

```
ejbActivate → ejbLoad → business method → ejbStore → ejbPassivate
```

If there is more than one transactional client concurrently accessing the same entity, the first client gets the ready instance and subsequent concurrent clients get new instances from the pool.

The Enterprise Server deployment descriptor has an element, `commit-option`, that specifies the commit option to be used. Based on the specified commit option, the appropriate handler is instantiated.

## Administration and Monitoring

An administrator can control a number of domain-level Transaction Service settings. For details, see [“Configuring the Transaction Service” on page 264](#).

The Transaction Timeout setting can be overridden by a bean. See [“Bean-Level Container-Managed Transaction Timeouts” on page 168](#).

In addition, the administrator can monitor transactions using statistics from the transaction manager that provide information on such activities as the number of transactions completed, rolled back, or recovered since server startup, and transactions presently being processed.

For information on administering and monitoring transactions, select the Transaction Service component under the relevant configuration in the Admin Console and click the Help button. Also see [Chapter 12, “Transactions,” in \*Sun GlassFish Enterprise Server 2.1 Administration Guide\*](#).

# Using Container-Managed Persistence

---

This chapter contains information on how EJB 2.1 container-managed persistence (CMP) works in the Sun GlassFish Enterprise Server in the following topics:

- “Enterprise Server Support for CMP” on page 187
- “CMP Mapping” on page 188
- “Automatic Schema Generation for CMP” on page 192
- “Schema Capture” on page 198
- “Configuring the CMP Resource” on page 199
- “Performance-Related Features” on page 199
- “Configuring Queries for 1.1 Finders” on page 202
- “CMP Restrictions and Optimizations” on page 206

## Enterprise Server Support for CMP

Enterprise Server support for EJB 2.1 CMP beans includes:

- Full support for the J2EE v1.4 specification’s CMP model. Extensive information on CMP is contained in chapters 10, 11, and 14 of the Enterprise JavaBeans Specification, v2.1. This includes the following:
  - Support for commit options B and C for transactions. See “Commit Options” on page 185.
  - The primary key class must be a subclass of `java.lang.Object`. This ensures portability, and is noted because some vendors allow primitive types (such as `int`) to be used as the primary key class.
- The Enterprise Server CMP implementation, which provides the following:
  - An Object/Relational (O/R) mapping tool that creates XML deployment descriptors for EJB JAR files that contain beans that use CMP.
  - Support for compound (multi-column) primary keys.
  - Support for sophisticated custom finder methods.

- Standards-based query language (EJB QL).
- CMP runtime support. See [“Configuring the CMP Resource” on page 199](#).
- Enterprise Server performance-related features, including the following:
  - Version column consistency checking
  - Relationship prefetching
  - Read-Only Beans

For details, see [“Performance-Related Features” on page 199](#).

## CMP Mapping

Implementation for entity beans that use CMP is mostly a matter of mapping CMP fields and CMR fields (relationships) to the database. This section addresses the following topics:

- [“Mapping Capabilities” on page 188](#)
- [“The Mapping Deployment Descriptor File” on page 188](#)
- [“Mapping Considerations” on page 189](#)

### Mapping Capabilities

*Mapping* refers to the ability to tie an object-based model to a relational model of data, usually the schema of a relational database. The CMP implementation provides the ability to tie a set of interrelated beans containing data and associated behaviors to the schema. This object representation of the database becomes part of the Java application. You can also customize this mapping to optimize these beans for the particular needs of an application. The result is a single data model through which both persistent database information and regular transient program data are accessed.

The mapping capabilities provided by the Enterprise Server include:

- Mapping a CMP bean to one or more tables
- Mapping CMP fields to one or more columns
- Mapping CMP fields to different column types
- Mapping tables with compound primary keys
- Mapping tables with unknown primary keys
- Mapping CMP relationships to foreign keys
- Mapping tables with overlapping primary and foreign keys

### The Mapping Deployment Descriptor File

Each module with CMP beans must have the following files:

- `ejb-jar.xml` – The J2EE standard file for assembling enterprise beans. For a detailed description, see the Enterprise JavaBeans Specification, v2.1.
- `sun-ejb-jar.xml` – The Enterprise Server standard file for assembling enterprise beans. For a detailed description, see “[The sun-ejb-jar.xml File](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.
- `sun-cmp-mappings.xml` – The *mapping deployment descriptor file*, which describes the mapping of CMP beans to tables in a database. For a detailed description, see “[The sun-cmp-mappings.xml File](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

The `sun-cmp-mappings.xml` file can be automatically generated and does not have to exist prior to deployment. For details, see “[Generation Options for CMP](#)” on page 195.

The `sun-cmp-mappings.xml` file maps CMP fields and CMR fields (relationships) to the database. A primary table must be selected for each CMP bean, and optionally, multiple secondary tables. CMP fields are mapped to columns in either the primary or secondary table(s). CMR fields are mapped to pairs of column lists (normally, column lists are the lists of columns associated with primary and foreign keys).

---

**Note** – Table names in databases can be case-sensitive. Make sure that the table names in the `sun-cmp-mappings.xml` file match the names in the database.

Relationships should always be mapped to the primary key field(s) of the related table.

---

The `sun-cmp-mappings.xml` file conforms to the `sun-cmp-mapping_1_2.dtd` file and is packaged with the user-defined bean classes in the EJB JAR file under the META-INF directory.

The Enterprise Server creates the mappings in the `sun-cmp-mappings.xml` file automatically during deployment if the file is not present.

To map the fields and relationships of your entity beans manually, edit the `sun-cmp-mappings.xml` deployment descriptor. Only do this if you are proficient in editing XML.

The mapping information is developed in conjunction with the database schema (`.dbschema`) file, which can be automatically captured when you deploy the bean (see “[Automatic Database Schema Capture](#)” on page 198). You can manually generate the schema using the `capture-schema` utility (“[Using the capture-schema Utility](#)” on page 198).

## Mapping Considerations

This section addresses the following topics:

- “[Join Tables and Relationships](#)” on page 190
- “[Automatic Primary Key Generation](#)” on page 190

- “Fixed Length CHAR Primary Keys” on page 190
- “Managed Fields” on page 190
- “BLOB Support” on page 191
- “CLOB Support” on page 192

The data types used in automatic schema generation are also suggested for manual mapping. These data types are described in [“Supported Data Types for CMP” on page 193](#).

## Join Tables and Relationships

Use of join tables in the database schema is supported for all types of relationships, not just many-to-many relationships. For general information about relationships, see section 10.3.7 of the Enterprise JavaBeans Specification, v2.1.

## Automatic Primary Key Generation

The Enterprise Server supports automatic primary key generation for EJB 1.1, 2.0, and 2.1 CMP beans. To specify automatic primary key generation, give the `prim-key-class` element in the `ejb-jar.xml` file the value `java.lang.Object`. CMP beans with automatically generated primary keys can participate in relationships with other CMP beans. The Enterprise Server does not support database-generated primary key values.

If the database schema is created during deployment, the Enterprise Server creates the schema with the primary key column, then generates unique values for the primary key column at runtime.

If the database schema is not created during deployment, the primary key column in the mapped table must be of type `NUMERIC` with a precision of 19 or more, and must not be mapped to any CMP field. The Enterprise Server generates unique values for the primary key column at runtime.

## Fixed Length CHAR Primary Keys

If an existing database table has a primary key column in which the values vary in length, but the type is `CHAR` instead of `VARCHAR`, the Enterprise Server automatically trims any extra spaces when retrieving primary key values. It is not a good practice to use a fixed length `CHAR` column as a primary key. Use this feature with schemas that cannot be changed, such as a schema inherited from a legacy application.

## Managed Fields

A managed field is a CMP or CMR field that is mapped to the same database column as another CMP or CMR field. CMP fields mapped to the same column and CMR fields mapped to exactly the same column lists always have the same value in memory. For CMR fields that share only a subset of their mapped columns, changes to the columns affect the relationship fields in memory differently. Basically, the Enterprise Server always tries to keep the state of the objects in memory synchronized with the database.

A managed field can have any fetched-with subelement. If the fetched-with subelement is <default/>, the `-DAllowManagedFieldsInDefaultFetchGroup` flag must be set to `true`. See “Default Fetch Group Flags” on page 202 and “fetched-with” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## BLOB Support

Binary Large Object (BLOB) is a data type used to store values that do not correspond to other types such as numbers, strings, or dates. Java fields whose types implement `java.io.Serializable` or are represented as `byte[]` can be stored as BLOBs.

If a CMP field is defined as `Serializable`, it is serialized into a `byte[]` before being stored in the database. Similarly, the value fetched from the database is deserialized. However, if a CMP field is defined as `byte[]`, it is stored directly instead of being serialized and deserialized when stored and fetched, respectively.

To enable BLOB support in the Enterprise Server environment, define a CMP field of type `byte[]` or a user-defined type that implements the `java.io.Serializable` interface. If you map the CMP bean to an existing database schema, map the field to a column of type BLOB.

To use BLOB or CLOB data types larger than 4 KB for CMP using the Inet Oraxo JDBC Driver for Oracle Databases, you must set the `streamsToLob` property value to `true`.

For a list of the JDBC drivers currently supported by the Enterprise Server, see the *Sun GlassFish Enterprise Server 2.1 Release Notes*. For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

For automatic mapping, you might need to change the default BLOB column length for the generated schema using the `schema-generator-properties` element in `sun-ejb-jar.xml`. See your database vendor documentation to determine whether you need to specify the length. For example:

```
<schema-generator-properties>
  <property>
    <name>Employee.voiceGreeting.jdbc-type</name>
    <value>BLOB</value>
  </property>
  <property>
    <name>Employee.voiceGreeting.jdbc-maximum-length</name>
    <value>10240</value>
  </property>
  ...
</schema-generator-properties>
```

## CLOB Support

Character Large Object (CLOB) is a data type used to store and retrieve very long text fields. CLOBs translate into long strings.

To enable CLOB support in the Enterprise Server environment, define a CMP field of type `java.lang.String`. If you map the CMP bean to an existing database schema, map the field to a column of type CLOB.

To use BLOB or CLOB data types larger than 4 KB for CMP using the Inet Oraxo JDBC Driver for Oracle Databases, you must set the `streamsToLob` property value to `true`.

For a list of the JDBC drivers currently supported by the Enterprise Server, see the [Sun GlassFish Enterprise Server 2.1 Release Notes](#). For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

For automatic mapping, you might need to change the default CLOB column length for the generated schema using the `schema-generator-properties` element in `sun-ejb-jar.xml`. See your database vendor documentation to determine whether you need to specify the length. For example:

```
<schema-generator-properties>
  <property>
    <name>Employee.resume.jdbc-type</name>
    <value>CLOB</value>
  </property>
  <property>
    <name>Employee.resume.jdbc-maximum-length</name>
    <value>10240</value>
  </property>
  ...
</schema-generator-properties>
```

## Automatic Schema Generation for CMP

The automatic schema generation feature provided in the Enterprise Server defines database tables based on the fields in entity beans and the relationships between the fields. This insulates developers from many of the database related aspects of development, allowing them to focus on entity bean development. The resulting schema is usable as-is or can be given to a database administrator for tuning with respect to performance, security, and so on.

This section addresses the following topics:

- “Supported Data Types for CMP” on page 193
- “Generation Options for CMP” on page 195



**Note** – Automatic schema generation is supported on an all-or-none basis: it expects that no tables exist in the database before it is executed. It is not intended to be used as a tool to generate extra tables or constraints.

Deployment won't fail if all tables are not created, and undeployment won't fail if not all tables are dropped. This is done to allow you to investigate the problem and fix it manually. You should not rely on the partially created database schema to be correct for running the application.

## Supported Data Types for CMP

CMP supports a set of JDBC data types that are used in mapping Java data fields to SQL types. Supported JDBC data types are as follows: BIGINT, BIT, BLOB, CHAR, CLOB, DATE, DECIMAL, DOUBLE, FLOAT, INTEGER, NUMERIC, REAL, SMALLINT, TIME, TIMESTAMP, TINYINT, VARCHAR.

The following table contains the mappings of Java types to JDBC types when automatic mapping is used.

TABLE 10-1 Java Type to JDBC Type Mappings for CMP

Java Type	JDBC Type	Nullability
boolean	BIT	No
java.lang.Boolean	BIT	Yes
byte	TINYINT	No
java.lang.Byte	TINYINT	Yes
double	DOUBLE	No
java.lang.Double	DOUBLE	Yes
float	REAL	No
java.lang.Float	REAL	Yes
int	INTEGER	No
java.lang.Integer	INTEGER	Yes
long	BIGINT	No
java.lang.Long	BIGINT	Yes
short	SMALLINT	No

TABLE 10-1 Java Type to JDBC Type Mappings for CMP (Continued)

Java Type	JDBC Type	Nullability
<code>java.lang.Short</code>	SMALLINT	Yes
<code>java.math.BigDecimal</code>	DECIMAL	Yes
<code>java.math.BigInteger</code>	DECIMAL	Yes
<code>char</code>	CHAR	No
<code>java.lang.Character</code>	CHAR	Yes
<code>java.lang.String</code>	VARCHAR or CLOB	Yes
<code>Serializable</code>	BLOB	Yes
<code>byte[]</code>	BLOB	Yes
<code>java.util.Date</code>	DATE (Oracle only) TIMESTAMP (all other databases)	Yes
<code>java.sql.Date</code>	DATE	Yes
<code>java.sql.Time</code>	TIME	Yes
<code>java.sql.Timestamp</code>	TIMESTAMP	Yes

**Note** – Java types assigned to CMP fields must be restricted to Java primitive types, Java `Serializable` types, `java.util.Date`, `java.sql.Date`, `java.sql.Time`, or `java.sql.Timestamp`. An entity bean local interface type (or a collection of such) can be the type of a CMR field.

The following table contains the mappings of JDBC types to database vendor-specific types when automatic mapping is used. For a list of the JDBC drivers currently supported by the Enterprise Server, see the [Sun GlassFish Enterprise Server 2.1 Release Notes](#). For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

TABLE 10-2 Mappings of JDBC Types to Database Vendor Specific Types for CMP

JDBC Type	Java DB, Derby, CloudScape	Oracle	DB2	Sybase ASE 12.5	MS-SQL Server
BIT	SMALLINT	SMALLINT	SMALLINT	TINYINT	BIT
TINYINT	SMALLINT	SMALLINT	SMALLINT	TINYINT	TINYINT
SMALLINT	SMALLINT	SMALLINT	SMALLINT	SMALLINT	SMALLINT

TABLE 10-2 Mappings of JDBC Types to Database Vendor Specific Types for CMP (Continued)

JDBC Type	Java DB, Derby, CloudScape	Oracle	DB2	Sybase ASE 12.5	MS-SQL Server
INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
BIGINT	BIGINT	NUMBER	BIGINT	NUMERIC	NUMERIC
REAL	REAL	REAL	FLOAT	FLOAT	REAL
DOUBLE	DOUBLE PRECISION	DOUBLE PRECISION	DOUBLE	DOUBLE PRECISION	FLOAT
DECIMAL(p, s)	DECIMAL(p, s)	NUMBER(p, s)	DECIMAL(p, s)	DECIMAL(p, s)	DECIMAL(p, s)
VARCHAR	VARCHAR	VARCHAR2	VARCHAR	VARCHAR	VARCHAR
DATE	DATE	DATE	DATE	DATETIME	DATETIME
TIME	TIME	DATE	TIME	DATETIME	DATETIME
TIMESTAMP	TIMESTAMP	TIMESTAMP(9)	TIMESTAMP	DATETIME	DATETIME
BLOB	BLOB	BLOB	BLOB	IMAGE	IMAGE
CLOB	CLOB	CLOB	CLOB	TEXT	NTEXT

## Generation Options for CMP

Deployment descriptor elements or `asadmin` command line options can control automatic schema generation by the following:

- Creating tables during deployment
- Dropping tables during undeployment
- Dropping and creating tables during redeployment
- Specifying the database vendor
- Specifying that table names are unique
- Specifying type mappings for individual CMP fields

---

**Note** – Before using these options, make sure you have a properly configured CMP resource. See [“Configuring the CMP Resource” on page 199](#).

For a read-only bean, do not create the database schema during deployment. Instead, work with your database administrator to populate the data into the tables. See [“Using Read-Only Beans” on page 178](#).

Automatic schema generation is not supported for beans with version column consistency checking. Instead, work with your database administrator to create the schema and add the required triggers. See [“Version Column Consistency Checking” on page 200](#).

---

The following optional data subelements of the `cmp-resource` element in the `sun-ejb-jar.xml` file control the automatic creation of database tables at deployment. For more information about the `cmp-resource` element, see [“`cmp-resource`” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#) and [“Configuring the CMP Resource” on page 199](#).

TABLE 10-3 The `sun-ejb-jar.xml` Generation Elements

Element	Default	Description
<code>create-tables-at-deploy</code>	false	If true, causes database tables to be created for beans that are automatically mapped by the EJB container. If false, does not create tables.
<code>drop-tables-at-undeploy</code>	false	If true, causes database tables that were automatically created when the bean(s) were last deployed to be dropped when the bean(s) are undeployed. If false, does not drop tables.
<code>database-vendor-name</code>	none	Specifies the name of the database vendor for which tables are created. Allowed values are <code>javadb</code> , <code>db2</code> , <code>mssql</code> , <code>oracle</code> , <code>postgresql</code> , <code>pointbase</code> , <code>derby</code> (also for CloudScape), and <code>sybase</code> , case-insensitive.  If no value is specified, a connection is made to the resource specified by the <code>jndi-name</code> subelement of the <code>cmp-resource</code> element in the <code>sun-ejb-jar.xml</code> file, and the database vendor name is read. If the connection cannot be established, or if the value is not recognized, SQL-92 compliance is presumed.
<code>schema-generator-properties</code>	none	Specifies field-specific column attributes in property subelements. Each property name is of the following format:  <i>bean-name.field-name.attribute</i>  For example:  <code>Employee.firstName.jdbc-type</code>  Also allows you to set the <code>use-unique-table-names</code> property. If true, this property specifies that generated table names are unique within each application server domain. The default is false.  For further information and an example, see <a href="#">“<code>schema-generator-properties</code>” in <i>Sun GlassFish Enterprise Server 2.1 Application Deployment Guide</i></a> .

The following options of the `asadmin deploy` or `asadmin deploydir` command control the automatic creation of database tables at deployment.

TABLE 10-4 The `asadmin deploy` and `asadmin deploydir` Generation Options for CMP

Option	Default	Description
<code>--createtables</code>	none	If true, causes database tables to be created for beans that need them. If false, does not create tables. If not specified, the value of the <code>create-tables-at-deploy</code> attribute in <code>sun-ejb-jar.xml</code> is used.

TABLE 10-4 The `asadmin deploy` and `asadmin deploydir` Generation Options for CMP (Continued)

Option	Default	Description
<code>--dropandcreatetables</code>	none	<p>If <code>true</code>, and if tables were automatically created when this application was last deployed, tables from the earlier deployment are dropped and fresh ones are created.</p> <p>If <code>true</code>, and if tables were <i>not</i> automatically created when this application was last deployed, no attempt is made to drop any tables. If tables with the same names as those that would have been automatically created are found, the deployment proceeds, but a warning indicates that tables could not be created.</p> <p>If <code>false</code>, settings of <code>create-tables-at-deploy</code> or <code>drop-tables-at-undeploy</code> in the <code>sun-ejb-jar.xml</code> file are overridden.</p>
<code>--uniquetablename</code>	none	<p>If <code>true</code>, specifies that table names are unique within each application server domain. If not specified, the value of the <code>use-unique-table-names</code> property in <code>sun-ejb-jar.xml</code> is used.</p>
<code>--dbvendorname</code>	none	<p>Specifies the name of the database vendor for which tables are created. Allowed values are <code>javadb</code>, <code>db2</code>, <code>mssql</code>, <code>oracle</code>, <code>postgresql</code>, <code>pointbase</code>, <code>derby</code> (also for CloudScape), and <code>sybase</code>, case-insensitive.</p> <p>If not specified, the value of the <code>database-vendor-name</code> attribute in <code>sun-ejb-jar.xml</code> is used.</p> <p>If no value is specified, a connection is made to the resource specified by the <code>jndi-name</code> subelement of the <code>cmp-resource</code> element in the <code>sun-ejb-jar.xml</code> file, and the database vendor name is read. If the connection cannot be established, or if the value is not recognized, SQL-92 compliance is presumed.</p>

If one or more of the beans in the module are manually mapped and you use any of the `asadmin deploy` or `asadmin deploydir` options, the deployment is not harmed in any way, but the options have no effect, and a warning is written to the server log.

The following options of the `asadmin undeploy` command control the automatic removal of database tables at undeployment.

TABLE 10-5 The `asadmin undeploy` Generation Options for CMP

Option	Default	Description
<code>--droptables</code>	none	<p>If <code>true</code>, causes database tables that were automatically created when the bean(s) were last deployed to be dropped when the bean(s) are undeployed. If <code>false</code>, does not drop tables.</p> <p>If not specified, the value of the <code>drop-tables-at-undeploy</code> attribute in <code>sun-ejb-jar.xml</code> is used.</p>

For more information about the `asadmin deploy`, `asadmin deploydir`, and `asadmin undeploy` commands, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

When command line and `sun-ejb-jar.xml` options are both specified, the `asadmin` options take precedence.

The `asant` tasks `sun-appserv-deploy` and `sun-appserv-undeploy` are equivalent to `asadmin deploy` and `asadmin undeploy`, respectively. These `asant` tasks also override the `sun-ejb-jar.xml` options. For details, see [Chapter 3](#), “The `asant` Utility.”

## Schema Capture

This section addresses the following topics:

- “Automatic Database Schema Capture” on page 198
- “Using the `capture-schema` Utility” on page 198

### Automatic Database Schema Capture

You can configure a CMP bean in Enterprise Server to automatically capture the database metadata and save it in a `.dbschema` file during deployment. If the `sun-cmp-mappings.xml` file contains an empty `<schema/>` entry, the `cmp-resource` entry in the `sun-ejb-jar.xml` file is used to get a connection to the database, and automatic generation of the schema is performed.

---

**Note** – Before capturing the database schema automatically, make sure you have a properly configured CMP resource. See “[Configuring the CMP Resource](#)” on page 199.

---

### Using the `capture-schema` Utility

You can use the `capture-schema` command to manually generate the database metadata (`.dbschema`) file. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

The `capture-schema` utility does *not* modify the schema in any way. Its only purpose is to provide the persistence engine with information about the structure of the database (the schema).

Keep the following in mind when using the `capture-schema` command:

- The name of a `.dbschema` file must be unique across all deployed modules in a domain.
- If more than one schema is accessible for the schema user, more than one table with the same name might be captured if the `-schemaname` parameter of `capture-schema` is not set.
- The schema name must be upper case.
- Table names in databases are case-sensitive. Make sure that the table name matches the name in the database.

- PostgreSQL databases internally convert all names to lower case. Before running the `capture-schema` command on a PostgreSQL database, make sure table and column names are lower case in the `sun-cmp-mappings.xml` file.
- An Oracle database user running the `capture-schema` command needs `ANALYZE ANY TABLE` privileges if that user does not own the schema. These privileges are granted to the user by the database administrator.

## Configuring the CMP Resource

An EJB module that contains CMP beans requires the JNDI name of a JDBC resource in the `jndi-name` subelement of the `cmp-resource` element in the `sun-ejb-jar.xml` file. Set `PersistenceManagerFactory` properties as properties of the `cmp-resource` element in the `sun-ejb-jar.xml` file. See “[cmp-resource](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

In the Admin Console, open the Resources component, then select JDBC. Click the Help button in the Admin Console for information on creating a new JDBC resource.

For a list of the JDBC drivers currently supported by the Enterprise Server, see the *Sun GlassFish Enterprise Server 2.1 Release Notes*. For configurations of supported and other drivers, see “[Configurations for Specific JDBC Drivers](#)” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

For example, if the JDBC resource has the JNDI name `jdbc/MyDatabase`, set the CMP resource in the `sun-ejb-jar.xml` file as follows:

```
<cmp-resource>
  <jndi-name>jdbc/MyDatabase</jndi-name>
</cmp-resource>
```

## Performance-Related Features

The Enterprise Server provides the following features to enhance performance or allow more fine-grained data checking. These features are supported only for entity beans with container managed persistence.

- “[Version Column Consistency Checking](#)” on page 200
- “[Relationship Prefetching](#)” on page 200
- “[Read-Only Beans](#)” on page 201
- “[Default Fetch Group Flags](#)” on page 202

---

**Note** – Use of any of these features results in a non-portable application.

---

## Version Column Consistency Checking

The version consistency feature saves the bean state at first transactional access and caches it between transactions. The state is copied from the cache instead of being read from the database. The bean state is verified by primary key and version column values at flush for custom queries (for dirty instances only) and at commit (for clean and dirty instances).

### ▼ To Use Version Consistency

- 1 **Create the version column in the primary table.**
- 2 **Give the version column a numeric data type.**
- 3 **Provide appropriate update triggers on the version column.**

These triggers must increment the version column on each update of the specified row.

- 4 **Specify the version column.**

This is specified in the `check-version-of-accessed-instances` subelement of the `consistency` element in the `sun-cmp-mappings.xml` file. See “[consistency](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

- 5 **Map the CMP bean to an existing schema.**

Automatic schema generation is not supported for beans with version column consistency checking. Instead, work with your database administrator to create the schema and add the required triggers.

## Relationship Prefetching

In many cases when an entity bean’s state is fetched from the database, its relationship fields are always accessed in the same transaction. Relationship prefetching saves database round trips by fetching data for an entity bean and those beans referenced by its CMR fields in a single database round trip.

To enable relationship prefetching for a CMR field, use the `default` subelement of the `fetch-with` element in the `sun-cmp-mappings.xml` file. By default, these CMR fields are prefetched whenever `findByPrimaryKey` or a custom finder is executed for the entity, or when the entity is navigated to from a relationship. (Recursive prefetching is not supported, because it does not usually enhance performance.) See “[fetch-with](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.



To disable prefetching for specific custom finders, use the `prefetch-disabled` element in the `sun-ejb-jar.xml` file. See “[prefetch-disabled](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

Multilevel relationship prefetching is supported for CMP 2.1 entity beans. To enable multilevel relationship prefetching, set the following property using the `asadmin create-jvm-options` command:

```
asadmin create-jvm-options -Dcom.sun.jdo.spi.persistence.support.sqlstore.MULTILEVEL_PREFETCH=true
```

## Read-Only Beans

Another feature that the Enterprise Server provides is the *read-only bean*, an entity bean that is never modified by an EJB client. Read-only beans avoid database updates completely.

---

**Note** – Read-only beans are specific to the Enterprise Server and are not part of the Enterprise JavaBeans Specification, v2.1. Use of this feature for an EJB 2.1 bean results in a non-portable application.

---

A read-only bean can be used to cache a database entry that is frequently accessed but rarely updated (externally by other beans). When the data that is cached by a read-only bean is updated by another bean, the read-only bean can be notified to refresh its cached data.

The Enterprise Server provides a number of ways by which a read-only bean’s state can be refreshed. By setting the `refresh-period-in-seconds` element in the `sun-ejb-jar.xml` file and the `trans-attribute` element (or `@TransactionAttribute` annotation) in the `ejb-jar.xml` file, it is easy to configure a read-only bean that is one of the following:

- Always refreshed
- Periodically refreshed
- Never refreshed
- Programmatically refreshed

Access to CMR fields of read-only beans is not supported. Deployment will succeed, but an exception will be thrown at runtime if a `get` or `set` method is invoked.

Read-only beans are best suited for situations where the underlying data never changes, or changes infrequently. For further information and usage guidelines, see “[Using Read-Only Beans](#)” on page 178.

## Default Fetch Group Flags

Using the following flags can improve performance.

Setting `-DAllowManagedFieldsInDefaultFetchGroup=true` allows CMP fields that by default cannot be placed into the default fetch group to be loaded along with all other fields that are fetched when the CMP state is loaded into memory. These could be multiple fields mapped to the same column in the database table, for example, an instance field and a CMR. By default this flag is set to `false`.

For additional information, see “level” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

Setting `-DAllowMediatedWriteInDefaultFetchGroup` specifies how updated CMP fields are written back to the database. If the flag is `false`, all fields in the CMP bean are written back to the database if at least one field in the default fetch group has been changed in a transaction. If the flag is `true`, only fields modified by the bean are written back to the database. Specifying `true` can improve performance, particularly on database tables with many columns that have not been updated. By default this flag is set to `false`.

To set one of these flags, use the `asadmin create-jvm-options` command. For example:

```
asadmin create-jvm-options --user adminuser -DAllowManagedFieldsInDefaultFetchGroup=true
```

## Configuring Queries for 1.1 Finders

This section contains the following topics:

- “About JDOQL Queries” on page 202
- “Query Filter Expression” on page 203
- “Query Parameters” on page 204
- “Query Variables” on page 204
- “JDOQL Examples” on page 205

### About JDOQL Queries

The Enterprise JavaBeans Specification, v1.1 does not specify the format of the finder method description. The Enterprise Server uses an extension of Java Data Objects Query Language (JDOQL) queries to implement finder and selector methods. You can specify the following elements of the underlying JDOQL query:

- **Filter expression** - A Java-like expression that specifies a condition that each object returned by the query must satisfy. Corresponds to the WHERE clause in EJB QL.
- **Query parameter declaration** - Specifies the name and the type of one or more query input parameters. Follows the syntax for formal parameters in the Java language.

- **Query variable declaration** - Specifies the name and type of one or more query variables. Follows the syntax for local variables in the Java language. A query filter might use query variables to implement joins.
- **Query ordering declaration** - Specifies the ordering expression of the query. Corresponds to the ORDER BY clause of EJB QL.

The Enterprise Server specific deployment descriptor (`sun-ejb-jar.xml`) provides the following elements to store the EJB 1.1 finder method settings:

```
query-filter
query-params
query-variables
query-ordering
```

The bean developer uses these elements to construct a query. When the finder method that uses these elements executes, the values of these elements are used to execute a query in the database. The objects from the JDOQL query result set are converted into primary key instances to be returned by the EJB 1.1 `ejbFind` method.

The JDO specification, [JSR 12](http://jcp.org/en/jsr/detail?id=12) (<http://jcp.org/en/jsr/detail?id=12>), provides a comprehensive description of JDOQL. The following information summarizes the elements used to define EJB 1.1 finders.

## Query Filter Expression

The filter expression is a String containing a Boolean expression evaluated for each instance of the candidate class. If the filter is not specified, it defaults to true. Rules for constructing valid expressions follow the Java language, with the following differences:

- Equality and ordering comparisons between primitives and instances of wrapper classes are valid.
- Equality and ordering comparisons of Date fields and Date parameters are valid.
- Equality and ordering comparisons of String fields and String parameters are valid.
- White space (non-printing characters space, tab, carriage return, and line feed) is a separator and is otherwise ignored.
- The following assignment operators are not supported.
  - Comparison operators such as `=`, `+=`, and so on
  - Pre- and post-increment
  - Pre- and post-decrement
- Methods, including object construction, are not supported, except for these methods.

```
Collection.contains(Object o)
Collection.isEmpty()
```

```
String.startsWith(String s)
String.endsWith(String e)
```

In addition, the Enterprise Server supports the following nonstandard JDOQL methods.

```
String.like(String pattern)
String.like(String pattern, char escape)
String.substring(int start, int length)
String.indexOf(String str)
String.indexOf(String str, int start)
String.length()
Math.abs(numeric n)
Math.sqrt(double d)
```

- Navigation through a null-valued field, which throws a `NullPointerException`, is treated as if the sub-expression returned `false`.

---

**Note** – Comparisons between floating point values are by nature inexact. Therefore, equality comparisons (`==` and `!=`) with floating point values should be used with caution. Identifiers in the expression are considered to be in the name space of the candidate class, with the addition of declared parameters and variables. As in the Java language, `this` is a reserved word, and refers to the current instance being evaluated.

---

The following expressions are supported.

- Relational operators (`==`, `!=`, `>`, `<`, `>=`, `<=`)
- Boolean operators (`&`, `&&`, `|`, `||`, `~`, `!`)
- Arithmetic operators (`+`, `-`, `*`, `/`)
- String concatenation, only for `String + String`
- Parentheses to explicitly mark operator precedence
- Cast operator
- Promotion of numeric operands for comparisons and arithmetic operations

The rules for promotion follow the Java rules extended by `BigDecimal`, `BigInteger`, and numeric wrapper classes. See the numeric promotions of the Java language specification.

## Query Parameters

The parameter declaration is a `String` containing one or more parameter type declarations separated by commas. This follows the Java syntax for method signatures.

## Query Variables

The type declarations follow the Java syntax for local variable declarations.

## JDOQL Examples

This section provides a few query examples.

### Example 1

The following query returns all players called Michael. It defines a filter that compares the name field with a string literal:

```
name == "Michael"
```

The finder element of the sun-ejb-jar.xml file looks like this:

```
<finder>
  <method-name>findPlayerByName</method-name>
  <query-filter>name == "Michael"</query-filter>
</finder>
```

### Example 2

This query returns all products in a specified price range. It defines two query parameters which are the lower and upper bound for the price: double low, double high. The filter compares the query parameters with the price field:

```
low < price && price < high
```

Query ordering is set to price ascending.

The finder element of the sun-ejb-jar.xml file looks like this:

```
<finder>
  <method-name>findInRange</method-name>
  <query-params>double low, double high</query-params>
  <query-filter>low &lt; price &amp;&amp; price &lt; high</query-filter>
  <query-ordering>price ascending</query-ordering>
</finder>
```

### Example 3

This query returns all players having a higher salary than the player with the specified name. It defines a query parameter for the name java.lang.String name. Furthermore, it defines a variable to which the player's salary is compared. It has the type of the persistence capable class that corresponds to the bean:

```
mypackage.PlayerEJB_170160966_JDOState player
```

The filter compares the salary of the current player denoted by the this keyword with the salary of the player with the specified name:

```
(this.salary > player.salary) && (player.name == name)
```

The finder element of the sun-ejb-jar.xml file looks like this:

```
<finder>
  <method-name>findByHigherSalary</method-name>
  <query-params>java.lang.String name</query-params>
  <query-filter>
    (this.salary > player.salary) && (player.name == name)
  </query-filter>
  <query-variables>
    mypackage.PlayerEJB_170160966_JD0State player
  </query-variables>
</finder>
```

## CMP Restrictions and Optimizations

This section discusses restrictions and performance optimizations that pertain to using CMP.

- [“Disabling ORDER BY Validation” on page 206](#)
- [“Setting the Heap Size on DB2” on page 207](#)
- [“Eager Loading of Field State” on page 207](#)
- [“Restrictions on Remote Interfaces” on page 207](#)
- [“PostgreSQL Case Insensitivity” on page 207](#)
- [“No Support for lock-when-loaded on Sybase” on page 208](#)
- [“Sybase Finder Limitation” on page 208](#)
- [“Date and Time Fields” on page 208](#)
- [“Set RECURSIVE\\_TRIGGERS to false on MSSQL” on page 209](#)
- [“MySQL Database Restrictions” on page 209](#)

### Disabling ORDER BY Validation

EJB QL as defined in the EJB 2.1 Specification defines certain restrictions for the SELECT clause of an ORDER BY query (see section 11.2.8 ORDER BY Clause). This ensures that a query does not order by a field that is not returned by the query. By default, the EJB QL compiler checks the above restriction and throws an exception if the query does not conform.

However, some databases support SQL statements with an ORDER BY column that is not included in the SELECT clause. To disable the validation of the ORDER BY clause against the SELECT clause, set the `DISABLE_ORDERBY_VALIDATION` JVM option as follows:

```
asadmin create-jvm-options --user adminuser
-Dcom.sun.jdo.spi.persistence.support.ejb.ejbqlc.DISABLE_ORDERBY_VALIDATION=true
```

The `DISABLE_ORDERBY_VALIDATION` option is set to `false` by default. Setting it to `true` results in a non-portable module or application.

## Setting the Heap Size on DB2

On DB2, the database configuration parameter APPLHEAPSZ determines the heap size. If you are using the Sun GlassFish or DataDirect database driver, set this parameter to at least 2048 for CMP. For more information, see <http://publib.boulder.ibm.com/infocenter/db2luw/v8/index.jsp?topic=/com.ibm.db2.udb.doc/opt/tsbp2024.htm>.

## Eager Loading of Field State

By default, the EJB container loads the state for all persistent fields (excluding relationship, BLOB, and CLOB fields) before invoking the `ejbLoad` method of the abstract bean. This approach might not be optimal for entity objects with large state if most business methods require access to only parts of the state.

Use the `fetched-with` element in `sun-cmp-mappings.xml` for fields that are used infrequently. See “`fetched-with`” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## Restrictions on Remote Interfaces

The following restrictions apply to the remote interface of an EJB 2.1 bean that uses CMP:

- Do not expose the `get` and `set` methods for CMR fields or the persistence collection classes that are used in container-managed relationships through the remote interface of the bean. However, you are free to expose the `get` and `set` methods that correspond to the CMP fields of the entity bean through the bean’s remote interface.
- Do not expose the container-managed collection classes that are used for relationships through the remote interface of the bean.
- Do not expose local interface types or local home interface types through the remote interface or remote home interface of the bean.

Dependent value classes can be exposed in the remote interface or remote home interface, and can be included in the client EJB JAR file.

## PostgreSQL Case Insensitivity

Case-sensitive behavior cannot be achieved for PostgreSQL databases. PostgreSQL databases internally convert all names to lower case, which makes the following workarounds necessary:

- In the CMP 2.1 runtime, PostgreSQL table and column names are not quoted, which makes these names case insensitive.
- Before running the `capture-schema` command on a PostgreSQL database, make sure table and column names are lower case in the `sun-cmp-mappings.xml` file.

## No Support for lock-when-loaded on Sybase

For EJB 2.1 beans, the lock-when-loaded consistency level is implemented by placing update locks on the data corresponding to a bean when the data is loaded from the database. There is no suitable mechanism available on Sybase databases to implement this feature. Therefore, the lock-when-loaded consistency level is not supported on Sybase databases. See [“consistency” in Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

## Sybase Finder Limitation

If a finder method with an input greater than 255 characters is executed and the primary key column is mapped to a VARCHAR column, Sybase attempts to convert type VARCHAR to type TEXT and generates the following error:

```
com.sybase.jdbc2.jdbc.SybSQLException: Implicit conversion from datatype
'TEXT' to 'VARCHAR' is not allowed. Use the CONVERT function to run this
query.
```

To avoid this error, make sure the finder method input is less than 255 characters.

## Date and Time Fields

If a field type is a Java date or time type (`java.util.Date`, `java.sql.Date`, `java.sql.Time`, `java.sql.Timestamp`), make sure that the field value exactly matches the value in the database.

For example, the following code uses a `java.sql.Date` type as a primary key field:

```
java.sql.Date myDate = new java.sql.Date(System.currentTimeMillis())
BeanA.create(myDate, ...);
```

For some databases, this code results in only the year, month, and date portion of the field value being stored in the database. Later if the client tries to find this bean by primary key as follows, the bean is not found in the database because the value does not match the one that is stored in the database.

```
myBean = BeanA.findByPrimaryKey(myDate);
```

Similar problems can happen if the database truncates the timestamp value while storing it, or if a custom query has a date or time value comparison in its WHERE clause.

For automatic mapping to an Oracle database, fields of type `java.util.Date`, `java.sql.Date`, and `java.sql.Time` are mapped to Oracle's DATE data type. Fields of type `java.sql.Timestamp` are mapped to Oracle's TIMESTAMP(9) data type.



## Set RECURSIVE\_TRIGGERS to false on MSSQL

For version consistency triggers on MSSQL, the property `RECURSIVE_TRIGGERS` must be set to `false`, which is the default. If set to `true`, triggers throw a `java.sql.SQLException`.

Set this property as follows:

```
EXEC sp_dboption 'database-name', 'recursive triggers', 'FALSE'
go
```

You can test this property as follows:

```
SELECT DATABASEPROPERTYEX('database-name', 'IsRecursiveTriggersEnabled')
go
```

## MySQL Database Restrictions

The following restrictions apply when you use a MySQL database with the Enterprise Server for persistence.

- MySQL treats `int1` and `int2` as reserved words. If you want to define `int1` and `int2` as fields in your table, use `'int1'` and `'int2'` field names in your SQL file.
- When `VARCHAR` fields get truncated, a warning is displayed instead of an error. To get an error message, start the MySQL database in strict SQL mode.
- The order of fields in a foreign key index must match the order in the explicitly created index on the primary table.
- The `CREATE TABLE` syntax in the SQL file must end with the following line.

```
) Engine=InnoDB;
```

InnoDB provides MySQL with a transaction-safe (ACID compliant) storage engine having commit, rollback, and crash recovery capabilities.

- For a `FLOAT` type field, the correct precision must be defined. By default, MySQL uses four bytes to store a `FLOAT` type that does not have an explicit precision definition. For example, this causes a number such as `12345.67890123` to be rounded off to `12345.7` during an `INSERT`. To prevent this, specify `FLOAT(10,2)` in the DDL file, which forces the database to use an eight-byte double-precision column. For more information, see <http://dev.mysql.com/doc/mysql/en/numeric-types.html>.
- To use `||` as the string concatenation symbol, start the MySQL server with the `--sql-mode="PIPES_AS_CONCAT"` option. For more information, see <http://dev.mysql.com/doc/refman/5.0/en/server-sql-mode.html> and <http://dev.mysql.com/doc/mysql/en/ansi-mode.html>.

- MySQL always starts a new connection when `autoCommit=true` is set. This ensures that each SQL statement forms a single transaction on its own. If you try to rollback or commit an SQL statement, you get an error message.

```
javax.transaction.SystemException: java.sql.SQLException:  
Can't call rollback when autocommit=true
```

```
javax.transaction.SystemException: java.sql.SQLException:  
Error open transaction is not closed
```

To resolve this issue, add `relaxAutoCommit=true` to the JDBC URL. For more information, see <http://forums.mysql.com/read.php?39,31326,31404>.

- Change the trigger create format from the following:

```
CREATE TRIGGER T_UNKNOWNPKVC1  
BEFORE UPDATE ON UNKNOWNPKVC1  
FOR EACH ROW  
    WHEN (NEW.VERSION = OLD.VERSION)  
BEGIN  
    :NEW.VERSION := :OLD.VERSION + 1;  
END;  
/
```

To the following:

```
DELIMITER |  
CREATE TRIGGER T_UNKNOWNPKVC1  
BEFORE UPDATE ON UNKNOWNPKVC1  
FOR EACH ROW  
    WHEN (NEW.VERSION = OLD.VERSION)  
BEGIN  
    :NEW.VERSION := :OLD.VERSION + 1;  
END  
|  
DELIMITER ;
```

For more information, see <http://dev.mysql.com/doc/mysql/en/create-trigger.html>.

- MySQL does not allow a DELETE on a row that contains a reference to itself. Here is an example that illustrates the issue.

```
create table EMPLOYEE (  
    empId int NOT NULL,  
    salary float(25,2) NULL,  
    mgrId int NULL,  
    PRIMARY KEY (empId),  
    FOREIGN KEY (mgrId) REFERENCES EMPLOYEE (empId)  
    ) ENGINE=InnoDB;
```

```
insert into Employee values (1, 1234.34, 1);
delete from Employee where empId = 1;
```

This example fails with the following error message.

```
ERROR 1217 (23000): Cannot delete or update a parent row:
a foreign key constraint fails
```

To resolve this issue, change the table creation script to the following:

```
create table EMPLOYEE (
    empId int NOT NULL,
    salary float(25,2) NULL,
    mgrId int NULL,
    PRIMARY KEY (empId),
    FOREIGN KEY (mgrId) REFERENCES EMPLOYEE (empId)
    ON DELETE SET NULL
) ENGINE=InnoDB;

insert into Employee values (1, 1234.34, 1);
delete from Employee where empId = 1;
```

This can be done only if the foreign key field is allowed to be null. For more information, see <http://bugs.mysql.com/bug.php?id=12449> and <http://dev.mysql.com/doc/mysql/en/innodb-foreign-key-constraints.html>.

- When an SQL script has foreign key constraints defined, `capture-schema` fails to capture the table information correctly. To work around the problem, remove the constraints and then run `capture-schema`. Here is an example that illustrates the issue.

```
CREATE TABLE ADDRESSBOOKBEANTABLE (ADDRESSBOOKNAME VARCHAR(255)
    NOT NULL PRIMARY KEY,
CONNECTEDUSERS BLOB NULL,
OWNER VARCHAR(256),
FK_FOR_ACCESSPRIVILEGES VARCHAR(256),
CONSTRAINT FK_ACCESSPRIVILEGE FOREIGN KEY (FK_FOR_ACCESSPRIVILEGES)
    REFERENCES ACCESSPRIVILEGESBEANTABLE (ROOT)
) ENGINE=InnoDB;
```

To resolve this issue, change the table creation script to the following:

```
CREATE TABLE ADDRESSBOOKBEANTABLE (ADDRESSBOOKNAME VARCHAR(255)
    NOT NULL PRIMARY KEY,
CONNECTEDUSERS BLOB NULL,
OWNER VARCHAR(256),
FK_FOR_ACCESSPRIVILEGES VARCHAR(256)
) ENGINE=InnoDB;
```



# Developing Java Clients

---

This chapter describes how to develop, assemble, and deploy Java clients in the following sections:

- “[Introducing the Application Client Container](#)” on page 213
- “[Developing Clients Using the ACC](#)” on page 215

## Introducing the Application Client Container

The Application Client Container (ACC) includes a set of Java classes, libraries, and other files that are required for and distributed with Java client programs that execute in their own Java Virtual Machine (JVM). The ACC manages the execution of Java EE application client components (application clients), which are used to access a variety of Java EE services (such as JMS resources, EJB components, web services, security, and so on.) from a JVM outside the Sun GlassFish Enterprise Server.

The ACC communicates with the Enterprise Server using RMI-IIOP protocol and manages the details of RMI-IIOP communication using the client ORB that is bundled with it. Compared to other Java EE containers, the ACC is lightweight.

For information about debugging application clients, see “[Application Client Debugging](#)” on page 71.

---

**Note** – Interoperability between application clients and Enterprise Servers running under different major versions is not supported.

---

## ACC Security

The ACC determines when authentication is needed. This typically occurs when the client refers to an EJB component or when annotations in the client's main class trigger injection

which, in turn, requires contact with the Enterprise Server's naming service. To authenticate the end user, the ACC prompts for any required information, such as a username and password. The ACC itself provides a very simple dialog box to prompt for and read these values.

The ACC integrates with the Enterprise Server's authentication system. It also supports SSL (Secure Socket Layer)/IIOP if configured and when necessary; see [“Using RMI/IIOP Over SSL” on page 224](#).

You can provide an alternate implementation to gather authentication information, tailored to the needs of the application client. To do so, include the class to perform these duties in the application client and identify the fully-qualified name of this class in the `callback-handler` element of the `application-client.xml` descriptor for the client. The ACC uses this class instead of its default class for asking for and reading the authentication information. The class must implement the `javax.security.auth.callback.CallbackHandler` interface. See the Java EE specification, section 9.2, *Application Clients: Security*, for more details.

Application clients can use [“Programmatic Login” on page 102](#).

For more information about security for application clients, see the Java EE 5 Specification, Section EE.9.7, “Java EE Application Client XML Schema.”

## ACC Naming

The client container enables the application clients to use the Java Naming and Directory Interface (JNDI) to look up Java EE services (such as JMS resources, EJB components, web services, security, and so on.) and to reference configurable parameters set at the time of deployment.

## ACC Annotation

Annotation is supported for application clients. For more information, see section 9.4 of the Java EE 5 Specification and [“Java EE Standard Annotation” in \*Sun GlassFish Enterprise Server 2.1 Application Deployment Guide\*](#).

## Java Web Start

Java Web Start allows your application client to be easily launched and automatically downloaded and updated. It is enabled for all application clients by default. For more information, see [“Using Java Web Start” on page 218](#).

## Developing Clients Using the ACC

This section describes the procedure to develop, assemble, and deploy client applications using the ACC. This section describes the following topics:

- “To Access an EJB Component From an Application Client” on page 215
- “To Access a JMS Resource From an Application Client” on page 217
- “Using Java Web Start” on page 218
- “Running an Application Client Using the `appClient Script`” on page 224
- “Using the `package-appClient Script`” on page 224
- “The `client.policy File`” on page 224
- “Using RMI/IIOP Over SSL” on page 224
- “Connecting to a Remote EJB Module Through a Firewall” on page 226

### ▼ To Access an EJB Component From an Application Client

- 1 **In your client code, reference the EJB component by using an `@EJB` annotation or by looking up the JNDI name as defined in the `ejb-jar.xml` file.**

For more information about annotations in application clients, see section 9.4 of the Java EE 5 Specification.

For more information about naming and lookups, see “[Accessing the Naming Context](#)” on page 269.

If load balancing is enabled as in [Step 7](#) and the EJB components being accessed are in a different cluster, the endpoint list must be included in the lookup, as follows:

```
corbaname: host1:port1, host2:port2, .../NameService#ejb/jndi-name
```

- 2 **Define the `@EJB` annotations or the `ejb-ref` elements in the `application-client.xml` file. Define the corresponding `ejb-ref` elements in the `sun-application-client.xml` file.**

For more information on the `application-client.xml` file, see the Java EE 5 Specification, Section EE.9.7, “Java EE Application Client XML Schema.”

For more information on the `sun-application-client.xml` file, see “[The `sun-application-client.xml` file](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*. For a general explanation of how to map JNDI names using reference elements, see “[Mapping References](#)” on page 274.

- 3 **Deploy the application client and EJB component together in an application.**

For more information on deployment, see the *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*. To get the client JAR file, use the `--retrieve` option of the `asadmin deploy` command.

To retrieve the stubs and ties whether or not you requested their generation during deployment, use the `asadmin get-client-stubs` command. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

#### 4 Ensure that the client JAR file includes the following files:

- A Java class to access the bean.
- `application-client.xml` - (optional) Java EE application client deployment descriptor. For information on the `application-client.xml` file, see the Java EE 5 Specification, Section EE.9.7, “Java EE Application Client XML Schema.”
- `sun-application-client.xml` - (optional) Enterprise Server specific client deployment descriptor. For information on the `sun-application-client.xml` file, see “[The sun-application-client.xml file](#)” in [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).
- The MANIFEST.MF file. This file contains a reference to the main class, which states the complete package prefix and class name of the Java client.

#### 5 Prepare the client machine. This step is not needed for Java Web Start.

If you are using the `appclient` script, either package the application client to run on a remote client system using the `package-appclient` script, or copy the following JAR files to the client machine manually and include them in the classpath on the client side:

- `appserv-rt.jar` - available at `as-install/lib`
- `javaee.jar` - available at `as-install/lib`
- The client JAR file

For more information, see “[Using the package-appclient Script](#)” on page 224.

#### 6 To access EJB components that are residing in a remote system, make the following changes to the `sun-acc.xml` file. This step is not needed for Java Web Start.

- Define the `target-server` element’s address attribute to reference the remote server machine. See “[target-server](#)” in [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).
- Define the `target-server` element’s port attribute to reference the ORB port on the remote server.

This information can be obtained from the `domain.xml` file on the remote system. For more information on `domain.xml` file, see the [Sun GlassFish Enterprise Server 2.1 Administration Reference](#).

#### 7 To set up load balancing and failover of remote EJB references, define at least two `target-server` elements in the `sun-acc.xml` file. This step is not needed for Java Web Start.



---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

If the Enterprise Server instance on which the application client is deployed participates in a cluster, the ACC finds all currently active IIOP endpoints in the cluster automatically. However, a client should have at least two endpoints specified for bootstrapping purposes, in case one of the endpoints has failed.

The target-server elements specify one or more IIOP endpoints used for load balancing. The address attribute is an IPv4 address or host name, and the port attribute specifies the port number. See “client-container” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

## 8 Run the application client.

See “Using Java Web Start” on page 218 or “Running an Application Client Using the appClient Script” on page 224.

## ▼ To Access a JMS Resource From an Application Client

### 1 Create a JMS client.

For detailed instructions on developing a JMS client, see “Chapter 33: The Java Message Service API” in the *Java EE 5 Tutorial* (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

### 2 Next, configure a JMS resource on the Enterprise Server.

For information on configuring JMS resources, see “Creating JMS Resources: Destinations and Connection Factories” on page 280.

### 3 Define the @Resource or @Resources annotations or the resource-ref elements in the application-client.xml file. Define the corresponding resource-ref elements in the sun-application-client.xml file.

For more information on the application-client.xml file, see the Java EE 5 Specification, Section EE.9.7, “Java EE Application Client XML Schema.”

For more information on the sun-application-client.xml file, see “The sun-application-client.xml file” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*. For a general explanation of how to map JNDI names using reference elements, see “Mapping References” on page 274.

### 4 Ensure that the client JAR file includes the following files:

- A Java class to access the resource.
- `application-client.xml` - (optional) Java EE application client deployment descriptor. For information on the `application-client.xml` file, see the Java EE 5 Specification, Section EE.9.7, “Java EE Application Client XML Schema.”
- `sun-application-client.xml` - (optional) Enterprise Server specific client deployment descriptor. For information on the `sun-application-client.xml` file, see “[The sun-application-client.xml file](#)” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.
- The MANIFEST.MF file. This file contains a reference to the main class, which states the complete package prefix and class name of the Java client.

## 5 Prepare the client machine. This step is not needed for Java Web Start.

If you are using the `appclient` script, either package the application client to run on a remote client system using the `package-appclient` script, or copy the following JAR files to the client machine manually and include them in the classpath on the client side:

- `appserv-rt.jar` - available at `as-install/lib`
- `javaee.jar` - available at `as-install/lib`
- `imqjmsra.jar` - available at `as-install/lib/install/applications/jmsra`
- The client JAR file

For more information, see “[Using the package-appclient Script](#)” on page 224.

## 6 Run the application client.

See “[Using Java Web Start](#)” on page 218 or “[Running an Application Client Using the appclient Script](#)” on page 224.

# Using Java Web Start

Java Web Start allows your application client to be easily launched and automatically downloaded and updated. General information about Java Web Start is available at <http://java.sun.com/products/javawebstart/reference/api/index.html>.

Java Web Start is discussed in the following topics:

- “[Enabling and Disabling Java Web Start](#)” on page 219
- “[Downloading and Launching an Application Client](#)” on page 219
- “[The Application Client URL](#)” on page 220
- “[Signing JAR Files Used in Java Web Start](#)” on page 221
- “[Error Handling](#)” on page 223
- “[Vendor Icon, Splash Screen, and Text](#)” on page 223

## Enabling and Disabling Java Web Start

Java Web Start is enabled for all application clients by default.

The application developer or deployer can specify that Java Web Start is always disabled for an application client by setting the value of the `eligible` element to `false` in the `sun-application-client.xml` file. See the [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

The Enterprise Server administrator can disable Java Web Start for a previously deployed eligible application client using the `asadmin set` command.

To disable Java Web Start for all eligible application clients in an application, use the following command:

```
asadmin set --user adminuser
domain1.applications.j2ee-application.app-name.java-web-start-enabled="false"
```

To disable Java Web Start for a stand-alone eligible application client, use the following command:

```
asadmin set --user adminuser
domain1.applications.appclient-module.module-name.java-web-start-enabled="false"
```

Setting `java-web-start-enabled="true"` re-enables Java Web Start for an eligible application client. For more information about the `asadmin set` command, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## Downloading and Launching an Application Client

If Java Web Start is enabled for your deployed application client, you can launch it for testing. Simply click on the Launch button next to the application client or application's listing on the App Client Modules page in the Admin Console.

On other machines, you can download and launch the application client using Java Web Start in the following ways:

- Using a web browser, directly enter the URL for the application client. See “[The Application Client URL](#)” on page 220.
- Click on a link to the application client from a web page.
- Use the Java Web Start command `j avaws`, specifying the URL of the application client as a command line argument.
- If the application has previously been downloaded using Java Web Start, you have additional alternatives.
  - Use the desktop icon that Java Web Start created for the application client. When Java Web Start downloads an application client for the first time it asks you if such an icon should be created.

- Use the Java Web Start control panel to launch the application client.

When you launch an application client, Java Web Start contacts the server to see if a newer client version is available. This means you can redeploy an application client without having to worry about whether client machines have the latest version.

## The Application Client URL

The default URL for an application or module generally is as follows:

```
http://host:port/context-root
```

The default URL for a stand-alone application client module is as follows:

```
http://host:port/appclient-module-id
```

The default URL for an application client module embedded within an application is as follows. Note that the relative path to the application client JAR file is included.

```
http://host:port/application-id/appclient-path
```

If the *context-root*, *appclient-module-id*, or *application-id* is not specified during deployment, the name of the JAR or EAR file without the extension is used. If the application client module or application is not in JAR or EAR file format, an *appclient-module-id* or *application-id* is generated.

Regardless of how the *context-root* or *id* is determined, it is written to the server log. For details about naming, see “Naming Standards” in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

To set a different URL for an application client, use the `context-root` subelement of the `java-web-start-access` element in the `sun-application-client.xml` file. This overrides the *appclient-module-id* or *application-id*. See *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

You can also pass arguments to the ACC or to the application client's `main` method as query parameters in the URL. If multiple application client arguments are specified, they are passed in the order specified.

A question mark separates the context root from the arguments. Ampersands (&) separate the arguments and their values. Each argument and each value must begin with `arg=`. Here is an example URL with a `-color` argument for a stand-alone application client. The `-color` argument is passed to the application client's `main` method.

```
http://localhost:8080/testClient?arg=-color&arg=red
```

---

**Note** – If you are using the `javaws URL` command to launch Java Web Start with a URL that contains arguments, enclose the URL in double quotes (") to avoid breaking the URL at the ampersand (&) symbol.

---

Ideally, you should build your production application clients with user-friendly interfaces that collect information which might otherwise be gathered as command-line arguments. This minimizes the degree to which users must customize the URLs that launch application clients using Java Web Start. Command-line argument support is useful in a development environment and for existing application clients that depend on it.

## Signing JAR Files Used in Java Web Start

Java Web Start enforces a security sandbox. By default it grants any application, including application clients, only minimal privileges. Because Java Web Start applications can be so easily downloaded, Java Web Start provides protection from potentially harmful programs that might be accessible over the network. If an application requires a higher privilege level than the sandbox permits, the code that needs privileges must be in a JAR file that was signed. When Java Web Start downloads such a signed JAR file, it displays information about the certificate that was used to sign the JAR, and it asks you whether you want to trust that signed code. If you agree, the code receives elevated permissions and runs. If you reject the signed code, Java Web Start does not start the downloaded application.

The Enterprise Server serves two types of signed JAR files in response to Java Web Start requests. One type is a JAR file installed as part of the Enterprise Server, which starts an application client during a Java Web Start launch: `as-install/lib/appserv-jwsacc.jar`.

The other type is a generated application client JAR file. As part of deployment, the Enterprise Server generates a new application client JAR file that contains classes, resources, and descriptors needed to run the application client on end-user systems. When you deploy an application with the `asadmin deploy` command's `--retrieve` option, use the `asadmin get-client-stubs` command, or select the Generate RMISTubs option from the EJB Modules deployment page in the Admin Console, this is the JAR file retrieved to your system. Because application clients need access beyond the minimal sandbox permissions to work in the Java Web Start environment, the generated application client JAR file must be signed before it can be downloaded to and executed on an end-user system.

A JAR file can be signed automatically or manually. The following sections describe the ways of signing JAR files.

- [“Automatically Signing JAR Files” on page 222](#)
- [“Manually Signing `appserv-jwsacc.jar`” on page 222](#)
- [“Manually Signing the Generated Application Client JAR File” on page 222](#)

## Automatically Signing JAR Files

The Enterprise Server automatically creates a signed version of the required JAR file if none exists. When a Java Web Start request for the `appserv-jwsacc.jar` file arrives, the Enterprise Server looks for `domain-dir/java-web-start/appserv-jwsacc.jar`. When a request for an application's generated application client JAR file arrives, the Enterprise Server looks in the directory `domain-dir/java-web-start/app-name` for a file with the same name as the generated JAR file created during deployment.

In either case, if the requested signed JAR file is absent or older than its unsigned counterpart, the Enterprise Server creates a signed version of the JAR file automatically and deposits it in the relevant directory. Whether the Enterprise Server just signed the JAR file or not, it serves the file from the `domain-dir/java-web-start` directory tree in response to the Java Web Start request.

To sign these JAR files, the Enterprise Server uses its self-signed certificate. When you create a new domain, either by installing the Enterprise Server or by using the `asadmin create-domain` command, the Enterprise Server creates a self-signed certificate and adds it to the domain's key store.

A self-signed certificate is generally untrustworthy because no certification authority vouches for its authenticity. The automatic signing feature uses the same certificate to create all required signed JAR files. To sign different JAR files with different certificates, do the signing manually.

## Manually Signing `appserv-jwsacc.jar`

You can sign the `appserv-jwsacc.jar` file manually any time after you have installed the Enterprise Server. Copy the unsigned file from `as-install/lib` to a different working directory and use the `jar signer` command provided with the JDK to create a signed version of exactly the same name using your certificate. Then manually copy the signed file into `domain-dir/java-web-start`. From then on, the Enterprise Server serves the JAR file signed with your certificate whenever a Java Web Start request asks that domain for the `appserv-jwsacc.jar` file. Note that you can sign each domain's `appserv-jwsacc.jar` file differently.

Remember that if you create a new domain and do not sign `appserv-jwsacc.jar` manually for that domain, the Enterprise Server creates an auto-signed version of it for use by the new domain. Also, if you create a domain-specific signed `appserv-jwsacc.jar`, delete the domain, and then create a new domain with the same name as the just-deleted domain, the Enterprise Server does not remember the earlier signed `appserv-jwsacc.jar`. You must recreate the manually signed version.

## Manually Signing the Generated Application Client JAR File

You can sign the generated application client JAR file for an application any time after you have deployed the application. As you deploy the application, you can specify the `asadmin deploy` command's `--retrieve` option or select the Generate RMISTubs option on the EJB Modules

deployment page in the Admin Console. Doing either of these tasks returns a copy of the generated application client JAR file to a directory you specify. Or, after you have deployed an application, you can download the generated application client JAR file using the `asadmin get-client-stubs` command.

Once you have a copy of the generated application client JAR file, you can sign it using the `jarsigner` tool and your certificate. Then place the signed JAR file in the `domain-dir/java-web-start/app-name` directory. You do not need to restart the server to start using the new signed JAR file.

## Error Handling

When an application client is launched using Java Web Start, any error that the application client logic does not catch and handle is written to `System.err` and displayed in a dialog box. This display appears if an error occurs even before the application client logic receives control. It also appears if the application client code does not catch and handle errors itself.

## Vendor Icon, Splash Screen, and Text

To specify a vendor-specific icon, splash screen, text string, or a combination of these for Java Web Start download and launch screens, use the `vendor` element in the `sun-application-client.xml` file. The complete format of this element's data is as follows:

```
<vendor>icon-image-URI::splash-screen-image-URI::vendor-text</vendor>
```

The following example `vendor` element contains an icon, a splash screen, and a text string:

```
<vendor>images/icon.jpg::otherDir/splash.jpg::MyCorp, Inc.</vendor>
```

The following example `vendor` element contains an icon and a text string:

```
<vendor>images/icon.jpg::MyCorp, Inc.</vendor>
```

The following example `vendor` element contains a splash screen and a text string; note the initial double colon:

```
<vendor>::otherDir/splash.jpg::MyCorp, Inc.</vendor>
```

The following example `vendor` element contains only a text string:

```
<vendor>MyCorp, Inc.</vendor>
```

The default value is the text string `Application Client`.

For more information about the `sun-application-client.xml` file, see the [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

## Running an Application Client Using the `appclient` Script

To run an application client that does *not* have Java Web Start enabled, you can launch the ACC using the `appclient` script. This is optional. This script is located in the `as-install/bin` directory. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Using the `package-appclient` Script

You can package an application client that does *not* have Java Web Start enabled into a single `appclient.jar` file using the `package-appclient` script. This is optional. This script is located in the `as-install/bin` directory. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## The `client.policy` File

The `client.policy` file is the J2SE policy file used by the application client. Each application client has a `client.policy` file. The default policy file limits the permissions of Java EE deployed application clients to the minimal set of permissions required for these applications to operate correctly. If an application client requires more than this default set of permissions, edit the `client.policy` file to add the custom permissions that your application client needs. Use the J2SE standard policy tool or any text editor to edit this file.

For more information on using the J2SE policy tool, see <http://java.sun.com/docs/books/tutorial/security1.2/tour2/index.html>.

For more information about the permissions you can set in the `client.policy` file, see <http://java.sun.com/javase/6/docs/technotes/guides/security/permissions.html>.

## Using RMI/IIOP Over SSL

You can configure RMI/IIOP over SSL in two ways: using a username and password, or using a client certificate.

To use a username and password, configure the `ior-security-config` element in the `sun-ejb-jar.xml` file. The following configuration establishes SSL between an application client and an EJB component using a username and password. The user has to login to the ACC using either the `sun-acc.xml` mechanism or the [“Programmatic Login” on page 102](#) mechanism.

```
<ior-security-config>
  <transport-config>
    <integrity>required</integrity>
```



```

    <confidentiality>required</confidentiality>
    <establish-trust-in-target>supported</establish-trust-in-target>
    <establish-trust-in-client>none</establish-trust-in-client>
</transport-config>
<as-context>
  <auth-method>username_password</auth-method>
  <realm>default</realm>
  <required>true</required>
</as-context>
<sas-context>
  <caller-propagation>none</caller-propagation>
</sas-context>
</ior-security-config>

```

For more information about the `sun-ejb-jar.xml` and `sun-acc.xml` files, see the *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

To use a client certificate, configure the `ior-security-config` element in the `sun-ejb-jar.xml` file. The following configuration establishes SSL between an application client and an EJB component using a client certificate.

```

<ior-security-config>
  <transport-config>
    <integrity>required</integrity>
    <confidentiality>required</confidentiality>
    <establish-trust-in-target>supported</establish-trust-in-target>
    <establish-trust-in-client>required</establish-trust-in-client>
  </transport-config>
  <as-context>
    <auth-method>none</auth-method>
    <realm>default</realm>
    <required>>false</required>
  </as-context>
  <sas-context>
    <caller-propagation>none</caller-propagation>
  </sas-context>
</ior-security-config>

```

To use a client certificate, you must also specify the system properties for the keystore and truststore to be used in establishing SSL. To use SSL with the Application Client Container (ACC), you need to set `VMARGS` environment variable in one of the following ways:

- Set the environment variable `VMARGS` in the shell. For example, in the `ksh` or `bash` shell, the command to set this environment variable would be as follows:

```

export VMARGS="-Djavax.net.ssl.keyStore=${keystore.db.file}
-Djavax.net.ssl.trustStore=${truststore.db.file}
-Djavax.net.ssl.keyStorePassword=${ssl.password}
-Djavax.net.ssl.trustStorePassword=${ssl.password}"

```

- Set the env element in the asant script (see [Chapter 3, “The asant Utility”](#)). For example:

```
<target name="runclient">
  <exec executable="${S1AS_HOME}/bin/appclient">
    <env key="VMARGS" value=" -Djavax.net.ssl.keyStore=${keystore.db.file}
      -Djavax.net.ssl.trustStore=${truststore.db.file}
      -Djavax.net.ssl.keyStorePassword=${ssl.password}
      -Djavax.net.ssl.trustStorePassword=${ssl.password}"/>
    <arg value="-client"/>
    <arg value="${appClient.jar}"/>
  </exec>
</target>
```

## Connecting to a Remote EJB Module Through a Firewall

To deploy and run an application client that connects to an EJB module on a Enterprise Server instance that is behind a firewall, you must set ORB Virtual Address Agent Implementation (ORBVAA) options. Use the `asadmin create-jvm-options` command as follows:

```
asadmin create-jvm-options --user adminuser -Dcom.sun.corba.ee.ORBVAAMHost=public-IP-adress
asadmin create-jvm-options --user adminuser -Dcom.sun.corba.ee.ORBVAAPort=public-port
asadmin create-jvm-options --user adminuser
-Dcom.sun.corba.ee.ORBUserConfigurators.com.sun.corba.ee.impl.plugin.howlb.VirtualAddressAgentImpl=x
```

Set the `ORBVAAMHost` and `ORBVAAPort` options to the host and port of the public address. The `ORBUserConfigurators` option tells the ORB to create an instance of the `VirtualAddressAgentImpl` class and invoke the `configure` method on the resulting object, which must implement the `com.sun.corba.ee.spi.orb.ORBConfigurator` interface. The `ORBUserConfigurators` value doesn't matter. Together, these options create an ORB that in turn creates `Object` references (the underlying implementation of remote EJB references) containing the public address, while the ORB listens on the private address specified for the IIOP port in the Enterprise Server configuration.

# Developing Connectors

---

This chapter describes Sun GlassFish Enterprise Server support for the J2EE™ 1.5 Connector Architecture (CA).

The J2EE Connector Architecture provides a Java solution to the problem of connectivity between multiple application servers and existing enterprise information systems (EISs). By using the J2EE Connector architecture, EIS vendors no longer need to customize their product for each application server. Application server vendors who conform to the J2EE Connector architecture do not need to write custom code to add connectivity to a new EIS.

This chapter uses the terms *connector* and *resource adapter* interchangeably. Both terms refer to a resource adapter module that is developed in conformance with the J2EE Connector Specification.

For more information about connectors, see [J2EE Connector Architecture](http://java.sun.com/j2ee/connector/) (<http://java.sun.com/j2ee/connector/>) and “Chapter 37: J2EE Connector Architecture” in the [Java EE 5 Tutorial](http://java.sun.com/javaee/5/docs/tutorial/doc/index.html) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

For connector examples, see [http://developers.sun.com/prodtech/appserver/reference/techart/as8\\_connectors](http://developers.sun.com/prodtech/appserver/reference/techart/as8_connectors).

This chapter includes the following topics:

- “Connector Support in the Enterprise Server” on page 228
- “Deploying and Configuring a Stand-Alone Connector Module” on page 229
- “Redeploying a Stand-Alone Connector Module” on page 230
- “Deploying and Configuring an Embedded Resource Adapter” on page 230
- “Advanced Connector Configuration Options” on page 231
- “Inbound Communication Support” on page 234
- “Configuring a Message Driven Bean to Use a Resource Adapter” on page 235

## Connector Support in the Enterprise Server

The Enterprise Server supports the development and deployment of resource adapters that are compatible with Connector specification (and, for backward compatibility, the Connector 1.0 specification).

The Connector 1.0 specification defines the outbound connectivity system contracts between the resource adapter and the Enterprise Server. The Connector 1.5 specification introduces major additions in defining system level contracts between the Enterprise Server and the resource adapter with respect to the following:

- Inbound connectivity from an EIS - Defines the transaction and message inflow system contracts for achieving inbound connectivity from an EIS. The message inflow contract also serves as a standard message provider pluggability contract, thereby allowing various providers of messaging systems to seamlessly plug in their products with any application server that supports the message inflow contract.
- Resource adapter life cycle management and thread management - These features are available through the lifecycle and work management contracts.

## Connector Architecture for JMS and JDBC

In the Admin Console, connector, JMS, and JDBC resources are handled differently, but they use the same underlying Connector architecture. In the Enterprise Server, all communication to an EIS, whether to a message provider or an RDBMS, happens through the Connector architecture. To provide JMS infrastructure to clients, the Enterprise Server uses the Sun GlassFish Message Queue software. To provide JDBC infrastructure to clients, the Enterprise Server uses its own JDBC system resource adapters. The application server automatically makes these system resource adapters available to any client that requires them.

For more information about JMS in the Enterprise Server, see [Chapter 18, “Using the Java Message Service.”](#) For more information about JDBC in the Enterprise Server, see [Chapter 15, “Using the JDBC API for Database Access.”](#)

## Connector Configuration

The Enterprise Server does not need to use `sun-ra.xml`, which previous Enterprise Server versions used, to store server-specific deployment information inside a Resource Adapter Archive (RAR) file. (However, the `sun-ra.xml` file is still supported for backward compatibility.) Instead, the information is stored in the server configuration. As a result, you can create multiple connector connection pools for a connection definition in a functional resource adapter instance, and you can create multiple user-accessible connector resources (that is, registering a resource with a JNDI name) for a connector connection pool. In addition, dynamic changes can be made to connector connection pools and the connector resource properties without restarting the Enterprise Server.

# Deploying and Configuring a Stand-Alone Connector Module

You can deploy a stand-alone connector module using the Admin Console or the `asadmin` command. For information about using the Admin Console, click the Help button in the Admin Console. For information about using the `asadmin` command, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

Deploying a stand-alone connector module allows multiple deployed Java EE applications to share the connector module. A resource adapter configuration is automatically created for the connector module.

## ▼ To Deploy and Configure a Stand-Alone Connector Module

### 1 Deploy the connector module in one of the following ways.

- In the Admin Console, open the Applications component and select Connector Modules. When you deploy the connector module, a resource adapter configuration is automatically created for the connector module.
- Use the `asadmin deploy` or `asadmin deploydir` command. To override the default configuration properties of a resource adapter, if necessary, use the `asadmin create-resource-adapter-config` command.

### 2 Configure connector connection pools for the deployed connector module in one of the following ways:

- In the Admin Console, open the Resources component, select Connectors, and select Connector Connection Pools.
- Use the `asadmin create-connector-connection-pool` command.

### 3 Configure connector resources for the connector connection pools in one of the following ways.

- In the Admin Console, open the Resources component, select Connectors, and select Connector Resources.
- Use the `asadmin create-connector-resource` command.

This associates a connector resource with a JNDI name.

### 4 Create an administered object for an inbound resource adapter, if necessary, in one of the following ways:

- In the Admin Console, open the Resources component, select Connectors, and select Admin Object Resources.
- Use the `asadmin create-admin-object` command.

## Redeploying a Stand-Alone Connector Module

Redeployment of a connector module maintains all connector connection pools, connector resources, and administered objects defined for the previously deployed connector module. You need not reconfigure any of these resources.

However, you should redeploy any dependent modules. A dependent module uses or refers to a connector resource of the redeployed connector module. Redeployment of a connector module results in the shared class loader reloading the new classes. Other modules that refer to the old resource adapter classes must be redeployed to gain access to the new classes. For more information about class loaders, see [Chapter 2, “Class Loaders.”](#)

During connector module redeployment, the server log provides a warning indicating that all dependent applications should be redeployed. Client applications or application components using the connector module’s resources may throw class cast exceptions if dependent applications are not redeployed after connector module redeployment.

To disable automatic redeployment, set the `--force` option to `false`. In this case, if the connector module has already been deployed, the Enterprise Server provides an error message.

## Deploying and Configuring an Embedded Resource Adapter

A connector module can be deployed as a Java EE component in a Java EE application. Such connectors are only visible to components residing in the same Java EE application. Simply deploy this Java EE application as you would any other Java EE application.

You can create new connector connection pools and connector resources for a connector module embedded within a Java EE application by prefixing the connector name with *app-name#*. For example, if an application `appX.ear` has `jdbcra.rar` embedded within it, the connector connection pools and connector resources refer to the connector module as `appX#jdbcra`.

However, an embedded connector module cannot be undeployed using the name *app-name#connector-name*. To undeploy the connector module, you must undeploy the application in which it is embedded.

The association between the physical JNDI name for the connector module in the Enterprise Server and the logical JNDI name used in the application component is specified in the Enterprise Server specific XML descriptor `sun-ejb-jar.xml`.

# Advanced Connector Configuration Options

You can use these advanced connector configuration options:

- “Thread Pools” on page 231
- “Security Maps” on page 231
- “Overriding Configuration Properties” on page 232
- “Testing a Connector Connection Pool” on page 232
- “Handling Invalid Connections” on page 233
- “Setting the Shutdown Timeout” on page 233
- “Using Last Agent Optimization of Transactions” on page 234

## Thread Pools

Connectors can submit work instances to the Enterprise Server for execution. By default, the Enterprise Server services work requests for all connectors from its default thread pool. However, you can associate a specific user-created thread pool to service work requests from a connector. A thread pool can service work requests from multiple resource adapters. To create a thread pool:

- In the Admin Console, select Thread Pools under the relevant configuration. For details, click the Help button in the Admin Console.
- Use the `asadmin create-threadpool` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

To associate a connector with a thread pool:

- In the Admin Console, open the Applications component and select Connector Modules. Deploy the module, or select the previously deployed module. Specify the name of the thread pool in the Thread Pool ID field. For details, click the Help button in the Admin Console.
- Use the `--threadpoolid` option of the `asadmin create-resource-adapter-config` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

If you create a resource adapter configuration for a connector module that is already deployed, the connector module deployment is restarted with the new configuration properties.

## Security Maps

Create a security map for a connector connection pool to map an application principal or a user group to a back end EIS principal. The security map is usually used in situations where one or more EIS back end principals are used to execute operations (on the EIS) initiated by various principals or user groups in the application.

To create or update security maps for a connector connection pool:

- In the Admin Console, open the Resources component, select Connectors, select Connector Connection Pools, and select the Security Maps tab. For details, click the Help button in the Admin Console.
- Use the `asadmin create-connector-security-map` command. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

If a security map already exists for a connector connection pool, the new security map is appended to the previous one. The connector security map configuration supports the use of the wildcard asterisk (\*) to indicate all users or all user groups.

When an application principal initiates a request to an EIS, the Enterprise Server first checks for an exact match to a mapped back end EIS principal using the security map defined for the connector connection pool. If there is no exact match, the Enterprise Server uses the wild card character specification, if any, to determine the mapped back end EIS principal.

## Overriding Configuration Properties

You can override the properties (`config-property` elements) specified in the `ra.xml` file of a resource adapter. Use the `asadmin create-resource-adapter-config` command to create a configuration for a resource adapter. Use this command's `--property` option to specify a name-value pair for a resource adapter property.

You can use the `asadmin create-resource-adapter-config` command either before or after resource adapter deployment. If it is executed after deploying the resource adapter, the existing resource adapter is restarted with the new properties. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

You can also use token replacement for overriding resource adapter configuration properties in individual server instances when the resource adapter is deployed to a cluster. For example, for a property called `inboundPort`, you can assign the value `${inboundPort}`. You can then assign a different value to this property for each server instance. Changes to system properties take effect upon server restart.

## Testing a Connector Connection Pool

After configuring a connector connection pool, use the `asadmin ping-connection-pool` command to test the health of the underlying connections. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).



## Handling Invalid Connections

If a resource adapter generates a `ConnectionErrorOccured` event, the Enterprise Server considers the connection invalid and removes the connection from the connection pool. Typically, a resource adapter generates a `ConnectionErrorOccured` event when it finds a `ManagedConnection` object unusable. Reasons can be network failure with the EIS, EIS failure, fatal problems with resource adapter, and so on. If the `fail-all-connections` property in the connection pool configuration is set to `true`, all connections are destroyed and the pool is recreated.

The `is-connection-validation-required` property specifies whether connections have to be validated before being given to the application. If a resource's validation fails, it is destroyed, and a new resource is created and returned.

You can set the `fail-all-connections` and `is-connection-validation-required` configuration properties during creation of a connector connection pool. Or, you can use the `asadmin set` command to dynamically reconfigure a previously set property. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

The interface `ValidatingManagedConnectionFactory` exposes the method `getInvalidConnections` to allow retrieval of the invalid connections. The Enterprise Server checks if the resource adapter implements this interface, and if it does, invalid connections are removed when the connection pool is resized.

## Setting the Shutdown Timeout

According to the Connector specification, while an application server shuts down, all resource adapters should be stopped. A resource adapter might hang during shutdown, since shutdown is typically a resource intensive operation. To avoid such a situation, you can set a timeout that aborts resource adapter shutdown if exceeded. The default timeout is 30 seconds per resource adapter module. To configure this timeout:

- In the Admin Console, select Connector Service under the relevant configuration and edit the shutdown Timeout field. For details, click the Help button in the Admin Console.
- Use the following command:

```
asadmin set --user adminuser server1.connector-service.shutdown-timeout-in-seconds="num-sec"
```

For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

The Enterprise Server deactivates all message-driven bean deployments before stopping a resource adapter.

## Using Last Agent Optimization of Transactions

Transactions that involve multiple resources or multiple participant processes are *distributed* or *global* transactions. A global transaction can involve one non-XA resource if last agent optimization is enabled. Otherwise, all resources must be XA. For more information about transactions in the Enterprise Server, see [Chapter 16, “Using the Transaction Service.”](#)

The Connector specification requires that if a resource adapter supports `XATransaction`, the `ManagedConnection` created from that resource adapter must support both distributed and local transactions. Therefore, even if a resource adapter supports `XATransaction`, you can configure its connector connection pools as non-XA or without transaction support for better performance. A non-XA resource adapter becomes the last agent in the transactions in which it participates.

The value of the connection pool configuration property `transaction-support` defaults to the value of the `transaction-support` property in the `ra.xml` file. The connection pool configuration property can override the `ra.xml` file property if the transaction level in the connection pool configuration property is lower. If the value in the connection pool configuration property is higher, it is ignored.

## Inbound Communication Support

The Connector specification defines the transaction and message inflow system contracts for achieving inbound connectivity from an EIS. The message inflow contract also serves as a standard message provider pluggability contract, thereby allowing various message providers to seamlessly plug in their products with any application server that supports the message inflow contract. In the inbound communication model, the EIS initiates all communication to an application. An application can be composed of enterprise beans (session, entity, or message-driven beans), which reside in an EJB container.

Incoming messages are received through a message endpoint, which is a message-driven bean. This message-driven bean asynchronously consumes messages from a message provider. An application can also synchronously send and receive messages directly using messaging style APIs.

A resource adapter supporting inbound communication provides an instance of an `ActivationSpec` JavaBean class for each supported message listener type. Each class contains a set of configurable properties that specify endpoint activation configuration information during message-driven bean deployment. The required `config-property` element in the `ra.xml` file provides a list of configuration property names required for each activation specification. An endpoint activation fails if the required property values are not specified. Values for the properties that are overridden in the message-driven bean's deployment descriptor are applied to the `ActivationSpec` JavaBean when the message-driven bean is deployed.

Administered objects can also be specified for a resource adapter, and these JavaBeans are specific to a messaging style or message provider. For example, some messaging styles may need applications to use special administered objects (such as `Queue` and `Topic` objects in JMS). Applications use these objects to send and synchronously receive messages using connection objects using messaging style APIs. For more information about administered objects, see [Chapter 18, “Using the Java Message Service.”](#)

## Configuring a Message Driven Bean to Use a Resource Adapter

The Connectors specification’s message inflow contract provides a generic mechanism to plug in a wide-range of message providers, including JMS, into a Java-EE-compatible application server. Message providers use a resource adapter and dispatch messages to message endpoints, which are implemented as message-driven beans.

The message-driven bean developer provides activation configuration information in the message-driven bean’s `ejb-jar.xml` file. Configuration information includes messaging-style-specific configuration details, and possibly message-provider-specific details as well. The message-driven bean deployer uses this configuration information to set up the activation specification JavaBean. The activation configuration properties specified in `ejb-jar.xml` override configuration properties in the activation specification definition in the `ra.xml` file.

According to the EJB specification, the messaging-style-specific descriptor elements contained within the activation configuration element are not specified because they are specific to a messaging provider. In the following sample message-driven bean `ejb-jar.xml`, a message-driven bean has the following activation configuration property names: `destinationType`, `SubscriptionDurability`, and `MessageSelector`.

```
<!-- A sample MDB that listens to a JMS Topic -->
<!-- message-driven bean deployment descriptor -->
...
<activation-config>
  <activation-config-property>
    <activation-config-property-name>
      destinationType
    </activation-config-property-name>
    <activation-config-property-value>
      javax.jms.Topic
    </activation-config-property-value>
  </activation-config-property>
  <activation-config-property>
    <activation-config-property-name>
      SubscriptionDurability
    </activation-config-property-name>
    <activation-config-property-value>
```

```

        Durable
    </activation-config-property-value>
</activation-config-property>
<activation-config-property>
    <activation-config-property-name>
        MessageSelector
    </activation-config-property-name>
    <activation-config-property-value>
        JMSType = 'car' AND color = 'blue'
    </activation-config-property-value>
</activation-config-property>
...
</activation-config>
...

```

When the message-driven bean is deployed, the value for the `resource-adapter-mid` element in the `sun-ejb-jar.xml` file is set to the resource adapter module name that delivers messages to the message endpoint (to the message-driven bean). In the following example, the `jmsra` JMS resource adapter, which is the bundled resource adapter for the Sun GlassFish Message Queue message provider, is specified as the resource adapter module identifier for the `SampleMDB` bean.

```

<sun-ejb-jar>
<enterprise-beans>
    <unique-id>1</unique-id>
    <ejb>
        <ejb-name>SampleMDB</ejb-name>
        <jndi-name>SampleQueue</jndi-name>
        <!-- JNDI name of the destination from which messages would be
            delivered from MDB needs to listen to -->
        ...
        <mdb-resource-adapter>
            <resource-adapter-mid>jmsra</resource-adapter-mid>
            <!-- Resource Adapter Module Id that would deliver messages to
                this message endpoint -->
        </mdb-resource-adapter>
        ...
    </ejb>
    ...
</enterprise-beans>
...
</sun-ejb-jar>

```

When the message-driven bean is deployed, the Enterprise Server uses the `resourceadapter-mid` setting to associate the resource adapter with a message endpoint through the message inflow contract. This message inflow contract with the application server gives the resource adapter a handle to the `MessageEndpointFactory` and the `ActivationSpec` JavaBean, and the adapter uses this handle to deliver messages to the message endpoint instances (which are created by the `MessageEndpointFactory`).

When a message-driven bean first created for use on the Enterprise Server 7 is deployed, the Connector runtime transparently transforms the previous deployment style to the current connector-based deployment style. If the deployer specifies neither a `resource-adapter-mid` property nor the Message Queue resource adapter's activation configuration properties, the Connector runtime maps the message-driven bean to the `jms ra` system resource adapter and converts the JMS-specific configuration to the Message Queue resource adapter's activation configuration properties.



# Developing Lifecycle Listeners

---

Lifecycle listener modules provide a means of running short or long duration Java-based tasks within the application server environment, such as instantiation of singletons or RMI servers. These modules are automatically initiated at server startup and are notified at various phases of the server life cycle.

All lifecycle module classes and interfaces are in the *as-install/lib/appserv-ext.jar* file.

For Javadoc tool pages relevant to lifecycle modules, go to <http://glassfish.dev.java.net/nonav/javaee5/api/index.html> and click on the `com.sun.appserv.server` package.

The following sections describe how to create and use a lifecycle listener module:

- “Server Life Cycle Events” on page 239
- “The LifecycleListener Interface” on page 240
- “The LifecycleEvent Class” on page 240
- “The Server Lifecycle Event Context” on page 241
- “Deploying a Lifecycle Module” on page 241
- “Considerations for Lifecycle Modules” on page 242

## Server Life Cycle Events

A lifecycle module listens for and performs its tasks in response to the following events in the server life cycle:

- After the `INIT_EVENT`, the server reads the configuration, initializes built-in subsystems (such as security and logging services), and creates the containers.
- After the `STARTUP_EVENT`, the server loads and initializes deployed applications.
- After the `READY_EVENT`, the server is ready to service requests.
- After the `SHUTDOWN_EVENT`, the server destroys loaded applications and stops.

- After the `TERMINATION_EVENT`, the server closes the containers, the built-in subsystems, and the server runtime environment.

These events are defined in the `LifecycleEvent` class.

The lifecycle modules that listen for these events implement the `LifecycleListener` interface.

## The LifecycleListener Interface

To create a lifecycle module is to configure a customized class that implements the `com.sun.appserv.server.LifecycleListener` interface. You can create and simultaneously execute multiple lifecycle modules.

The `LifecycleListener` interface defines this method:

```
public void handleEvent(com.sun.appserv.server.LifecycleEvent event)
throws ServerLifecycleException
```

This method responds to a lifecycle event and throws a `com.sun.appserv.server.ServerLifecycleException` if an error occurs.

A sample implementation of the `LifecycleListener` interface is the `LifecycleListenerImpl.java` file, which you can use for testing lifecycle events.

## The LifecycleEvent Class

The `com.sun.appserv.server.LifecycleEvent` class defines a server life cycle event. The following methods are associated with the event:

- `public java.lang.Object getData()`  
This method returns an instance of `java.util.Properties` that contains the properties defined for the lifecycle module in the `domain.xml` file. For more information about the `domain.xml` file, see the [Sun GlassFish Enterprise Server 2.1 Administration Reference](#).
- `public int getEventType()`  
This method returns the type of the last event, which is `INIT_EVENT`, `STARTUP_EVENT`, `READY_EVENT`, `SHUTDOWN_EVENT`, or `TERMINATION_EVENT`.
- `public com.sun.appserv.server.LifecycleEventContext getLifecycleEventContext()`  
This method returns the lifecycle event context, described next.

A `LifecycleEvent` instance is passed to the `LifecycleListener.handleEvent` method.



## The Server Lifecycle Event Context

The `com.sun.appserv.server.LifecycleEventContext` interface exposes runtime information about the server. The lifecycle event context is created when the `LifecycleEvent` class is instantiated at server initialization. The `LifecycleEventContext` interface defines these methods:

- `public java.lang.String[] getCmdLineArgs()`  
This method returns the server startup command-line arguments.
- `public java.lang.String getInstallRoot()`  
This method returns the server installation root directory.
- `public java.lang.String getInstanceName()`  
This method returns the server instance name.
- `public javax.naming.InitialContext getInitialContext()`  
This method returns the initial JNDI naming context. The naming environment for lifecycle modules is installed after the `STARTUP_EVENT`. A lifecycle module can look up any resource by its `jndi-name` attribute after the `READY_EVENT`.

If a lifecycle module needs to look up resources, it can do so after the `READY_EVENT`. It can use the `getInitialContext()` method to get the initial context to which all the resources are bound.

## Deploying a Lifecycle Module

You can deploy a lifecycle module using the following tools:

- In the Admin Console, open the Applications component and go to the Lifecycle Modules page. For details, click the Help button in the Admin Console.
- Use the `asadmin create-lifecycle-module` command. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

You do not need to specify a classpath for the lifecycle module if you place it in the `domain-dir/lib` or `domain-dir/lib/classes` directory for the Domain Administration Server. Do not place it in the `lib` directory for a particular instance, or it will be deleted when that instance synchronizes with the Domain Administration Server.

After you deploy a lifecycle module, you must restart the server to activate it. The server instantiates it and registers it as a lifecycle event listener at server initialization.

---

**Note** – If the `is-failure-fatal` setting is set to `true` (the default is `false`), lifecycle module failure prevents server initialization or startup, but not shutdown or termination.

---

## Considerations for Lifecycle Modules

The resources allocated at initialization or startup should be freed at shutdown or termination. The lifecycle module classes are called synchronously from the main server thread, therefore it is important to ensure that these classes don't block the server. Lifecycle modules can create threads if appropriate, but these threads must be stopped in the shutdown and termination phases.

The `LifeCycleModule` class loader is the parent class loader for lifecycle modules. Each lifecycle module's `classpath` in `domain.xml` is used to construct its class loader. All the support classes needed by a lifecycle module must be available to the `LifeCycleModule` class loader or its parent, the `Connector` class loader.

You must ensure that the `server.policy` file is appropriately set up, or a lifecycle module trying to perform a `System.exec()` might cause a security access violation. For details, see [“The server.policy File” on page 89](#).

The configured properties for a lifecycle module are passed as properties after the `INIT_EVENT`. The JNDI naming context is not available before the `STARTUP_EVENT`. If a lifecycle module requires the naming context, it can get this after the `STARTUP_EVENT`, `READY_EVENT`, or `SHUTDOWN_EVENT`.

# Developing Custom MBeans

---

An MBean is a managed Java object, similar to a `JavaBean™`, that follows the design patterns set forth in the instrumentation level of the `Java™ Management Extensions (JMX™)` specification. An MBean can represent a device, an application, or any resource that needs to be managed. MBeans expose a management interface: a set of readable and/or writable attributes and a set of invocable operations, along with a self-description. The actual runtime interface of an MBean depends on the type of that MBean. MBeans can also emit notifications when certain defined events occur. Unlike other components, MBeans have no annotations or deployment descriptors.

The Sun GlassFish Enterprise Server supports the development of custom MBeans as part of the self-management infrastructure or as separate applications. All types of MBeans (standard, dynamic, open, and model) are supported. For more about self-management, see [Chapter 20, “Using the Application Server Management Extensions,”](#) and [Chapter 19, “Configuring Management Rules,”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

For general information about JMX technology, including how to download the JMX specification, see <http://java.sun.com/products/JavaManagement/index.jsp>.

For a useful overview of JMX technology, see <http://java.sun.com/javase/6/docs/technotes/guides/jmx/overview/JMXoverviewTOC.html>.

For a tutorial of JMX technology, see <http://java.sun.com/javase/6/docs/technotes/guides/jmx/tutorial/tutorialTOC.html>.

This chapter includes the following topics:

- “The MBean Life Cycle” on page 244
- “MBean Class Loading” on page 245
- “Creating, Deleting, and Listing MBeans” on page 245
- “The MBeanServer in the Enterprise Server” on page 247
- “Enabling and Disabling MBeans” on page 248
- “Handling MBean Attributes” on page 248

# The MBean Life Cycle

The MBean life cycle proceeds as follows:

1. The MBean's class files are installed in the Enterprise Server. See [“MBean Class Loading” on page 245](#).
2. The MBean is deployed using the `asadmin create-mbean` command or the Admin Console. See [“Creating, Deleting, and Listing MBeans” on page 245](#).
3. The MBean class is loaded. This also results in loading of other classes. The delegation model is used. See the class loader diagram in [“The Class Loader Hierarchy” on page 33](#).
4. The MBean is instantiated. Its default constructor is invoked reflectively. This is why the MBean class must have a default constructor.
5. The MBean's `ObjectName` is determined according to the following algorithm.

- If you specify the `ObjectName`, it is used as is. The domain must be `user:.` The property name `server` is reserved and cannot be used.

The Enterprise Server automatically appends `server=target` to the `ObjectName` when the MBean is registered, where the `target` is the name of the server instance or cluster to which the MBean is deployed.

- If the MBean implements the `MBeanRegistration` interface, it must provide an `ObjectName` in its `preregister()` method that follows the same rules.
  - If the `ObjectName` is not specified directly or through the `MBeanRegistration` interface, the default is `user:type=impl-class-name`.
6. All attributes are set using `setAttribute` calls in the order in which the attributes are specified. Attempting to specify a read-only attribute results in an error.

If attribute values are set during MBean deployment, these values are passed in as `String` objects. Therefore, attribute types must be Java classes having constructors that accept `String` objects. If you specify an attribute that does not have such a constructor, an error is reported.

Attribute values specified during MBean deployment are persisted to the Enterprise Server configuration. Changes to attributes after registration through a JMX connector such as JConsole do not affect the Enterprise Server configuration. To change an attribute value in the Enterprise Server configuration, use the `asadmin set` command. See [“Handling MBean Attributes” on page 248](#).

7. If the MBean is enabled, the `MBeanServer.registerMBean(Object, ObjectName)` method is used to register the MBean in the `MBeanServer`. This is the only method called by the Enterprise Server runtime. See [“The MBeanServer in the Enterprise Server” on page 247](#).

MBeans are enabled by default. Disabling an MBean deregisters it. See [“Enabling and Disabling MBeans” on page 248](#).

8. The MBean is automatically loaded, instantiated, and registered upon each server restart.

9. When the MBean is deleted using the `asadmin delete-mbean` command or the Admin Console, the MBean is first deregistered if it is enabled, then the MBean definition is deleted from the configuration. The class files are not deleted, however.

## MBean Class Loading

After you develop a custom MBean, copy its class files (or JAR file) into the MBean class loader directory, `domain-dir/applications/mbeans`. You have two choices of where to place any dependent classes:

- Common class loader – Copy the classes as JAR files into the `domain-dir/lib` directory, or copy the classes as `.class` files into the `domain-dir/lib/classes` directory. The classes are loaded when you restart the Enterprise Server. The classes are available to all other MBeans, applications, and modules deployed on servers that share the same configuration.
- MBean class loader – Copy the classes into the `domain-dir/applications/mbeans` directory. No restart is required. The classes are available to all other MBeans deployed on servers that share the same configuration, but *not* to applications and modules.

After copying the classes, register the MBean using the `asadmin create-mbean` command. See “[The `asadmin create-mbean` Command](#)” on page 245.

For general information about Enterprise Server class loaders, see [Chapter 2, “Class Loaders.”](#)

## Creating, Deleting, and Listing MBeans

This section describes the following commands:

- Use the `asadmin create-mbean` command to deploy, or *register*, an MBean.
- Use the `asadmin delete-mbean` command to undeploy an MBean.
- Use the `asadmin list-mbeans` command to list deployed MBeans.

To perform these tasks using the Admin Console, open the Custom MBeans component. For details, click the Help button in the Admin Console.

### The `asadmin create-mbean` Command

After installing the MBean classes as explained in “[MBean Class Loading](#)” on page 245, use the `asadmin create-mbean` command to deploy the MBean. This registers the MBean in the MBeanServer that is part of the Enterprise Server runtime environment. For more information about the MBeanServer, see “[The MBeanServer in the Enterprise Server](#)” on page 247.

Here is a simple example of an `asadmin create-mbean` command in which `TextPatterns` is the implementation class. The `--attributes` and `--target` options are not required.

```
asadmin create-mbean --user adminuser --target server1 --attributes color=red;font=Times TextPatterns
```

Other options not included in the example are as follows:

- `--name` defaults to the implementation class name
- `--objectname` is explained in [“The MBean Life Cycle” on page 244](#)
- `--enabled` defaults to `true` and is explained in [“Enabling and Disabling MBeans” on page 248](#)

All options must precede the implementation class.

For full details on the `asadmin create-mbean` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

For more information about MBean attributes, see [“Handling MBean Attributes” on page 248](#).

---

**Note** – To redeploy an MBean, simply install its new classes into the Enterprise Server as described in [“MBean Class Loading” on page 245](#). Then either restart the server or use `asadmin delete-mbean` followed by `asadmin create-mbean`.

---

## The `asadmin delete-mbean` Command

To undeploy an MBean, use the `asadmin delete-mbean` command. This removes its registration from the `MBeanServer`, but does not delete its code. Here is an example `asadmin delete-mbean` command in which `TextPatterns` is the implementation class. The `--target` option is not required.

```
asadmin delete-mbean --user adminuser --target server1 TextPatterns
```

For full details on the `asadmin delete-mbean` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## The `asadmin list-mbeans` Command

To list MBeans that have been deployed, use the `asadmin list-mbeans` command. Note that this command only lists the MBean definitions and not the MBeans registered in the `MBeanServer`. Here is an example `asadmin list-mbeans` command. The `--target` option is not required.

```
asadmin list-mbeans --user adminuser --target server1
```

The output of the `asadmin list-mbeans` command lists the following information:

- Implementation class – The name of the implementation class without the extension.

- Name – The name of the registered MBean, which defaults to but may be different from the implementation class name.
- Object name – The ObjectName of the MBean, which is explained in “[The MBean Life Cycle](#)” on page 244.
- Object type – For custom MBeans, the object type is always user. System MBeans have other object types.
- Enabled – Whether the MBean is enabled. MBeans are enabled by default. See “[Enabling and Disabling MBeans](#)” on page 248.

For full details on the `asadmin list-mbeans` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## The MBeanServer in the Enterprise Server

Custom MBeans are registered in the PlatformMBeanServer returned by the `java.lang.management.ManagementFactory.getPlatformMBeanServer()` method. This MBeanServer is associated with a standard JMX connector server.

You can use any JMX connector to look up MBeans in this MBeanServer just as you would any other MBeanServer. If your JMX connector is remote, you can connect to this MBeanServer using the following information:

- Host name of the Enterprise Server machine
- MBeanServer port, which is 8686 by default
- Administrator username
- Administrator password

For example, if you use JConsole, you can enter this information under the Remote tab. JConsole is a generic JMX connector you can use to look up and manage MBeans. For more information about JConsole, see <http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html>, the JMX tutorial at <http://java.sun.com/javase/6/docs/technotes/guides/jmx/tutorial/tutorialTOC.html>, and “Using JConsole” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

The connection to this MBeanServer is non-SSL by default for the developer profile and SSL by default for the cluster profile.

If SSL is enabled, you must provide the location of the truststore that contains the server certificate that the JMX connector should trust. For example, if you are using JConsole, you supply this location at the command line as follows:

```
jconsole -J-Djavax.net.ssl.trustStore=home-directory/.asadmintruststore
```

Look up the MBean by its name. By default, the name is the same as the implementation class.

You can reconfigure the JMX connector server's naming service port in one of the following ways:

- In the Admin Console, open the Admin Service component under the relevant configuration, select the system subcomponent, edit the Port field, and select Save. For details, click the Help button in the Admin Console.
- Use the `asadmin set` command as in the following example:

```
asadmin set --user adminuser server1.admin-service.jmx-connector.system.port=8687
```

For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Enabling and Disabling MBeans

A custom MBean is enabled by default. You can disable an MBean during deployment by using the `asadmin create-mbean` command's optional `--enabled=false` option. See “[The `asadmin create-mbean` Command](#)” on page 245.

After deployment, you can disable an MBean using the `asadmin set` command. For example:

```
asadmin set --user adminuser server1.applications.mbean.TextPatterns.enabled=false
```

If the MBean name is different from the implementation class, you must use the name in the `asadmin set` command. In this example, the name is `TextPatterns`.

For full details on the `asadmin set` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Handling MBean Attributes

You can set MBean attribute values that are not read-only in the following ways:

- In the MBean code itself, which does not affect the Enterprise Server configuration
- During deployment using the `asadmin create-mbean` command
- During deployment using the Custom MBeans component in the Admin Console
- Using the `asadmin set` command
- Using a JMX connector such as JConsole, which does not affect the Enterprise Server configuration

In the Enterprise Server configuration, MBean attributes are stored as properties. Therefore, using the `asadmin set` command means editing properties. For example:

```
asadmin set --user adminuser server1.applications.mbean.TextPatterns.property.color=blue
```



If the MBean name is different from the implementation class, you must use the MBean name in the `asadmin set` command. In this example, the name is `TextPatterns`.

For full details on the `asadmin set` command, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).



PART III

Using Services and APIs



# Using the JDBC API for Database Access

---

This chapter describes how to use the Java™ Database Connectivity (JDBC™) API for database access with the Sun GlassFish Enterprise Server. This chapter also provides high level JDBC implementation instructions for servlets and EJB components using the Enterprise Server. If the JDK version 1.6 is used, the Enterprise Server supports the JDBC 4.0 API, which encompasses the JDBC 3.0 API and the JDBC 2.0 Optional Package API.

The JDBC specifications are available at  
<http://java.sun.com/products/jdbc/download.html>.

A useful JDBC tutorial is located at  
<http://java.sun.com/docs/books/tutorial/jdbc/index.html>.

---

**Note** – The Enterprise Server does not support connection pooling or transactions for an application’s database access if it does not use standard Java EE DataSource objects.

---

This chapter discusses the following topics:

- “General Steps for Creating a JDBC Resource” on page 253
- “Creating Applications That Use the JDBC API” on page 255
- “Restrictions and Optimizations” on page 259

## General Steps for Creating a JDBC Resource

To prepare a JDBC resource for use in Java EE applications deployed to the Enterprise Server, perform the following tasks:

- “Integrating the JDBC Driver” on page 254
- “Creating a Connection Pool” on page 254
- “Testing a JDBC Connection Pool” on page 255
- “Creating a JDBC Resource” on page 255

For information about how to configure some specific JDBC drivers, see “Configurations for Specific JDBC Drivers” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

## Integrating the JDBC Driver

To use JDBC features, you must choose a JDBC driver to work with the Enterprise Server, then you must set up the driver. This section covers these topics:

- “Supported Database Drivers” on page 254
- “Making the JDBC Driver JAR Files Accessible” on page 254

### Supported Database Drivers

Supported JDBC drivers are those that have been fully tested by Sun. For a list of the JDBC drivers currently supported by the Enterprise Server, see the *Sun GlassFish Enterprise Server 2.1 Release Notes*. For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

**Note** – Because the drivers and databases supported by the Enterprise Server are constantly being updated, and because database vendors continue to upgrade their products, always check with Sun technical support for the latest database support information.

---

### Making the JDBC Driver JAR Files Accessible

To integrate the JDBC driver into a Enterprise Server domain, copy the JAR files into the *domain-dir/lib* directory, then restart the server. This makes classes accessible to all applications or modules deployed on servers that share the same configuration. For more information about Enterprise Server class loaders, see [Chapter 2, “Class Loaders.”](#)

## Creating a Connection Pool

When you create a connection pool that uses JDBC technology (a *JDBC connection pool*) in the Enterprise Server, you can define many of the characteristics of your database connections.

You can create a JDBC connection pool in one of these ways:

- In the Admin Console, open the Resources component, open the JDBC component, and select Connection Pools. For details, click the Help button in the Admin Console.
- Use the `asadmin create-jdbc-connection-pool` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

For a complete description of JDBC connection pool features, see the *Sun GlassFish Enterprise Server 2.1 Administration Guide*

## Testing a JDBC Connection Pool

You can test a JDBC connection pool for usability in one of these ways:

- In the Admin Console, open the Resources component, open the JDBC component, select Connection Pools, and select the connection pool you want to test. Then select the Ping button in the top right corner of the page. For details, click the Help button in the Admin Console.
- Use the `asadmin ping-connection-pool` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

Both these commands fail and display an error message unless they successfully connect to the connection pool.

For information about how to tune a connection pool, see the *Sun GlassFish Enterprise Server 2.1 Performance Tuning Guide*.

## Creating a JDBC Resource

A JDBC resource, also called a data source, lets you make connections to a database using `getConnection()`. Create a JDBC resource in one of these ways:

- In the Admin Console, open the Resources component, open the JDBC component, and select JDBC Resources. For details, click the Help button in the Admin Console.
- Use the `asadmin create-jdbc-resource` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Creating Applications That Use the JDBC API

An application that uses the JDBC API is an application that looks up and connects to one or more databases. This section covers these topics:

- “Sharing Connections” on page 255
- “Obtaining a Physical Connection From a Wrapped Connection” on page 256
- “Marking Bad Connections” on page 256
- “Using Non-Transactional Connections” on page 257
- “Using JDBC Transaction Isolation Levels” on page 258
- “Allowing Non-Component Callers” on page 259

## Sharing Connections

When multiple connections acquired by an application use the same JDBC resource, the connection pool provides connection sharing within the same transaction scope. For example,

suppose Bean A starts a transaction and obtains a connection, then calls a method in Bean B. If Bean B acquires a connection to the same JDBC resource with the same sign-on information, and if Bean A completes the transaction, the connection can be shared.

Connections obtained through a resource are shared only if the resource reference declared by the Java EE component allows it to be shareable. This is specified in a component's deployment descriptor by setting the `res-sharing-scope` element to `Shareable` for the particular resource reference. To turn off connection sharing, set `res-sharing-scope` to `Unshareable`.

For general information about connections and JDBC URLs, see [Chapter 3, "JDBC Resources,"](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

## Obtaining a Physical Connection From a Wrapped Connection

The `DataSource` implementation in the Enterprise Server provides a `getConnection` method that retrieves the JDBC driver's `SQLConnection` from the Enterprise Server's `Connection` wrapper. The method signature is as follows:

```
public java.sql.Connection getConnection(java.sql.Connection con)
throws java.sql.SQLException
```

For example:

```
InitialContext ctx = new InitialContext();
com.sun.appserv.jdbc.DataSource ds = (com.sun.appserv.jdbc.DataSource)
    ctx.lookup("jdbc/MyBase");
Connection con = ds.getConnection();
Connection drivercon = ds.getConnection(con);
// Do db operations.
// Do not close driver connection.
con.close(); // return wrapped connection to pool.
```

## Marking Bad Connections

The `DataSource` implementation in the Enterprise Server provides a `markConnectionAsBad` method. A marked bad connection is removed from its connection pool when it is closed. The method signature is as follows:

```
public void markConnectionAsBad(java.sql.Connection con)
```

For example:



```

com.sun.appserv.jdbc.DataSource ds=
    (com.sun.appserv.jdbc.DataSource)context.lookup("dataSource");
Connection con = ds.getConnection();
Statement stmt = null;
try{
    stmt = con.createStatement();
    stmt.executeUpdate("Update");
}
catch (BadConnectionException e){
    dataSource.markConnectionAsBad(con) //marking it as bad for removal
}
finally{
    stmt.close();
    con.close(); //Connection will be destroyed during close.
}

```

## Using Non-Transactional Connections

You can specify a non-transactional database connection in any of these ways:

- Check the Non-Transactional Connections box on the JDBC Connection Pools page in the Admin Console. The default is unchecked. For more information, click the Help button in the Admin Console.
- Specify the `--nontransactionalconnections` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.
- Use the `DataSource` implementation in the Enterprise Server, which provides a `getNonTxConnection` method. This method retrieves a JDBC connection that is not in the scope of any transaction. There are two variants.

```
public java.sql.Connection getNonTxConnection() throws java.sql.SQLException
```

```
public java.sql.Connection getNonTxConnection(String user, String password)
    throws java.sql.SQLException
```

- Create a resource with the JNDI name ending in `__nontx`. This forces all connections looked up using this resource to be non transactional.

Typically, a connection is enlisted in the context of the transaction in which a `getConnection` call is invoked. However, a non-transactional connection is not enlisted in a transaction context even if a transaction is in progress.

The main advantage of using non-transactional connections is that the overhead incurred in enlisting and delisting connections in transaction contexts is avoided. However, use such connections carefully. For example, if a non-transactional connection is used to query the database while a transaction is in progress that modifies the database, the query retrieves the

unmodified data in the database. This is because the in-progress transaction hasn't committed. For another example, if a non-transactional connection modifies the database and a transaction that is running simultaneously rolls back, the changes made by the non-transactional connection are not rolled back.

Here is a typical use case for a non-transactional connection: a component that is updating a database in a transaction context spanning over several iterations of a loop can refresh cached data by using a non-transactional connection to read data before the transaction commits.

## Using JDBC Transaction Isolation Levels

For general information about transactions, see [Chapter 16, "Using the Transaction Service,"](#) and [Chapter 12, "Transactions,"](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide*. For information about last agent optimization, which can improve performance, see ["Transaction Scope" on page 262](#).

Not all database vendors support all transaction isolation levels available in the JDBC API. The Enterprise Server permits specifying any isolation level your database supports. The following table defines transaction isolation levels.

TABLE 15-1 Transaction Isolation Levels

Transaction Isolation Level	Description
TRANSACTION_READ_UNCOMMITTED	Dirty reads, non-repeatable reads, and phantom reads can occur.
TRANSACTION_READ_COMMITTED	Dirty reads are prevented; non-repeatable reads and phantom reads can occur.
TRANSACTION_REPEATABLE_READ	Dirty reads and non-repeatable reads are prevented; phantom reads can occur.
TRANSACTION_SERIALIZABLE	Dirty reads, non-repeatable reads and phantom reads are prevented.

Note that you cannot call `setTransactionIsolation()` during a transaction.

You can set the default transaction isolation level for a JDBC connection pool. For details, see ["Creating a Connection Pool" on page 254](#).

To verify that a level is supported by your database management system, test your database programmatically using the `supportsTransactionIsolationLevel()` method in `java.sql.DatabaseMetaData`, as shown in the following example:

```
InitialContext ctx = new InitialContext();
DataSource ds = (DataSource)
ctx.lookup("jdbc/MyBase");
Connection con = ds.getConnection();
DatabaseMetaData dbmd = con.getMetaData();
if (dbmd.supportsTransactionIsolationLevel(TRANSACTION_SERIALIZABLE))
{ Connection.setTransactionIsolation(TRANSACTION_SERIALIZABLE); }
```

For more information about these isolation levels and what they mean, see the JDBC API specification.

---

**Note** – Applications that change the isolation level on a pooled connection programmatically risk polluting the pool, which can lead to errors.

---

## Allowing Non-Component Callers

You can allow non-Java-EE components, such as servlet filters, lifecycle modules, and third party persistence managers, to use this JDBC connection pool. The returned connection is automatically enlisted with the transaction context obtained from the transaction manager. Standard Java EE components can also use such pools. Connections obtained by non-component callers are not automatically closed at the end of a transaction by the container. They must be explicitly closed by the caller.

You can enable non-component callers in the following ways:

- Check the Allow Non Component Callers box on the JDBC Connection Pools page in the Admin Console. The default is `false`. For more information, click the Help button in the Admin Console.
- Specify the `--allownoncomponentcallers` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).
- Create a JDBC resource with a `__pm` suffix.

## Restrictions and Optimizations

This section discusses restrictions and performance optimizations that affect using the JDBC API.

### Disabling Stored Procedure Creation on Sybase

By default, DataDirect and Sun GlassFish JDBC drivers for Sybase databases create a stored procedure for each parameterized `PreparedStatement`. On the Enterprise Server, exceptions are thrown when primary key identity generation is attempted. To disable the creation of these stored procedures, set the property `PrepareMethod=direct`.



# Using the Transaction Service

---

The Java EE platform provides several abstractions that simplify development of dependable transaction processing for applications. This chapter discusses Java EE transactions and transaction support in the Sun GlassFish Enterprise Server.

This chapter contains the following sections:

- “Transaction Resource Managers” on page 261
- “Transaction Scope” on page 262
- “Distributed Transaction Recovery” on page 263
- “Configuring the Transaction Service” on page 264
- “The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction” on page 264
- “Transaction Logging” on page 265
- “Storing Transaction Logs in a Database” on page 265
- “Recovery Workarounds” on page 266

For more information about the Java™ Transaction API (JTA) and Java Transaction Service (JTS), see Chapter 12, “Transactions,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide* and the following sites: <http://java.sun.com/products/jta/> and <http://java.sun.com/products/jts/>.

You might also want to read “Chapter 35: Transactions” in the *Java EE 5 Tutorial* (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

## Transaction Resource Managers

There are three types of transaction resource managers:

- Databases - Use of transactions prevents databases from being left in inconsistent states due to incomplete updates. For information about JDBC transaction isolation levels, see “Using JDBC Transaction Isolation Levels” on page 258.

The Enterprise Server supports a variety of JDBC XA drivers. For a list of the JDBC drivers currently supported by the Enterprise Server, see the [Sun GlassFish Enterprise Server 2.1 Release Notes](#). For configurations of supported and other drivers, see “Configurations for Specific JDBC Drivers” in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

- Java Message Service (JMS) Providers - Use of transactions ensures that messages are reliably delivered. The Enterprise Server is integrated with Sun GlassFish Message Queue, a fully capable JMS provider. For more information about transactions and the JMS API, see [Chapter 18, “Using the Java Message Service.”](#)
- J2EE Connector Architecture (CA) components - Use of transactions prevents legacy EIS systems from being left in inconsistent states due to incomplete updates. For more information about connectors, see [Chapter 12, “Developing Connectors.”](#)

For details about how transaction resource managers, the transaction service, and applications interact, see [Chapter 12, “Transactions,”](#) in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

## Transaction Scope

A *local* transaction involves only one non-XA resource and requires that all participating application components execute within one process. Local transaction optimization is specific to the resource manager and is transparent to the Java EE application.

In the Enterprise Server, a JDBC resource is non-XA if it meets any of the following criteria:

- In the JDBC connection pool configuration, the DataSource class does not implement the `javax.sql.XADataSource` interface.
- The Global Transaction Support box is not checked, or the Resource Type setting does not exist or is not set to `javax.sql.XADataSource`.

A transaction remains local if the following conditions remain true:

- One and only one non-XA resource is used. If any additional non-XA resource is used, the transaction is aborted.
- No transaction importing or exporting occurs.

Transactions that involve multiple resources or multiple participant processes are *distributed* or *global* transactions. A global transaction can involve one non-XA resource if last agent optimization is enabled. Otherwise, all resources must be XA. The `use-last-agent-optimization` property is set to `true` by default. For details about how to set this property, see “[Configuring the Transaction Service](#)” on page 264.

If only one XA resource is used in a transaction, one-phase commit occurs, otherwise the transaction is coordinated with a two-phase commit protocol.

A two-phase commit protocol between the transaction manager and all the resources enlisted for a transaction ensures that either all the resource managers commit the transaction or they all abort. When the application requests the commitment of a transaction, the transaction manager issues a `PREPARE_TO_COMMIT` request to all the resource managers involved. Each of these resources can in turn send a reply indicating whether it is ready for commit (`PREPARED`) or not (`NO`). Only when all the resource managers are ready for a commit does the transaction manager issue a commit request (`COMMIT`) to all the resource managers. Otherwise, the transaction manager issues a rollback request (`ABORT`) and the transaction is rolled back.

## Distributed Transaction Recovery

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see [“Usage Profiles” in \*Sun GlassFish Enterprise Server 2.1 Administration Guide\*](#).

---

To enable cluster-wide automatic recovery, you must first facilitate storing of transaction logs in a shared file system. You can do this in one of these ways:

- Set the Enterprise Server's `log-root` attribute to a shared file system base directory and set the transaction service's `tx-log-dir` attribute to a relative path.
- Set `tx-log-dir` to an absolute path to a shared file system directory, in which case `log-root` is ignored for transaction logs.
- Set a system-property called `TX-LOG-DIR` in the `domain.xml` file to a shared file system directory.

```
<server config-ref="server-config" name="server">
  <system-property name="TX-LOG-DIR"
    value="/net/tulsa/nodeagents/na/instance1/logs" />
</server>
```

Next, you must set the transaction service's `delegated-recovery` property to `true` (the default is `false`).

For information about setting `tx-log-dir` and `delegated-recovery`, see [“Configuring the Transaction Service” on page 264](#). For information about setting `log-root` and other general logging settings, see [Chapter 17, “Configuring Logging” in \*Sun GlassFish Enterprise Server 2.1 Administration Guide\*](#). For information about system-property and the `domain.xml` file, see the [Sun GlassFish Enterprise Server 2.1 Administration Reference](#).

## Configuring the Transaction Service

You can configure the transaction service in the Enterprise Server in the following ways:

- To configure the transaction service using the Admin Console, open the Transaction Service component under the relevant configuration. For details, click the Help button in the Admin Console.
- To configure the transaction service, use the `asadmin set` command to set the following attributes.

```
server.transaction-service.automatic-recovery = false
server.transaction-service.heuristic-decision = rollback
server.transaction-service.keypoint-interval = 2048
server.transaction-service.retry-timeout-in-seconds = 600
server.transaction-service.timeout-in-seconds = 0
server.transaction-service.tx-log-dir = domain-dir/logs
```

You can also set these properties:

```
server.transaction-service.property.oracle-xa-recovery-workaround = false
server.transaction-service.property.disable-distributed-transaction-logging = false
server.transaction-service.property.xaresource-txn-timeout = 600
server.transaction-service.property.pending-txn-cleanup-interval = 60
server.transaction-service.property.use-last-agent-optimization = true
server.transaction-service.property.db-logging-resource = jdbc/TxnDS
server.transaction-service.property.delegated-recovery = false
server.transaction-service.property.wait-time-before-recovery-insec = 60
server.transaction-service.property.xa-servername = myserver
```

You can use the `asadmin get` command to list all the transaction service attributes and properties. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction

You can access the Enterprise Server transaction manager, a `javax.transaction.TransactionManager` implementation, using the JNDI subcontext `java:comp/TransactionManager` or `java:appserver/TransactionManager`. You can access the Enterprise Server transaction synchronization registry, a `javax.transaction.TransactionSynchronizationRegistry` implementation, using the JNDI subcontext `java:comp/TransactionSynchronizationRegistry` or `java:appserver/TransactionSynchronizationRegistry`. You can also request injection of a `TransactionManager` or `TransactionSynchronizationRegistry` object using the `@Resource` annotation. Accessing the transaction synchronization registry is recommended. For details, see [Java Specification Request \(JSR\) 907 \(http://www.jcp.org/en/jsr/detail?id=907\)](http://www.jcp.org/en/jsr/detail?id=907).



You can also access `java:comp/UserTransaction`.

## Transaction Logging

The transaction service writes transactional activity into transaction logs so that transactions can be recovered. You can control transaction logging in these ways:

- Set the location of the transaction log files using the Transaction Log Location setting in the Admin Console, or set the `tx-log-dir` attribute using the `asadmin set` command.
- Turn off transaction logging by setting the `disable-distributed-transaction-logging` property to `true` and the `automatic-recovery` attribute to `false`. Do this *only* if performance is more important than transaction recovery.

## Storing Transaction Logs in a Database

For multi-core machines, logging transactions to a database may be more efficient.

To log transactions to a database, follow these steps:

1. Create a JDBC connection Pool, and set the `non-transactional-connections` attribute to `true`.
2. Create a JDBC resource that uses the connection pool and note the JNDI name of the JDBC resource.
3. Create a table named `txn_log_table` with the schema shown in [Table 16–1](#).
4. Add the `db-logging-resource` property to the transaction service. For example:

```
asadmin set --user adminuser server1.transaction-service.property.db-logging-resource="jdbc/TxnDS"
```

The property's value should be the JNDI name of the JDBC resource configured previously.

5. To disable file synchronization, use the following `asadmin create-jvm-options` command:

```
asadmin create-jvm-options --user adminuser -Dcom.sun.appserv.transaction.nofdsync
```

6. Restart the server.

For information about JDBC connection pools and resources, see [Chapter 15, “Using the JDBC API for Database Access.”](#) For more information about the `asadmin create-jvm-options` command, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

TABLE 16-1 Schema for txn\_log\_table

Column Name	JDBC Type
LOCALTID	BIGINT
SERVERNAME	VARCHAR(n)
GTRID	VARBINARY

The size of the SERVERNAME column should be at least the length of the Enterprise Server host name plus 10 characters.

The size of the GTRID column should be at least 64 bytes.

To define the SQL used by the transaction manager when it is storing its transaction logs in the database, use the following flags:

```
-Dcom.sun.jts.dblogging.insertquery=sql statement
-Dcom.sun.jts.dblogging.deletequery=sql statement
```

The default statements are as follows:

```
-Dcom.sun.jts.dblogging.insertquery=insert into txn_log_table values ( ?, ? , ? )
-Dcom.sun.jts.dblogging.deletequery=delete from txn_log_table where localtid = ? and servername = ?
```

To set one of these flags using the `asadmin create-jvm-options` command, you must quote the statement. For example:

```
create-jvm-options '-Dcom.sun.jts.dblogging.deletequery=delete from txn_log_table where gtrid = ?'
```

You can also set JVM options in the Admin Console. In the developer profile, select the Application Server component and the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. These flags and their statements must also be quoted in the Admin Console. For example:

```
'-Dcom.sun.jts.dblogging.deletequery=delete from txn_log_table where gtrid = ?'
```

## Recovery Workarounds

The Enterprise Server provides workarounds for some known issues with the recovery implementations of the following JDBC drivers. These workarounds are used unless explicitly disabled.

In the Oracle thin driver, the `XAResource.recover` method repeatedly returns the same set of in-doubt Xids regardless of the input flag. According to the XA specifications, the Transaction Manager initially calls this method with `TMSTARTSCAN` and then with `TMNOFLAGS` repeatedly until no Xids are returned. The `XAResource.commit` method also has some issues.

To disable the Enterprise Server workaround, set the `oracle-xa-recovery-workaround` property value to `false`. For details about how to set this property, see [“Configuring the Transaction Service” on page 264](#).

---

**Note** – These workarounds do not imply support for any particular JDBC driver.

---



# Using the Java Naming and Directory Interface

---

A *naming service* maintains a set of bindings, which relate names to objects. The Java EE naming service is based on the Java Naming and Directory Interface™ (JNDI) API. The JNDI API allows application components and clients to look up distributed resources, services, and EJB components. For general information about the JNDI API, see <http://java.sun.com/products/jndi/>.

You can also see the JNDI tutorial at <http://java.sun.com/products/jndi/tutorial/>.

This chapter contains the following sections:

- “Accessing the Naming Context” on page 269
- “Configuring Resources” on page 272
- “Using a Custom `jndi.properties` File” on page 273
- “Mapping References” on page 274

## Accessing the Naming Context

The Enterprise Server provides a naming environment, or *context*, which is compliant with standard Java EE requirements. A `Context` object provides the methods for binding names to objects, unbinding names from objects, renaming objects, and listing the bindings. The `InitialContext` is the handle to the Java EE naming service that application components and clients use for lookups.

The JNDI API also provides subcontext functionality. Much like a directory in a file system, a subcontext is a context within a context. This hierarchical structure permits better organization of information. For naming services that support subcontexts, the `Context` class also provides methods for creating and destroying subcontexts.

The rest of this section covers these topics:

- “Global JNDI Names” on page 270
- “Accessing EJB Components Using the `CosNaming` Naming Context” on page 271

- “Accessing EJB Components in a Remote Application Server” on page 271
- “Naming Environment for Lifecycle Modules” on page 272

---

**Note** – Each resource within a server instance must have a unique name. However, two resources in different server instances or different domains can have the same name.

---

## Global JNDI Names

Global JNDI names are assigned according to the following precedence rules:

1. A global JNDI name assigned in the `sun-ejb-jar.xml`, `sun-web.xml`, or `sun-application-client.xml` deployment descriptor file has the highest precedence. See “Mapping References” on page 274.
2. A global JNDI name assigned in a `mapped-name` element in the `ejb-jar.xml`, `web.xml`, or `application-client.xml` deployment descriptor file has the second highest precedence. The following elements have `mapped-name` subelements: `resource-ref`, `resource-env-ref`, `ejb-ref`, `message-destination`, `message-destination-ref`, `session`, `message-driven`, and `entity`.
3. A global JNDI name assigned in a `mappedName` attribute of an annotation has the third highest precedence. The following annotations have `mappedName` attributes: `@javax.annotation.Resource`, `@javax.ejb.EJB`, `@javax.ejb.Stateless`, `@javax.ejb.Stateful`, and `@javax.ejb.MessageDriven`.
4. A default global JNDI name is assigned in some cases if no name is assigned in deployment descriptors or annotations.
  - For an EJB 2.x dependency or a session or entity bean with a remote interface, the default is the fully qualified name of the home interface.
  - For an EJB 3.0 dependency or a session bean with a remote interface, the default is the fully qualified name of the remote business interface.
  - If both EJB 2.x and EJB 3.0 remote interfaces are specified, or if more than one 3.0 remote interface is specified, there is no default, and the global JNDI name must be specified.
  - For all other component dependencies that must be mapped to global JNDI names, the default is the name of the dependency relative to `java:comp/env`. For example, in the `@Resource(name="jdbc/Foo") DataSource ds;` annotation, the global JNDI name is `jdbc/Foo`.

## Accessing EJB Components Using the CosNaming Naming Context

The preferred way of accessing the naming service, even in code that runs outside of a Java EE container, is to use the no-argument `InitialContext` constructor. However, if EJB client code explicitly instantiates an `InitialContext` that points to the `CosNaming` naming service, it is necessary to set the `java.naming.factory.initial` property to `com.sun.jndi.cosnaming.CNCTXFactory` in the client JVM when accessing EJB components. You can set this property as a command-line argument, as follows:

```
-Djava.naming.factory.initial=com.sun.jndi.cosnaming.CNCTXFactory
```

Or you can set this property in the code, as follows:

```
Properties properties = null;
try {
    properties = new Properties();
    properties.put("java.naming.factory.initial",
        "com.sun.jndi.cosnaming.CNCTXFactory");
    ...
}
```

The `java.naming.factory.initial` property applies to only one instance; it is not cluster-aware.

## Accessing EJB Components in a Remote Application Server

The recommended approach for looking up an EJB component in a remote Enterprise Server from a client that is a servlet or EJB component is to use the Interoperable Naming Service syntax. Host and port information is prepended to any global JNDI names and is automatically resolved during the lookup. The syntax for an interoperable global name is as follows:

```
corbaname:iiop:host:port#a/b/name
```

This makes the programming model for accessing EJB components in another Enterprise Server exactly the same as accessing them in the same server. The deployer can change the way the EJB components are physically distributed without having to change the code.

For Java EE components, the code still performs a `java:comp/env` lookup on an EJB reference. The only difference is that the deployer maps the `ejb-reference` element to an interoperable name in an Enterprise Server deployment descriptor file instead of to a simple global JNDI name.

For example, suppose a servlet looks up an EJB reference using `java:comp/env/ejb/Foo`, and the target EJB component has a global JNDI name of `a/b/Foo`.

The `ejb-ref` element in `sun-web.xml` looks like this:

```
<ejb-ref>
  <ejb-ref-name>ejb/Foo</ejb-ref-name>
  <jndi-name>corbaname:iiop:host:port#a/b/Foo</jndi-name>
</ejb-ref>
```

The code looks like this:

```
Context ic = new InitialContext();
Object o = ic.lookup("java:comp/env/ejb/Foo");
```

For a client that doesn't run within a Java EE container, the code just uses the interoperable global name instead of the simple global JNDI name. For example:

```
Context ic = new InitialContext();
Object o = ic.lookup("corbaname:iiop:host:port#a/b/Foo");
```

Objects stored in the interoperable naming context and component-specific (`java:comp/env`) naming contexts are transient. On each server startup or application reloading, all relevant objects are re-bound to the namespace.

## Naming Environment for Lifecycle Modules

Lifecycle listener modules provide a means of running short or long duration tasks based on Java technology within the application server environment, such as instantiation of singletons or RMI servers. These modules are automatically initiated at server startup and are notified at various phases of the server life cycle. For details about lifecycle modules, see [Chapter 13, “Developing Lifecycle Listeners.”](#)

The configured properties for a lifecycle module are passed as properties during server initialization (the `INIT_EVENT`). The initial JNDI naming context is not available until server initialization is complete. A lifecycle module can get the `InitialContext` for lookups using the method `LifecycleEventContext.getInitialContext()` during, and only during, the `STARTUP_EVENT`, `READY_EVENT`, or `SHUTDOWN_EVENT` server life cycle events.

## Configuring Resources

The Enterprise Server exposes the following special resources in the naming environment. Full administration details are provided in the following sections:

- “External JNDI Resources” on page 273
- “Custom Resources” on page 273



## External JNDI Resources

An external JNDI resource defines custom JNDI contexts and implements the `javax.naming.spi.InitialContextFactory` interface. There is no specific JNDI parent context for external JNDI resources, except for the standard `java:comp/env/`.

Create an external JNDI resource in one of these ways:

- To create an external JNDI resource using the Admin Console, open the Resources component, open the JNDI component, and select External Resources. For details, click the Help button in the Admin Console.
- To create an external JNDI resource, use the `asadmin create-jndi-resource` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Custom Resources

A custom resource specifies a custom server-wide resource object factory that implements the `javax.naming.spi.ObjectFactory` interface. There is no specific JNDI parent context for external JNDI resources, except for the standard `java:comp/env/`.

Create a custom resource in one of these ways:

- To create a custom resource using the Admin Console, open the Resources component, open the JNDI component, and select Custom Resources. For details, click the Help button in the Admin Console.
- To create a custom resource, use the `asadmin create-custom-resource` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## Using a Custom `jndi.properties` File

To use a custom `jndi.properties` file, specify the path to the file in one of the following ways:

- Use the Admin Console. In the developer profile, select the Enterprise Server component and select the JVM Settings tab. In the cluster profile, select the JVM Settings component under the relevant configuration. Then select the Path Settings tab and edit the Classpath Prefix field. For details, click the Help button in the Admin Console.
- Edit the `classpath-prefix` attribute of the `java-config` element in the `domain.xml` file. For details about `domain.xml`, see the *Sun GlassFish Enterprise Server 2.1 Administration Reference*.

This adds the `jndi.properties` file to the Shared Chain class loader. For more information about class loading, see [Chapter 2, “Class Loaders.”](#)

For each property found in more than one `jndi.properties` file, the Java EE naming service either uses the first value found or concatenates all of the values, whichever makes sense.

## Mapping References

The following XML elements in the Enterprise Server deployment descriptors map resource references in application client, EJB, and web application components to JNDI names configured in the Enterprise Server:

- `resource-env-ref` - Maps the `@Resource` or `@Resources` annotation (or the `resource-env-ref` element in the corresponding Java EE XML file) to the absolute JNDI name configured in the Enterprise Server.
- `resource-ref` - Maps the `@Resource` or `@Resources` annotation (or the `resource-ref` element in the corresponding Java EE XML file) to the absolute JNDI name configured in the Enterprise Server.
- `ejb-ref` - Maps the `@EJB` annotation (or the `ejb-ref` element in the corresponding Java EE XML file) to the absolute JNDI name configured in the Enterprise Server.

JNDI names for EJB components must be unique. For example, appending the application name and the module name to the EJB name is one way to guarantee unique names. In this case, `mycompany.pkging.pkgingEJB.MyEJB` would be the JNDI name for an EJB in the module `pkgingEJB.jar`, which is packaged in the `pkging.ear` application.

These elements are part of the `sun-web.xml`, `sun-ejb-ref.xml`, and `sun-application-client.xml` deployment descriptor files. For more information about how these elements behave in each of the deployment descriptor files, see [Appendix A, “Deployment Descriptor Files,”](#) in *Sun GlassFish Enterprise Server 2.1 Application Deployment Guide*.

The rest of this section uses an example of a JDBC resource lookup to describe how to reference resource factories. The same principle is applicable to all resources (such as JMS destinations, JavaMail sessions, and so on).

The `@Resource` annotation in the application code looks like this:

```
@Resource(name="jdbc/helloDbDs") javax.sql.DataSource ds;
```

This references a resource with the JNDI name of `java:comp/env/jdbc/helloDbDs`. If this is the JNDI name of the JDBC resource configured in the Enterprise Server, the annotation alone is enough to reference the resource.

However, you can use an Enterprise Server specific deployment descriptor to override the annotation. For example, the `resource-ref` element in the `sun-web.xml` file maps the `res-ref-name` (the name specified in the annotation) to the JNDI name of another JDBC resource configured in the Enterprise Server.

```
<resource-ref>
  <res-ref-name>jdbc/helloDbDs</res-ref-name>
  <jndi-name>jdbc/helloDbDataSource</jndi-name>
</resource-ref>
```

# Using the Java Message Service

---

This chapter describes how to use the Java™ Message Service (JMS) API. The Sun Java System Enterprise Server has a fully integrated JMS provider: the Sun Java System Message Queue software.

For general information about the JMS API, see “Chapter 31: The Java Message Service API” in the *Java EE 5 Tutorial* (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

For detailed information about JMS concepts and JMS support in the Enterprise Server, see Chapter 4, “Configuring Java Message Service Resources,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

This chapter contains the following sections:

- “The JMS Provider” on page 276
- “Message Queue Resource Adapter” on page 277
- “Generic Resource Adapter” on page 277
- “Administration of the JMS Service” on page 277
- “Restarting the JMS Client After JMS Configuration” on page 281
- “JMS Connection Features” on page 281
- “Load-Balanced Message Inflow” on page 282
- “JMS Service High Availability” on page 283
- “Transactions and Non-Persistent Messages” on page 284
- “Authentication With ConnectionFactory” on page 284
- “Message Queue varhome Directory” on page 284
- “Delivering SOAP Messages Using the JMS API” on page 285

## The JMS Provider

The Enterprise Server support for JMS messaging, in general, and for message-driven beans, in particular, requires messaging middleware that implements the JMS specification: a JMS provider. The Enterprise Server uses the Sun GlassFish Message Queue software as its native JMS provider. The Message Queue software is tightly integrated into the Enterprise Server, providing transparent JMS messaging support. This support is known within Enterprise Server as the *JMS Service*. The JMS Service requires only minimal administration.

The relationship of the Message Queue software to the Enterprise Server can be one of these types: EMBEDDED, LOCAL, or REMOTE. The effects of these choices on the Message Queue broker life cycle are as follows:

- If the type is EMBEDDED, the Enterprise Server and Message Queue software run in the same JVM, and the networking stack is bypassed. The Message Queue broker is started and stopped automatically by the Enterprise Server. This is the default for the Domain Administration Server (DAS).

Lazy initialization starts the default embedded broker on the first access of JMS services rather than at Enterprise Server startup. EMBEDDED mode is not a supported configuration for a cluster.

- If the type is LOCAL, the Message Queue broker starts when the Enterprise Server starts. This is the default for all Enterprise Server instances except the DAS.

The LOCAL setting implicitly sets up a 1:1 relationship between an Enterprise Server instance and a Message Queue broker. When you create an Enterprise Server cluster, a Message Queue cluster is automatically created as well. During cluster creation, each instance in the Enterprise Server cluster is automatically configured with a broker in the Message Queue cluster, and a unique broker port is determined.

The first Enterprise Server instance's Message Queue broker is set as the master broker. If you delete the first Enterprise Server instance, you must use Message Queue administration tools to migrate the master broker. For details, see “[Managing a Conventional Cluster's Configuration Change Record](#)” in *Sun Java System Message Queue 4.3 Administration Guide*. However, if HADB is installed, there is no master broker. For more information about HADB, see “[JMS Service High Availability](#)” on page 283.

- If the type is REMOTE, the Message Queue broker must be started separately. For information about starting the broker, see the *Sun Java System Message Queue 4.3 Administration Guide*.

For more information about setting the type and the default JMS host, see “[Configuring the JMS Service](#)” on page 278.

For more information about the Message Queue software, refer to the documentation at <http://docs.sun.com/coll/1343.8>.

For general information about the JMS API, see the JMS web page at <http://java.sun.com/products/jms/index.html>.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

## Message Queue Resource Adapter

The Sun GlassFish Message Queue software is integrated into the Enterprise Server using a resource adapter that is compliant with the Connector specification. The module name of this system resource adapter is `jmsra`. Every JMS resource is converted to a corresponding connector resource of this resource adapter as follows:

- **Connection Factory** – A connector connection pool with a `max-pool-size` of 250 and a corresponding connector resource
- **Destination (Topic or Queue)** – A connector administered object

You use connector configuration tools to manage JMS resources. For more information, see Chapter 12, “Developing Connectors.”

## Generic Resource Adapter

The Enterprise Server provides a generic resource adapter for JMS, for those who want to use a JMS provider other than Sun GlassFish Message Queue. For details, see <http://genericjmsra.dev.java.net/> and “Configuring the Generic Resource Adapter for JMS” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

## Administration of the JMS Service

To configure the JMS Service and prepare JMS resources for use in applications deployed to the Enterprise Server, you must perform these tasks:

- “Configuring the JMS Service” on page 278
- “The Default JMS Host” on page 278
- “Creating JMS Hosts” on page 279
- “Checking Whether the JMS Provider Is Running” on page 279
- “Creating Physical Destinations” on page 279
- “Creating JMS Resources: Destinations and Connection Factories” on page 280

For more information about JMS administration tasks, see Chapter 4, “Configuring Java Message Service Resources,” in *Sun GlassFish Enterprise Server 2.1 Administration Guide* and the *Sun Java System Message Queue 4.3 Administration Guide*.

## Configuring the JMS Service

The JMS Service configuration is available to all inbound and outbound connections pertaining to the Enterprise Server cluster or instance. You can edit the JMS Service configuration in the following ways:

- To edit the JMS Service configuration using the Admin Console, open the Java Message Service component under the relevant configuration. For details, click the Help button in the Admin Console.
- To configure the JMS service, use the `asadmin set` command to set the following attributes:

```
server.jms-service.init-timeout-in-seconds = 60
server.jms-service.type = EMBEDDED
server.jms-service.start-args =
server.jms-service.default-jms-host = default_JMS_host
server.jms-service.reconnect-interval-in-seconds = 60
server.jms-service.reconnect-attempts = 3
server.jms-service.reconnect-enabled = true
server.jms-service.addresslist-behavior = random
server.jms-service.addresslist-iterations = 3
server.jms-service.mq-scheme = mq
server.jms-service.mq-service = jms
```

You can also set these properties:

```
server.jms-service.property.instance-name = imqbroker
server.jms-service.property.instance-name-suffix =
server.jms-service.property.append-version = false
server.jms-service.property.user-name =
server.jms-service.property.password =
```

You can use the `asadmin get` command to list all the JMS service attributes and properties. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

You can override the JMS Service configuration using JMS connection factory settings. For details, see [Chapter 4, “Configuring Java Message Service Resources,”](#) in [Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

---

**Note** – The Enterprise Server instance must be restarted after configuration of the JMS Service.

---

## The Default JMS Host

A JMS host refers to a Sun GlassFish Message Queue broker. A default JMS host for the JMS service is provided, named `default_JMS_host`. This is the JMS host that the Enterprise Server uses for performing all Message Queue broker administrative operations, such as creating and deleting JMS destinations.

If you have created a multi-broker cluster in the Message Queue software, delete the default JMS host, then add the Message Queue cluster's brokers as JMS hosts. In this case, the default JMS host becomes the first JMS host in the `AddressList`. For more information about the `AddressList`, see [“JMS Connection Features” on page 281](#). You can also explicitly set the default JMS host; see [“Configuring the JMS Service” on page 278](#).

When the Enterprise Server uses a Message Queue cluster, it executes Message Queue specific commands on the default JMS host. For example, when a physical destination is created for a Message Queue cluster of three brokers, the command to create the physical destination is executed on the default JMS host, but the physical destination is used by all three brokers in the cluster.

## Creating JMS Hosts

You can create additional JMS hosts in the following ways:

- Use the Admin Console. Open the Java Message Service component under the relevant configuration, then select the JMS Hosts component. For details, click the Help button in the Admin Console.
- Use the `asadmin create-jms-host` command. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

For machines having more than one host, use the Host field in the Admin Console or the `--mqhost` option of `create-jms-host` to specify the address to which the broker binds.

## Checking Whether the JMS Provider Is Running

You can use the `asadmin jms-ping` command to check whether a Sun GlassFish Message Queue instance is running. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

## Creating Physical Destinations

Produced messages are delivered for routing and subsequent delivery to consumers using *physical destinations* in the JMS provider. A physical destination is identified and encapsulated by an administered object (a `Topic` or `Queue` destination resource) that an application component uses to specify the destination of messages it is producing and the source of messages it is consuming.

If a message-driven bean is deployed and the `Queue` physical destination it listens to doesn't exist, the Enterprise Server automatically creates the physical destination and sets the value of the property `maxNumActiveConsumers` to `-1` (see [“Load-Balanced Message Inflow” on page 282](#)). However, it is good practice to create the `Queue` physical destination beforehand.

You can create a JMS physical destination in the following ways:

- Use the Admin Console. Open the Resources component, open the JMS Resources component, then select Physical Destinations. For details, click the Help button in the Admin Console.
- Use the `asadmin create-jmsdest` command. This command acts on the default JMS host of its target. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

To purge all messages currently queued at a physical destination, use the `asadmin flush-jmsdest` command. This deletes the messages before they reach any message consumers. For details, see the [Sun GlassFish Enterprise Server 2.1 Reference Manual](#).

To create a destination resource, see “Creating JMS Resources: Destinations and Connection Factories” on page 280.

## Creating JMS Resources: Destinations and Connection Factories

You can create two kinds of JMS resources in the Enterprise Server:

- **Connection Factories** – administered objects that implement the `ConnectionFactory`, `QueueConnectionFactory`, or `TopicConnectionFactory` interfaces.
- **Destination Resources** – administered objects that implement the `Queue` or `Topic` interfaces.

In either case, the steps for creating a JMS resource are the same. You can create a JMS resource in the following ways:

- To create a JMS resource using the Admin Console, open the Resources component, then open the JMS Resources component. Click Connection Factories to create a connection factory, or click Destination Resources to create a queue or topic. For details, click the Help button in the Admin Console.
- A JMS resource is a type of connector. To create a JMS resource using the command line, see “Deploying and Configuring a Stand-Alone Connector Module” on page 229.

---

**Note** – All JMS resource properties that used to work with version 7 of the Enterprise Server are supported for backward compatibility.

---



## Restarting the JMS Client After JMS Configuration

When a JMS client accesses a JMS administered object for the first time, the client JVM retrieves the JMS service configuration from the Enterprise Server. Further changes to the configuration are not available to the client JVM until the client is restarted.

## JMS Connection Features

The Sun GlassFish Message Queue software supports the following JMS connection features:

- [“Connection Pooling” on page 281](#)
- [“Connection Failover” on page 282](#)

Both these features use the `AddressList` configuration, which is populated with the hosts and ports of the JMS hosts defined in the Enterprise Server. The `AddressList` is updated whenever a JMS host configuration changes. The `AddressList` is inherited by any JMS resource when it is created and by any MDB when it is deployed.

---

**Note** – In the Sun GlassFish Message Queue software, the `AddressList` property is called `imqAddressList`.

---

## Connection Pooling

The Enterprise Server pools JMS connections automatically.

To dynamically modify connection pool properties using the Admin Console, go to either the Connection Factories page (see [“Creating JMS Resources: Destinations and Connection Factories” on page 280](#)) or the Connector Connection Pools page (see [“Deploying and Configuring a Stand-Alone Connector Module” on page 229](#)).

To use the command line, use the `asadmin create-connector-connection-pool` command to manage the pool (see [“Deploying and Configuring a Stand-Alone Connector Module” on page 229](#)).

For the developer profile, the `addresslist-behavior` JMS service attribute is set to `random` by default. This means that each `ManagedConnection` (physical connection) created from the `ManagedConnectionFactory` selects its primary broker in a random way from the `AddressList`.

For the cluster and enterprise profiles, the `addresslist-behavior` JMS service attribute is set to `priority` by default. This means that the first broker in the `AddressList` is selected first. This first broker is the local colocated Message Queue broker. If this broker is unavailable, connection attempts are made to brokers in the order in which they are listed in the `AddressList`. This ensures colocated production and consumption of messages and equitable load distribution across the Message Queue broker cluster.

When a JMS connection pool is created, there is one `ManagedConnectionFactory` instance associated with it. If you configure the `AddressList` as a `ManagedConnectionFactory` property, the `AddressList` configuration in the `ManagedConnectionFactory` takes precedence over the one defined in the Enterprise Server.

## Connection Failover

To specify whether the Enterprise Server tries to reconnect to the primary broker if the connection is lost, set the `reconnect-enabled` attribute in the JMS service. To specify the number of retries and the time between retries, set the `reconnect-attempts` and `reconnect-interval-in-seconds` attributes, respectively.

If reconnection is enabled and the primary broker goes down, the Enterprise Server tries to reconnect to another broker in the `AddressList`. The `AddressList` is updated whenever a JMS host configuration changes. The logic for scanning is decided by two JMS service attributes, `addresslist-behavior` and `addresslist-iterations`.

You can override these settings using JMS connection factory settings. For details, see [Chapter 4, “Configuring Java Message Service Resources,” in \*Sun GlassFish Enterprise Server 2.1 Administration Guide\*](#).

The Sun GlassFish Message Queue software transparently transfers the load to another broker when the failover occurs. JMS semantics are maintained during failover.

## Load-Balanced Message Inflow

You can configure `ActivationSpec` properties of the `jmsra` resource adapter in the `sun-ejb-jar.xml` file for a message-driven bean using `activation-config-property` elements. Whenever a message-driven bean (`EndPointFactory`) is deployed, the connector runtime engine finds these properties and configures them accordingly in the resource adapter. See “`activation-config-property`” in [Sun GlassFish Enterprise Server 2.1 Application Deployment Guide](#).

The Enterprise Server transparently enables messages to be delivered in random fashion to message-driven beans having same `ClientID`. The `ClientID` is required for durable subscribers.

For nondurable subscribers in which the `ClientID` is not configured, all instances of a specific message-driven bean that subscribe to same topic are considered equal. When a message-driven bean is deployed to multiple instances of the Enterprise Server, only one of the message-driven beans receives the message. If multiple distinct message-driven beans subscribe to same topic, one instance of each message-driven bean receives a copy of the message.

To support multiple consumers using the same queue, set the `maxNumActiveConsumers` property of the physical destination to a large value. If this property is set, the Sun GlassFish Message Queue software allows multiple message-driven beans to consume messages from same queue. The message is delivered randomly to the message-driven beans. If `maxNumActiveConsumers` is set to `-1`, there is no limit to the number of consumers.

To ensure that local delivery is preferred, set `addresslist-behavior` to `priority`. This setting specifies that the first broker in the `AddressList` is selected first. This first broker is the local colocated Message Queue instance. If this broker is unavailable, connection attempts are made to brokers in the order in which they are listed in the `AddressList`. This setting is the default for Enterprise Server instances that belong to a cluster.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see [“Usage Profiles” in Sun GlassFish Enterprise Server 2.1 Administration Guide](#).

---

## JMS Service High Availability

There are two levels of availability for JMS components:

- *Service availability* – At this level, availability of the JMS service matters, but it's not important if messages are not available for a while. As long as a connection gets failed over to a new available instance providing the service, the JMS component considers that the service is available and functions normally. This level of availability is described in [“Connection Failover” on page 282](#).
- *Data availability* – At this level, both availability of the service and persistent messages are required. JMS semantics of once and only once delivery and message ordering are also handled at this level.

You can enable data availability in the Sun GlassFish Message Queue cluster that comprises the Java Message Service (JMS). Messages are saved to the high-availability database (HADB) if HADB is installed, data availability is enabled, and the enterprise profile is selected. For information about profiles, see [“Usage Profiles” in Sun GlassFish Enterprise Server 2.1 Administration Guide](#). You must enable availability for the Enterprise Server instances before you can enable data availability for the corresponding brokers.

---

**Note** – Individual applications and modules cannot control or override JMS availability.

---

To enable data availability, select the Availability Service component under the relevant configuration in the Admin Console. Check the Availability Service box. To enable availability for the JMS service, select the JMS Availability tab, then check the Availability Service box. All

instances in an Enterprise Server cluster should have the same instance availability and JMS availability settings to ensure consistent behavior. For details, see the *Sun GlassFish Enterprise Server 2.1 High Availability Administration Guide*.

---

**Note** – Some topics in the documentation pertain to features that are available only in domains that are configured to support clusters. Examples of domains that support clusters are domains that are created with the cluster profile or the enterprise profile. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

## Transactions and Non-Persistent Messages

During transaction recovery, non-persistent messages might be lost. If the broker fails between the transaction manager’s prepare and commit operations, any non-persistent message in the transaction is lost and cannot be delivered. A message that is not saved to a persistent store is not available for transaction recovery.

## Authentication With ConnectionFactory

If your web, EJB, or client module has `res-auth` set to `Container`, but you use the `ConnectionFactory.createConnection("user", "password")` method to get a connection, the Enterprise Server searches the container for authentication information before using the supplied user and password. Version 7 of the Enterprise Server threw an exception in this situation.

## Message Queue varhome Directory

The Sun GlassFish Message Queue software uses a default directory for storing data such as persistent messages and its log file. This directory is called `varhome`. The Enterprise Server uses `domain-dir/imq` as the `varhome` directory if the type of relationship between the Enterprise Server and the Message Queue software is `LOCAL` or `EMBEDDED`. If the relationship type is `REMOTE`, the Message Queue software determines the `varhome` location. For more information about the types of relationships between the Enterprise Server and Message Queue, see “The JMS Provider” on page 276.

When executing Message Queue scripts such as `as-install/imq/bin/imqusermgr`, use the `-varhome` option to point the scripts to the Message Queue data if the relationship type is `LOCAL` or `EMBEDDED`. For example:

```
imqusermgr -varhome $AS_INSTALL/domains/domain1/imq add -u testuser
```

For more information about the Message Queue software, refer to the documentation at <http://docs.sun.com/coll/1343.8>.

# Delivering SOAP Messages Using the JMS API

Web service clients use the Simple Object Access Protocol (SOAP) to communicate with web services. SOAP uses a combination of XML-based data structuring and Hyper Text Transfer Protocol (HTTP) to define a standardized way of invoking methods in objects distributed in diverse operating environments across the Internet.

For more information about SOAP, see the Apache SOAP web site at <http://xml.apache.org/soap/index.html>.

You can take advantage of the JMS provider's reliable messaging when delivering SOAP messages. You can convert a SOAP message into a JMS message, send the JMS message, then convert the JMS message back into a SOAP message. The following sections explain how to do these conversions:

- “To Send SOAP Messages Using the JMS API” on page 285
- “To Receive SOAP Messages Using the JMS API” on page 286

## ▼ To Send SOAP Messages Using the JMS API

### 1 Import the `MessageTransformer` library.

```
import com.sun.messaging.xml.MessageTransformer;
```

This is the utility whose methods you use to convert SOAP messages to JMS messages and the reverse. You can then send a JMS message containing a SOAP payload as if it were a normal JMS message.

### 2 Initialize the `TopicConnectionFactory`, `TopicConnection`, `TopicSession`, and `publisher`.

```
tcf = new TopicConnectionFactory();
tc = tcf.createTopicConnection();
session = tc.createTopicSession(false, Session.AUTO_ACKNOWLEDGE);
topic = session.createTopic(topicName);
publisher = session.createPublisher(topic);
```

### 3 Construct a SOAP message using the SOAP with Attachments API for Java (SAAJ).

```
/*construct a default soap MessageFactory */
MessageFactory mf = MessageFactory.newInstance();
* Create a SOAP message object.*/
SOAPMessage soapMessage = mf.createMessage();
/** Get SOAP part.*/
SOAPPart soapPart = soapMessage.getSOAPPart();
/* Get SOAP envelope. */
SOAPEnvelope soapEnvelope = soapPart.getEnvelope();
/* Get SOAP body.*/
SOAPBody soapBody = soapEnvelope.getBody();
```

```
/* Create a name object. with name space */
/* http://www.sun.com/imq. */
Name name = soapEnvelope.createName("HelloWorld", "hw",
    "http://www.sun.com/imq");
* Add child element with the above name. */
SOAPElement element = soapBody.addChildElement(name)
/* Add another child element.*/
element.addTextNode( "Welcome to Sun Java System Web Services." );
/* Create an attachment with activation API.*/
URL url = new URL ("http://java.sun.com/webservices/");
DataHandler dh = new DataHandler (url);
AttachmentPart ap = soapMessage.createAttachmentPart(dh);
/*set content type/ID. */
ap.setContentType("text/html");
ap.setContentId("cid-001");
/** add the attachment to the SOAP message.*/
soapMessage.addAttachmentPart(ap);
soapMessage.saveChanges();
```

**4 Convert the SOAP message to a JMS message by calling the `MessageTransformer.SOAPMessageIntoJMSMessage()` method.**

```
Message m = MessageTransformer.SOAPMessageIntoJMSMessage (soapMessage,
    session );
```

**5 Publish the JMS message.**

```
publisher.publish(m);
```

**6 Close the JMS connection.**

```
tc.close();
```

## ▼ To Receive SOAP Messages Using the JMS API

**1 Import the `MessageTransformer` library.**

```
import com.sun.messaging.xml.MessageTransformer;
```

This is the utility whose methods you use to convert SOAP messages to JMS messages and the reverse. The JMS message containing the SOAP payload is received as if it were a normal JMS message.

**2 Initialize the `TopicConnectionFactory`, `TopicConnection`, `TopicSession`, `TopicSubscriber`, and `Topic`.**

```
messageFactory = MessageFactory.newInstance();
tcf = new com.sun.messaging.TopicConnectionFactory();
tc = tcf.createTopicConnection();
```

```
session = tc.createTopicSession(false, Session.AUTO_ACKNOWLEDGE);
topic = session.createTopic(topicName);
subscriber = session.createSubscriber(topic);
subscriber.setMessageListener(this);
tc.start();
```

- 3 Use the `onMessage` method to receive the message. Use the `SOAPMessageFromJMSMessage` method to convert the JMS message to a SOAP message.**

```
public void onMessage (Message message) {
    SOAPMessage soapMessage =
        MessageTransformer.SOAPMessageFromJMSMessage( message,
            messageFactory ); }
```

- 4 Retrieve the content of the SOAP message.**





# Using the JavaMail API

---

This chapter describes how to use the JavaMail™ API, which provides a set of abstract classes defining objects that comprise a mail system.

This chapter contains the following sections:

- “Introducing JavaMail” on page 289
- “Creating a JavaMail Session” on page 290
- “JavaMail Session Properties” on page 290
- “Looking Up a JavaMail Session” on page 290
- “Sending and Reading Messages Using JavaMail” on page 291

## Introducing JavaMail

The JavaMail API defines classes such as `Message`, `Store`, and `Transport`. The API can be extended and can be subclassed to provide new protocols and to add functionality when necessary. In addition, the API provides concrete subclasses of the abstract classes. These subclasses, including `MimeMessage` and `MimeBodyPart`, implement widely used Internet mail protocols and conform to the RFC822 and RFC2045 specifications. The JavaMail API includes support for the IMAP4, POP3, and SMTP protocols.

The JavaMail architectural components are as follows:

- The *abstract layer* declares classes, interfaces, and abstract methods intended to support mail handling functions that all mail systems support.
- The *internet implementation layer* implements part of the abstract layer using the RFC822 and MIME internet standards.
- JavaMail uses the *JavaBeans Activation Framework* (JAF) to encapsulate message data and to handle commands intended to interact with that data.

For more information, see [Chapter 5, “Configuring JavaMail Resources,”](#) in *Sun GlassFish Enterprise Server 2.1 Administration Guide* and the JavaMail specification at

<http://java.sun.com/products/javamail/>. A useful JavaMail tutorial is located at <http://java.sun.com/developer/onlineTraining/JavaMail/>.

## Creating a JavaMail Session

You can create a JavaMail session in the following ways:

- In the Admin Console, open the Resources component and select JavaMail Sessions. For details, click the Help button in the Admin Console.
- Use the `asadmin create-javamail-resource` command. For details, see the *Sun GlassFish Enterprise Server 2.1 Reference Manual*.

## JavaMail Session Properties

You can set properties for a JavaMail Session object. Every property name must start with a `mail-` prefix. The Enterprise Server changes the dash (-) character to a period (.) in the name of the property and saves the property to the `MailConfiguration` and `JavaMailSession` objects. If the name of the property doesn't start with `mail-`, the property is ignored.

For example, if you want to define the property `mail.from` in a JavaMail Session object, first define the property as follows:

- Name – `mail-from`
- Value – `john.doe@sun.com`

## Looking Up a JavaMail Session

The standard Java Naming and Directory Interface (JNDI) subcontext for JavaMail sessions is `java:comp/env/mail`.

Registering JavaMail sessions in the `mail` naming subcontext of a JNDI namespace, or in one of its child subcontexts, is standard. The JNDI namespace is hierarchical, like a file system's directory structure, so it is easy to find and nest references. A JavaMail session is bound to a logical JNDI name. The name identifies a subcontext, `mail`, of the root context, and a logical name. To change the JavaMail session, you can change its entry in the JNDI namespace without having to modify the application.

The resource lookup in the application code looks like this:

```
InitialContext ic = new InitialContext();
String snName = "java:comp/env/mail/MyMailSession";
Session session = (Session)ic.lookup(snName);
```

For more information about the JNDI API, see [Chapter 17, “Using the Java Naming and Directory Interface.”](#)

## Sending and Reading Messages Using JavaMail

The following sections describe how to send and read messages using the JavaMail API:

- “To Send a Message Using JavaMail” on page 291
- “To Read a Message Using JavaMail” on page 292

### ▼ To Send a Message Using JavaMail

#### 1 Import the packages that you need.

```
import java.util.*;
import javax.activation.*;
import javax.mail.*;
import javax.mail.internet.*;
import javax.naming.*;
```

#### 2 Look up the JavaMail session.

```
InitialContext ic = new InitialContext();
String snName = "java:comp/env/mail/MyMailSession";
Session session = (Session)ic.lookup(snName);
```

For more information, see [“Looking Up a JavaMail Session” on page 290.](#)

#### 3 Override the JavaMail session properties if necessary.

For example:

```
Properties props = session.getProperties();
props.put("mail.from", "user2@mailserver.com");
```

#### 4 Create a MimeMessage.

The `msgRecipient`, `msgSubject`, and `msgTxt` variables in the following example contain input from the user:

```
Message msg = new MimeMessage(session);
msg.setSubject(msgSubject);
msg.setSentDate(new Date());
msg.setFrom();
msg.setRecipients(Message.RecipientType.TO,
    InternetAddress.parse(msgRecipient, false));
msg.setText(msgTxt);
```

**5 Send the message.**

```
Transport.send(msg);
```

**▼ To Read a Message Using JavaMail****1 Import the packages that you need.**

```
import java.util.*;
import javax.activation.*;
import javax.mail.*;
import javax.mail.internet.*;
import javax.naming.*;
```

**2 Look up the JavaMail session.**

```
InitialContext ic = new InitialContext();
String snName = "java:comp/env/mail/MyMailSession";
Session session = (javax.mail.Session)ic.lookup(snName);
```

For more information, see [“Looking Up a JavaMail Session”](#) on page 290.

**3 Override the JavaMail session properties if necessary.**

For example:

```
Properties props = session.getProperties();
props.put("mail.from", "user2@mailserver.com");
```

**4 Get a Store object from the Session, then connect to the mail server using the Store object’s connect() method.**

You must supply a mail server name, a mail user name, and a password.

```
Store store = session.getStore();
store.connect("MailServer", "MailUser", "secret");
```

**5 Get the INBOX folder.**

```
Folder folder = store.getFolder("INBOX");
```

**6 It is efficient to read the Message objects (which represent messages on the server) into an array.**

```
Message[] messages = folder.getMessages();
```

# Using the Application Server Management Extensions

---

Sun GlassFish Enterprise Server uses [Enterprise Server Management eXtensions \(AMX\)](http://glassfish.dev.java.net/javaee5/amx/index.html) (<http://glassfish.dev.java.net/javaee5/amx/index.html>) for management and monitoring purposes. AMX technology exposes managed resources for remote management as the Java™ Management eXtensions (JMX™) API.

The Enterprise Server incorporates the [JMX 1.2 Reference Implementation](http://java.sun.com/products/JavaManagement/index.jsp) (<http://java.sun.com/products/JavaManagement/index.jsp>), which was developed by the Java Community Process as [Java Specification Request \(JSR\) 3](http://jcp.org/en/jsr/detail?id=3) (<http://jcp.org/en/jsr/detail?id=3>), and the [JMX Remote API 1.0 Reference Implementation](http://jcp.org/en/jsr/detail?id=160), which is [JSR 160](http://jcp.org/en/jsr/detail?id=160) (<http://jcp.org/en/jsr/detail?id=160>).

This chapter assumes some familiarity with the JMX technology, but the AMX interfaces can be used for the most part without understanding JMX. For more information about JMX, see the [JMX specifications and Reference Implementations](http://java.sun.com/products/JavaManagement/download.html) (<http://java.sun.com/products/JavaManagement/download.html>).

For information about creating custom MBeans, see [Chapter 14](#), “Developing Custom MBeans.”

This chapter contains the following topics:

- “About AMX” on page 294
- “AMX MBeans” on page 295
- “Dynamic Client Proxies” on page 298
- “Connecting to the Domain Administration Server” on page 298
- “Examining AMX Code Samples” on page 299
- “Running the AMX Samples” on page 302

## About AMX

AMX is an API that exposes all of the Enterprise Server configuration, monitoring and JSR 77 MBeans as easy-to-use client-side dynamic proxies implementing the AMX interfaces. To understand the design and implementation of the AMX API, you can get started with this [white paper](http://glassfish.dev.java.net/nonav/javaee5/amx/amx.html) (<http://glassfish.dev.java.net/nonav/javaee5/amx/amx.html>).

Complete API documentation for AMX is provided in the [Enterprise Server package](http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/index.html) (<http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/index.html>).

`com.sun.appserv.management`

The code samples in this section are taken from the package:

`com.sun.appserv.management.sample`

The Enterprise Server is based around the concept of *administration domains*. Each domain consists of one or more *managed resources*. A managed resource can be an Enterprise Server instance, a cluster of such instances, or a manageable entity within a server instance. A managed resource is of a particular type, and each resource type exposes a set of attributes and administrative operations that change the resource's state.

Managed resources are exposed as JMX *management beans*, or *MBeans*. While the MBeans can be accessed using standard JMX APIs (for example, `MBeanServerConnection`), most users find the use of the AMX client-side dynamic proxies much more convenient.

Virtually all components of the Enterprise Server are visible for monitoring and management through AMX. You can use third-party tools to perform all common administrative tasks programmatically, based on the JMX and JMX Remote API standards.

The AMX API consists of a set of interfaces. The interfaces are implemented by [client-side dynamic proxies](http://glassfish.dev.java.net/nonav/javaee5/amx/amx.html#AMXDynamicClientProxy) (<http://glassfish.dev.java.net/nonav/javaee5/amx/amx.html#AMXDynamicClientProxy>), each of which is associated with a server-side MBean in the Domain Administration Server (DAS). AMX provides routines to obtain proxies for MBeans, starting with the `DomainRoot` interface (see <http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/DomainRoot.html>).

---

**Note** – The term AMX interface in the context of this document should be understood as synonymous with a client-side dynamic proxy implementing that interface.

---

You can navigate generically through the MBean hierarchy using the `com.sun.appserv.management.base.Container` interface (see <http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/base/Container.html>).

When using AMX, the interfaces defined are implemented by client-side dynamic proxies, but they also implicitly define the `MBeanInfo` that is made available by the `MBean` or `MBeans` corresponding to it. Certain operations defined in the interface might have a different return type or a slightly different name when accessed through the `MBean` directly. This results from the fact that direct access to JMX requires the use of `ObjectName`, whereas the AMX interfaces use strongly typed proxies implementing the interface(s).

## AMX MBeans

All AMX MBeans are represented as interfaces in a subpackage of `com.sun.appserv.management` (see <http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/package-summary.html>) and are implemented by dynamic proxies on the client-side. Note that client-side means any client, wherever it resides. AMX may be used within the server itself such as in a custom `MBean`. While you can access AMX MBeans directly through standard JMX APIs, most users find the use of AMX interface (proxy) classes to be most convenient.

An AMX `MBean` belongs to an Enterprise Server domain. There is exactly one domain per DAS. Thus all MBeans accessible through the DAS belong to a single Enterprise Server administrative domain. All MBeans in an Enterprise Server administrative domain, and hence within the DAS, belong to the JMX domain `amx`. All AMX MBeans can be reached by navigating through the `DomainRoot`.

---

**Note** – Any MBeans that do not have the JMX domain `amx` are not part of AMX, and are neither documented nor supported for use by clients.

---

AMX defines different types of `MBean`, namely, *configuration* MBeans, *monitoring* MBeans, *utility* MBeans and *Java EE management JSR 77* (<http://jcp.org/en/jsr/detail?id=77>) MBeans. These MBeans are logically related in the following ways:

- They all implement the `com.sun.appserv.management.base.AMX` interface (see <http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/base/AMX.html>).
- They all have a `j2eeType` and `name` property within their `ObjectName`. See `com.sun.appserv.management.base.XTypes` (<http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/base/XTypes.html>) and `com.sun.appserv.management.j2ee.J2EETypes` (<http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/j2ee/J2EETypes.html>) for the available values of the `j2eeType` property.
- All MBeans that logically contain other MBeans implement the `com.sun.appserv.management.base.Container` interface.

- JSR 77 MBeans that have a corresponding configuration or monitoring peer expose it using `getConfigPeer()` or `getMonitoringPeer()`. However, there are many configuration and monitoring MBeans that do not correspond to JSR 77 MBeans.

## Configuration MBeans

Configuration information for a given Enterprise Server domain is stored in a central repository that is shared by all instances in that domain. The central repository can only be written to by the DAS. However, configuration information in the central repository is made available to administration clients through AMX MBeans.

The configuration MBeans are those that modify the underlying `domain.xml` or related files. Collectively, they form a model representing the configuration and deployment repository and the operations that can be performed on them.

The Group Attribute of configuration MBeans, obtained from `getGroup()`, has a value of `com.sun.appserv.management.base.AMX.GROUP_CONFIGURATION`.

## Monitoring MBeans

Monitoring MBeans provide transient monitoring information about all the vital components of the Enterprise Server.

The Group Attribute of monitoring MBeans, obtained from `getGroup()`, has a value of `com.sun.appserv.management.base.AMX.GROUP_MONITORING`.

## Utility MBeans

Utility MBeans provide commonly used services to the Enterprise Server.

The Group Attribute of utility MBeans, obtained from `getGroup()`, has a value of `com.sun.appserv.management.base.AMX.GROUP_UTILITY`.

## Java EE Management MBeans

The Java EE management MBeans implement, and in some cases extend, the management hierarchy as defined by JSR 77 (<http://jcp.org/en/jsr/detail?id=77>), which specifies the management model for the whole Java EE platform.

The AMX JSR 77 MBeans offer access to configuration and monitoring MBeans using the `getMonitoringPeer()` and `getConfigPeer()` methods.



The Group Attribute of Java EE management MBeans, obtained from `getGroup()`, has a value of `com.sun.appserv.management.base.AMX.GROUP_JSR77`.

## Other MBeans

MBeans that do not fit into one of the above four categories have the value `com.sun.appserv.management.base.AMX.GROUP_OTHER`. One such example is `com.sun.appserv.management.deploy.DeploymentMgr` (see <http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/deploy/DeploymentMgr.html>).

## MBean Notifications

All AMX MBeans that emit Notifications place a `java.util.Map` within the `UserData` field of a standard JMX Notification, which can be obtained using `Notification.getUserData()`. Within the map are one or more items, which vary according to the Notification type. Each Notification type, and the data available within the Notification, is defined in the Javadoc of the MBean (AMX interface) that emits it.

Note that certain standard Notifications, such as `javax.management.AttributeChangeNotification` (see <http://java.sun.com/javase/6/docs/api/javax/management/AttributeChangeNotification.html>) do not and cannot follow this behavior.

## Access to MBean Attributes

An AMX MBean Attribute is accessible in three ways:

- Dotted names using `MonitoringDottedNames` and `ConfigDottedNames`
- Attributes on MBeans using `getAttribute(s)` and `setAttributes(s)` (from the standard JMX API)
- Getters/setters within the MBean's interface class, for example, `getPort()`, `setPort()`, and so on

All dotted names that are accessible through the command line interface are available as Attributes within a single MBean. This includes properties, which are provided as Attributes beginning with the prefix `property.`, for example, `server.property.myproperty`.

---

**Note** – Certain attributes that ought to be of a specific type, such as `int`, are declared as `java.lang.String`. This is because the value of the attribute may be a template of a form such as `${HTTP_LISTENER_PORT}`.

---

## Dynamic Client Proxies

Dynamic Client Proxies are an important part of the AMX API, and enhance ease-of-use for the programmer.

JMX MBeans can be used directly by an `MBeanServerConnection` (see <http://java.sun.com/javase/6/docs/api/javax/management/MBeanServerConnection.html>) to the server. However, client proxies greatly simplify access to Attributes and operations on MBeans, offering `get/set` methods and type-safe invocation of operations. Compiling against the AMX interfaces means that compile-time checking is performed, as opposed to server-side runtime checking, when invoked generically through `MBeanServerConnection`.

See the API documentation for the `com.sun.appserv.management` package and its sub-packages for more information about using proxies. The API documentation explains the use of AMX with proxies. If you are using JMX directly (for example, by using `MBeanServerConnection`), the return type, argument types, and method names might vary as needed for the difference between a strongly-typed proxy interface and generic `MBeanServerConnection/ObjectName` interface.

## Connecting to the Domain Administration Server

As stated in “[Configuration MBeans](#)” on page 296, the AMX API allows client applications to connect to Enterprise Server instances using the DAS. All AMX connections are established to the DAS only: AMX does not support direct connections to individual server instances. This makes it simple to interact with all servers, clusters, and so on, with a single connection.

Sample code for connecting to the DAS is shown in “[Connecting to the DAS](#)” on page 299. The `com.sun.appserv.management.helper.Connect` class (see <http://glassfish.dev.java.net/nonav/javaee5/amx/javadoc/com/sun/appserv/management/helper/Connect.html>) is also available.

## Examining AMX Code Samples

An overview of the AMX API and code samples that demonstrate various uses of the AMX API can be found at

<http://glassfish.dev.java.net/nonav/javaee5/amx/samples/javadoc/index.html> and  
<http://glassfish.dev.java.net/nonav/javaee5/amx/samples/javadoc/amxsamples/Samples.html>.

The sample implementation is based around the `SampleMain` class. The principal uses of AMX demonstrated by `SampleMain` are the following:

- “Starting an Enterprise Server” on page 300
- “Deploying an Archive” on page 301
- “Displaying the AMX MBean Hierarchy” on page 301
- “Setting Monitoring States” on page 301
- “Accessing AMX MBeans” on page 301
- “Accessing and Displaying the Attributes of an AMX MBean” on page 301
- “Listing AMX MBean Properties” on page 301
- “Performing Queries” on page 301
- “Monitoring Attribute Changes” on page 302
- “Undeploying Modules” on page 302
- “Stopping an Enterprise Server” on page 302

All of these actions are performed by commands that you give to `SampleMain`. Although these commands are executed by `SampleMain`, they are defined as methods of the class `Samples`, which is also found in the `com.sun.appserv.management.sample` package.

## The SampleMain Class

The `SampleMain` class creates a connection to a DAS, and creates an interactive loop in which you can run the various commands defined in `Samples` that demonstrate different uses of AMX.

## Connecting to the DAS

The connection to the DAS is shown in the following code.

```
[...]  
public static AppserverConnectionSource  
    connect(  
        final String host,  
        final int port,  
        final String user,  
        final String password,  
        final TLSParams tlsParams )
```

```

throws IOException
{
    final String info = "host=" + host + ", port=" + port +
        ", user=" + user + ", password=" + password +
        ", tls=" + (tlsParams != null);

    SampleUtil.println( "Connecting..." + info );

    final AppserverConnectionSource conn =
        new AppserverConnectionSource(
            AppserverConnectionSource.PROTOCOL_RMI,
            host, port, user, password, tlsParams, null);

    conn.getJMXConnector( false );

    SampleUtil.println( "Connected: " + info );

    return( conn );
}
[...]
```

A connection to the DAS is obtained using an instance of the `com.sun.appserv.management.client.AppserverConnectionSource` class. For the connection to be established, you must know the name of the host and port number on which the DAS is running, and have the correct user name, password and TLS parameters.

After the connection to the DAS is established, `DomainRoot` is obtained as follows:

```
DomainRoot domainRoot = appserverConnectionSource.getDomainRoot();
```

This `DomainRoot` instance is a client-side dynamic proxy to the MBean `amx:j2eeType=X-DomainRoot,name=amx`.

See the API documentation for `com.sun.appserv.management.client.AppserverConnectionSource` for further details about connecting to the DAS using the `AppserverConnectionSource` class.

However, if you prefer to work with standard JMX, instead of getting `DomainRoot`, you can get the `MBeanServerConnection` or `JMXConnector`, as shown:

```
MBeanServerConnection conn =
appserverConnectionSource.getMBeanServerConnection( false );
JMXConnector jmxConn =
appserverConnectionSource.getJMXConnector( false );
```

## Starting an Enterprise Server

The `Samples.startServer` method demonstrates how to start an Enterprise Server.

In this sample AMX implementation, all the tasks are performed by the command `start-server` when you run `SampleMain`. See the `startServer` method to see how this command is implemented. Click the method name to see the source code.

## Deploying an Archive

The `Samples.uploadArchive()` and `deploy` methods demonstrate how to upload and deploy a Java EE archive file.

## Displaying the AMX MBean Hierarchy

The `Samples.displayHierarchy` method demonstrates how to display the AMX MBean hierarchy.

## Setting Monitoring States

The `Samples.setMonitoring` method demonstrates how to set monitoring states.

## Accessing AMX MBeans

The `Samples.handleList` method demonstrates how to access many (but not all) configuration elements.

## Accessing and Displaying the Attributes of an AMX MBean

The `Samples.displayAllAttributes` method demonstrates how to access and display the attributes of an AMX MBean.

## Listing AMX MBean Properties

The `Samples.displayAllProperties` method demonstrates how to list AMX MBean properties.

## Performing Queries

The `Samples.demoQuery` method demonstrates how to perform queries.

The `demoQuery()` method uses other methods that are defined by `Samples`, namely `displayWild()`, and `displayJ2EEType()`.

## Monitoring Attribute Changes

The `Samples.demoJMXMonitor` method demonstrates how to monitor attribute changes.

## Undeploying Modules

The `Samples.undeploy` method demonstrates how to undeploy a module.

## Stopping an Enterprise Server

The `Samples.stopServer` method demonstrates how to stop an Enterprise Server. The `stopServer` method simply calls the `Samples.getJ2EEServer` method on a given server instance, and then calls `J2EEServer.stop`.

# Running the AMX Samples

The following section lists the steps to run the AMX samples.

## ▼ To Run the AMX Sample

- 1 **Ensure that the JAR file `appserv-ext.jar` has been added to your classpath. Some examples also require that `j2ee.jar` be present.**
- 2 **Define a `SampleMain.properties` file, which provides the parameters required by `AppserverConnectionSource` to connect to the DAS.**

The file `SampleMain.properties` file should use the following format:

```
connect.host=localhost
connect.port=8686
connect.user=admin
connect.password=admin123
connect.truststore=sample-truststore
connect.truststorePassword=changeme
connect.useTLS=true
```

---

**Note** – If you are running the Enterprise Server enterprise profile, TLS must be used (provided it has not been turned off). The useTLS property must therefore be set to useTLS=true. For information about profiles, see “Usage Profiles” in *Sun GlassFish Enterprise Server 2.1 Administration Guide*.

---

**3 Scripts are provided in the `com.sun.appserv.management.sample` package to run the AMX samples.**

Start `SampleMain` by running the appropriate script for your platform:

- `run-samples.sh` on UNIX or Linux platforms
- `run-samples.bat` on Microsoft Windows platforms

**4 After `SampleMain` is running, you can interact with it by typing the commands examined above:**

- Enter Command> **start-server** *serverName*
- Enter Command> **list-attributes**

You see output like this:

```
--- Attributes for X-DomainRoot=amx ---
AttributeNames=[...]
BulkAccessObjectName=amx:j2eeType=X-BulkAccess,name=na
DomainConfigObjectName=amx:j2eeType=X-DomainConfig,name=na
MBeanInfoIsInvariant=true
J2EEDomainObjectName=amx:j2eeType=J2EEDomain,name=amx
AppserverDomainName=amx
ObjectName=amx:j2eeType=X-DomainRoot,name=amx
[...]
```

- Enter Command> **show-hierarchy**

You see output like this:

```
X-DomainRoot=amx
X-ConfigDottedNames
X-SystemInfo
X-QueryMgr
X-DeploymentMgr
X-UploadDownloadMgr
X-BulkAccess
X-MonitoringDottedNames
X-JMXMonitorMgr
X-Sample
X-DomainConfig
X-WebModuleConfig=adminui
X-WebModuleConfig=adminapp
X-WebModuleConfig=com_sun_web_ui
```

```
X-JDBCResourceConfig=jdbc/__default
X-JDBCResourceConfig=jdbc/__TimerPool
X-J2EEApplicationConfig=MEjbApp
[...]
```

- Enter Command> **list**

You see output like this:

```
--- Top-level ---
ConfigConfig: [server2-config, default-config, server-config,
server3-config]
ServerConfig: [server3, server, server2]
StandaloneServerConfig: [server3, server, server2]
ClusteredServerConfig: []
ClusterConfig: []
[...]
```

- Enter Command> **list-properties**

You see output like this:

```
Properties for:
amx:j2eeType=X-JDBCConnectionPoolConfig,name=DerbyPool
Password=pbPublic
DatabaseName=jdbc:derby://localhost:9092/sun-appserv-samples
User=pbPublic
[...]
```

- Enter Command> **query**

You see output like this:

```
--- Queried for j2eeType=X-*ResourceConfig ---
j2eeType=X-JDBCResourceConfig,name=jdbc/__default
j2eeType=X-JDBCResourceConfig,name=jdbc/__TimerPool
[...]
```

- And so on for the other commands:

Enter Command> **demo-jmx-monitor**

Enter Command> **set-monitoring** *monitoringLevel* (one of HIGH, LOW or OFF)

Enter Command> **stop-server** *serverName*

Enter Command> **quit**



# Index

---

## Numbers and Symbols

@OrderBy and session cache sharing, 127

## A

ACC, 213-214

ACC clients

  appclient script, 224

  failover, 216

  invoking a JMS resource, 217-218

  invoking an EJB component, 215-217

  Java Web Start, 218-223

  load balancing, 216

  making a remote call, 216

  package-appclient script, 224

  running, 218-223, 224

  SSL, 213-214, 224-226

ACC

  annotation, 214

  naming, 214

  security, 213-214, 224-226

action attribute, 51, 55

activation-config-property element, 282-283

ActivationSpec properties, 282-283

AddressList

  and connections, 281-282

  and default JMS host, 279

Admin Console, 30

  Admin Object Resources page, 229

  Admin Service page, 248

  App Client Modules page, 219

Admin Console (*Continued*)

  Audit Modules page, 87

  Classpath Prefix and Suffix fields, 35

  Classpath Prefix for jndi.properties, 273

  CMP resource configuration, 199

  Connector Connection Pools page, 229

  Connector Modules page, 229

  Connector Resources page, 229

  Connector Service page

    Shutdown Timeout field, 233

  connector thread pool assignment, 231

  Custom MBeans page, 245

  Debug Enabled field, 70

  Default Virtual Server field, 154

  Generate RMIS stubs field, 221

  HPROF configuration, 73

  JACC Providers page, 87

  JavaMail Sessions page, 290

  JDBC Connection Pools page, 254

    Allow Non Component Callers field, 259

    Non-Transactional Connections field, 257

    Ping button, 255

  JDBC Resources page, 255

  JMS Hosts page, 279

  JMS Resources page, 280

  JMS Service page, 278

  JNDI page

    Custom Resources page, 273

    External Resources page, 273

  JProbe configuration, 75

  Libraries field, 38

  Lifecycle Modules page, 241

- Admin Console (*Continued*)
  - Locale field, 152
  - Logging tab, 72, 156
  - Message Security page
    - creating providers, 94
    - enabling providers, 93
  - Monitor tab, 156
  - online help for, 30
  - Physical Destinations page, 280
  - Realms page, 84
  - role mapping configuration, 83
  - Security Manager Enabled field, 92
  - Security Maps tab, 232
  - System Classpath field, 35, 40
  - Thread Pools page, 231
  - Transaction Log Location field, 265
  - Transaction Service page, 264
  - Virtual Servers page, 153, 154
  - Web Services page
    - Publish tab, 109
    - Registry tab, 109
    - Test button, 110
  - Write to System Log field, 135
- administered objects, 280
  - and connectors, 229
- allow-concurrent-access element, 181
- AllowManagedFieldsInDefaultFetchGroup flag, 202
- AllowMediatedWriteInDefaultFetchGroup flag, 202
- alternate document roots, 158-159
- AMX
  - about, 294-295
  - MBeans, 295-298
  - proxies, 298
  - samples, 299-302
    - running, 302-304
- annotation
  - application clients, 214
  - EJB 3.0 specification, 165
  - JNDI names, 270
  - message layer, 92
  - schema generation, 120-121
  - security, 81
- Ant, 30, 43-68
- ANT\_HOME environment variable, 43
- Apache Ant, 30, 43-68
- appclient script, 224
- Application class loader, 36
- Application Client Container, *See* ACC
- Application Server Management eXtensions, *See* AMX
- applications
  - disabling, 54-56
  - examples, 32
- appserv-ext.jar file, 239
- appserv-jwsacc.jar file, 221
- appserv-tags.jar file, 140
- appserv-tags.tld file, 140-141
- AppservPasswordLoginModule class, 86
- AppservRealm class, 86
- asadmin command, 30
  - create-admin-object, 229
  - create-audit-module, 87
  - create-auth-realm, 84
  - create-connector-connection-pool, 229, 281
  - create-connector-resource, 229
  - create-connector-security-map, 232
  - create-custom-resource, 273
  - create-domain, 222
  - create-javamail-resource, 290
  - create-jdbc-connection-pool, 254
    - allownoncomponentcallers option, 259
    - nontransactionalconnections option, 257
  - create-jdbc-resource, 255
  - create-jms-host, 279
  - create-jmsdest, 280
  - create-jndi-resource, 273
  - create-jvm-options, 180, 202
    - com.sun.appserv.transaction.nofdsync option, 265
    - java.security.debug option, 91
  - create-lifecycle-module, 241
  - create-mbean, 245-246
  - create-message-security-provider, 94
  - create-resource-adapter-config, 229, 231, 232
  - create-threadpool, 231
  - delete-jvm-options
    - java.security.manager option, 92
  - delete-mbean, 246
  - deploy-jbi-service-assembly, 111

- asadmin command (*Continued*)
    - deploy
      - and connectors, 229
      - availabilityenabled option, 175
    - deploydir
      - and connectors, 229
      - availabilityenabled option, 175
      - schema generation, 124, 196
    - deploy
      - libraries option, 38
      - precompilejsp option, 144
      - retrieve option, 215, 221
      - schema generation, 124, 196
    - flush-jmsdest, 280
    - generate-jvm-report, 71
    - get, 264, 278
    - get-client-stubs, 216, 221
    - jms-ping, 279
    - list-mbeans, 246-247
    - list-timers, 170
    - migrate-timers, 170
    - ping-connection-pool, 232, 255
    - publish-to-registry, 109
    - set
      - custom MBean attributes, 248
      - custom MBean disabling, 248
      - default message security provider, 93
      - default principal settings, 83
      - java-web-start-enabled attribute, 219
      - jbi-enabled property, 112
      - JMS service settings, 278
      - JMX connector port, 248
      - transaction service settings, 264
    - undeploy
      - schema generation, 125, 197
  - asant script, 30, 43-68
    - Application Server specific tasks, 44-62
    - disabling deployed applications and modules, 54-56
    - updating deployed applications and modules, 59-60
    - using for deployment, 44-48
    - using for JSP precompilation, 58-59
    - using for server administration, 51-54, 57-58
    - using for undeployment, 48-51
  - asinstalldir attribute
    - sun-appserv-admin task, 57
    - sun-appserv-component task, 55
    - sun-appserv-deploy task, 47
    - sun-appserv-instance task, 52
    - sun-appserv-jspc task, 59
    - sun-appserv-undeploy task, 50
  - audit modules, 87-89
  - AuditModule class, 88-89
  - authentication
    - application clients, 213-214
    - audit modules, 88
    - JAAS, 85-87
    - JMS, 284
    - message-level, 99
    - P-asserted identity, 80
    - programmatic login, 102
    - realms, 84
    - single sign-on, 104-105
  - authorization
    - audit modules, 88
    - JAAS, 85-87
    - JACC, 87
    - roles, 82-84
  - automatic schema generation
    - for CMP, 192-198
    - Java Persistence options, 122-125
  - availability
    - configuring HTTP session persistence, 149-150, 150-151
  - availabilityenabled attribute, 46
  - availability
    - feature summary, 29
    - for ACC clients, 216
    - for stateful session beans, 172-177
    - for web modules, 145-147
    - of message-driven beans, 282-283
- B**
- bin directory, 43
  - binding attribute, 62
  - BLOB support, 191
  - Bootstrap class loader, 35

build.xml file, 30, 32

## C

cache for servlets

- default configuration, 136
- example configuration, 137
- helper class, 136, 138

cache sharing and @OrderBy, 127

cache tag, 142-143

CacheHelper interface, 138

cacheKeyGeneratorAttrName property, 138

caching

- a bean's state using version consistency, 200
- data using a non-transactional connection, 258
- EJB components, 167
- entities, 185
- JSP files, 140-143
- read-only beans, 179
- servlet results, 135-138
- stateful session beans, 171
- using a read-only bean for, 166, 180, 201

capture-schema command, 198-199

cascade attribute, 49

Catalina listeners, defining custom, 157-158

catalog attribute, 62

certificate realm, 84

checkpoint-at-end-of-method element, 176

checkpointing, 172

- selecting methods for, 176

class-loader element, 37, 155

class loaders, 33-42

- application-specific, 38-39
- circumventing isolation, 39-42
- delegation hierarchy, 33-36
- isolation, 38

classpath attribute, 58, 60

classpath-prefix attribute, 35

classpath-suffix attribute, 35

classpath, changing, 35

classpathref attribute, 58

client JAR file, 41

client.policy file, 224

CLOB support, 192

cluster attribute, 52

CMP, *See* container-managed persistence

cmp-resource element, 199

cmt-max-runtime-exceptions property, 184

Comet support, 160

command attribute, 57

command-line server configuration, *See* asadmin  
command

commit options, 185

Common class loader, 35

- using to circumvent isolation, 40

compiling JSP files, 144

component subelement, 65-68

config attribute, 52

connection factory, 182

ConnectionFactory interface, 280

Connector class loader, 36, 242

connectors, 227-237

- administered objects, 229
- and JDBC, 228
- and JMS, 228
- and message-driven beans, 235-237
- and transactions, 262
- configuration options, 231-234
- configuring, 228
- connection pools, 229
- deployment, 229
- embedded, 230
- generic JMS, 277
- inbound connectivity, 234-235
- invalid connections, 233
- last agent optimization, 234
- redeployment, 230
- resources, 229
- shutdown timeout, 233
- Sun Java System Application Server support, 228
- testing connection pools, 232
- thread pools, 231

container-managed persistence

- configuring 1.1 finders, 202-203
- data types for mapping, 193-195
- deployment descriptor, 188-189
- mapping, 188
- performance features, 199-201

- container-managed persistence (*Continued*)
    - prefetching, 200-201
    - resource manager, 199
    - restrictions, 206-211
    - support, 187-188
    - version consistency, 200
  - context root, 134
  - context.xml file, 160
  - context, for JNDI naming, 269-272
  - contextroot attribute, 45, 66
  - CosNaming naming service, 271
  - cp attribute, 60
  - create-admin-object command, 229
  - create-audit-module command, 87
  - create-auth-realm command, 84
  - create-connector-connection-pool command, 229, 281
  - create-connector-resource command, 229
  - create-connector-security-map command, 232
  - create-custom-resource command, 273
  - create-domain command, 222
  - create-javamail-resource command, 290
  - create-jdbc-connection-pool command, 254
    - allownoncomponentcallers option, 259
    - nontransactionalconnections option, 257
  - create-jdbc-resource command, 255
  - create-jms-host command, 279
  - create-jmsdest command, 280
  - create-jndi-resource command, 273
  - create-jvm-options command, 180, 202
    - com.sun.appserv.transaction.nofdsync option, 265
    - java.security.debug option, 91
  - create-lifecycle-module command, 241
  - create-mbean command, 245-246
  - create-message-security-provider command, 94
  - create-resource-adapter-config command, 229, 231, 232
  - create-threadpool command, 231
  - createtables attribute, 46
  - custom MBeans
    - deployment or registration, 245-246
    - enabling and disabling, 248
    - handling attributes of, 248-249
    - life cycle, 244-245
    - listing information about, 246-247
  - custom MBeans (*Continued*)
    - location and classloading, 245
    - redeployment, 246
    - the MBeanServer, 247-248
    - undeployment, 246
  - custom resource, 273
- ## D
- DAS, connecting to, 298
  - data types
    - for CMP mapping, 193-195
    - for schema generation, 121-122
  - database properties, 118
  - databases
    - as transaction resource managers, 261
    - CMP resource manager, 199
    - properties, 118
    - schema capture, 198
    - specifying for Java Persistence, 116-117
    - supported, 254
  - dbvendorname attribute, 46
  - debug attribute, 52, 64
  - debugging, 69-76
    - enabling, 69-70
    - generating a stack trace, 71
    - JPDA options, 70-71
  - DeclareRoles annotation, 82-84
  - default virtual server, 154
  - default web module, 134, 154
  - default-web.xml file, 155-156
  - delegation, class loader, 37
  - delete-jvm-options command, java.security.manager option, 92
  - delete-mbean command, 246
  - demoQuery method, 301-302
  - deploy command
    - and connectors, 229
    - availabilityenabled option, 175
    - libraries option, 38
    - precompilejsp option, 144
    - retrieve option, 215, 221
    - schema generation, 124, 196
  - deploy-jbi-service-assembly command, 111

- deploydir command
    - and connectors, 229
    - availabilityenabled option, 175
    - schema generation, 124, 196
  - deployment descriptor files, 274
  - deployment
    - disabling deployed applications and modules, 54-56
  - deploymentplan attribute, 46
  - deployment
    - read-only beans, 181
    - undeploying an application or module, 48
    - using asant script, 44-48
  - destdir attribute, 58, 60, 62
  - destinations
    - destination resources, 280
    - physical, 279-280
  - destroy method, 138
  - development environment
    - creating, 27-32
    - tools for developers, 29-32
  - digest authentication, 84
  - directory listings, disabling, 155
  - displayHierarchy method, 301
  - distributable web application, 145
  - distributed HTTP sessions, 145-147
  - document root, 153, 154
  - document roots, alternate, 158-159
  - doGet method, 139
  - Domain Administration Server, *See* DAS
  - domain attribute, 60
  - domain.xml file
    - configuring single sign-on, 105
    - Shared Chain class loader, 273
    - System class loader, 35, 40
  - doPost method, 139
  - dropandcreatetables attribute, 46
  - droptables attribute, 49
- E**
- EJB 3.0
    - Java Persistence, 115-131
    - summary of changes, 165
  - EJB components
    - caching, 167-168
    - calling from a different application, 41
    - flushing, 169
    - pooling, 167-168, 171
    - remote bean invocations, 169
    - security, 82
    - thread pools, 169
  - EJB QL queries, 202-203
  - ejb-ref element, 274
  - ejb-ref mapping, using JNDI name instead, 42
  - EJB reference failover, 216
  - EJB Timer Service, 169
  - ejbPassivate, 179
  - enabled attribute, 46
  - encoding, of servlets, 152-153
  - endorsed standards override mechanism, 37-38
  - Enterprise Service Bus (ESB), 111-113
  - env-classpath-ignored attribute, 35
  - events, server life cycle, 239
  - example applications, 32
  - explicitcommand attribute, 57
  - extension attribute, 61, 62
  - external JNDI resource, 273
- F**
- fail-all-connections property, 233
  - failover
    - for ACC clients, 216
    - JMS connection, 282
    - object types supported for, 146-147, 173
    - of stateful session bean state, 172-177
    - of web module sessions, 145-147
  - fetch group, options for, 202
  - file attribute
    - component element, 66
    - sun-appserv-component task, 55
    - sun-appserv-deploy task, 45
    - sun-appserv-undeploy task, 49
    - sun-appserv-update task, 60
  - file realm, 84
  - fileset subelement, 68
  - finder limitation for Sybase, 129, 208

finder methods, 202-203  
 flat transactions, 185  
 flush-jmsdest command, 280  
 flush tag, 143  
 flushing of EJB components, 169  
 force attribute, 45, 66

## G

generate-jvm-report command, 71  
 generatermistubs attribute, 46  
 generic JMS resource adapter, 277  
 genwsdl attribute, 61  
 get-client-stubs command, 216, 221  
 get command, 264, 278  
 getCharacterEncoding method, 153  
 getCmdLineArgs method, 241  
 getConnection method, 256  
 getData method, 240  
 getEventType method, 240  
 getHeaders method, 157  
 getInitialContext method, 241, 272  
 getInstallRoot method, 241  
 getInstanceName method, 241  
 getLifecycleEventContext method, 240  
 GlassFish project, 28

## H

HADB  
   and HTTP session persistence, 150-151  
   and stateful session bean failover, 174  
   using for JMS data, 283-284  
 handling requests, 139  
 header management, 157  
 help for Admin Console tasks, 30  
 high availability, *See* availability  
 high-availability database, *See* HADB  
 host attribute  
   server element, 63  
   sun-appserv-component task, 55  
   sun-appserv-deploy task, 47  
   sun-appserv-instance task, 52

host attribute (*Continued*)  
   sun-appserv-undeploy task, 50  
 HPROF profiler, 73-74  
 HTTP sessions, 144-151  
   cookies, 144  
   distributed, 145-147  
   object types supported for failover, 146-147  
   session managers, 147-151  
   URL rewriting, 144  
 HttpServletRequest, 136

## I

idempotent requests, 156  
 IMAP4 protocol, 289-290  
 inbound connectivity, 234-235  
 Inet Oracle JDBC driver, 128, 191, 192  
 INIT\_EVENT, 239  
 init method, 138  
 InitialContext naming service handle, 269-272  
 installation, 27-28  
 instance attribute, 52, 63  
 instanceport attribute, 63  
 instantiating servlets, 138  
 internationalization, 152  
 Interoperable Naming Service, 271-272  
 is-connection-validation-required property, 233  
 is-failure-fatal attribute, 242  
 is-read-only-bean element, 181  
 isolation of class loaders, 38, 39-42

## J

J2EE Connector architecture, 227-237  
 J2SE policy file, 224  
 JACC, 87  
 JAR file, client for a deployed application, 41  
 Java Authentication and Authorization Service (JAAS), 85-87  
 Java Authorization Contract for Containers, *See* JACC  
 Java Business Integration (JBI), 111-113  
   Ant tasks for, 68  
 java-config element, 35

- Java Database Connectivity, *See* JDBC
- Java DB database, 116-117
- Java Debugger (jdb), 69
- Java EE Service Engine, 111-113
- Java EE tutorial, 133
- Java EE, security model, 80
- Java Management Extensions
  - See* JMX
- Java Message Service
  - See* JMS
- Java Naming and Directory Interface, *See* JNDI
- Java optional package mechanism, 37
- Java Persistence, 115-131
  - annotation for schema generation, 120-121
  - changing the provider, 126
  - data types for schema generation, 121-122
  - database for, 116-117
  - deployment options for schema
    - generation, 122-125
  - restrictions, 127-131
- Java Platform Debugger Architecture, *See* JPDA
- Java Servlet API, 133
- Java Transaction API (JTA), 261-267
- Java Transaction Service (JTS), 261-267
- Java Web Start, 218-223
  - signing client JAR files, 221-223
- JavaBeans, 139
- JavaMail
  - and JNDI lookups, 290-291
  - architecture, 289
  - creating sessions, 290
  - defined, 289-292
  - messages
    - reading, 292
    - sending, 291-292
  - session properties, 290
  - specification, 289
- JConsole, 247
- jdbc realm, 84
- JDBC
  - connection pool creation, 254
  - Connection wrapper, 256
  - creating resources, 255
  - integrating driver JAR files, 40, 254
- JDBC (*Continued*)
  - non-component callers, 259
  - non-transactional connections, 257-258
  - restrictions, 259
  - sharing connections, 255-256
  - specification, 253
  - supported drivers, 254
  - transaction isolation levels, 258
  - tutorial, 253
- JDOQL, 202-203
- JMS, 182, 275-287
  - jms-ping command, 279
- JMS
  - and transactions, 262
  - authentication, 284
  - checking if provider is running, 279
  - configuring, 278
  - connection failover, 282
  - connection pooling, 281-282
  - creating hosts, 279
  - creating resources, 280
  - debugging, 71
  - default host, 278-279
  - generic resource adapter, 277
  - high availability, 283-284
  - JMS Service administration, 277-280
  - load balancing, 282-283
  - provider, 276-277
- jmsra system JMS connector, 277
- JMS
  - restarting the client, 281
  - SOAP messages, 285-287
  - system connector for, 277
  - transactions and non-persistent messages, 284
- JMX, 243-249, 293-304
- JNDI
  - and EJB components, 274
  - and JavaMail, 290-291
  - and lifecycle modules, 241, 242, 272
  - custom resource, 273
  - defined, 269-274
  - external JNDI resources, 273
  - for message-driven beans, 182
  - global names, 270



**JNDI** (*Continued*)

- mapping references, 274
- name for container-managed persistence, 199
- tutorial, 269
- using instead of ejb-ref mapping, 42
- join tables, 190
- JPDA debugging options, 70-71
- JProbe profiler, 74-76
- JSP Engine class loader, 36
- JSP files
  - caching, 140-143
  - command-line compiler, 144
  - precompiling, 45, 58-59, 144
  - specification, 139
  - tag libraries, 140
- jspc command, 144
- JSR 109, 107
- JSR 115, 80, 87, 88
- JSR 12, 203
- JSR 160, 293
- JSR 181, 108
- JSR 196, 80, 93
- JSR 220, 115, 165
- JSR 224, 107
- JSR 3, 293
- JSR 77, 295-298
- JSR 907, 264-265

**K**

- keep attribute, 61, 62
- key attribute
  - of cache tag, 142
  - of flush tag, 143

**L**

- last agent optimization, 234, 262
- ldap realm, 84
- lib directory
  - and the Common class loader, 35
  - for a web application, 41
- libraries, 38-39, 39

- lifecycle modules, 239
  - allocating and freeing resources, 242
  - and class loaders, 242
  - and the server.policy file, 242
  - deployment, 241-242
  - naming environment, 272
- LifecycleEvent class, 240
- LifecycleEventContext interface, 241
- LifecycleListener interface, 240
- LifecycleListenerImpl.java file, 240
- LifeCycleModule class loader, 36, 242
- list-mbeans command, 246-247
- list-timers command, 170
- listeners, Catalina, defining custom, 157-158
- load balancing
  - and idempotent requests, 156
  - of ACC clients, 216
  - of message-driven beans, 282-283
- load-on-startup element in web.xml, 160
- locale, setting default, 152
- lock-when-loaded consistency level, 208
- logging, 72
  - in the web container, 156
- login method, 103
- LoginModule, 86
- login, programmatic, 102

**M**

- main.xml file, 32
- managed fields, 190-191
- mapping for container-managed persistence
  - considerations, 189-192
  - data types, 193-195
  - features, 188
- mapping resource references, 274
- markConnectionAsBad method, 256-257
- MBean class loader, 35
- MBeans
  - accessing, 301
  - AMX, 294-295, 295-298
  - attributes, 297-298
  - configuration, 296

MBeans (*Continued*)

- custom
  - See* custom MBeans
- definition, 243-249
- displaying attributes, 301
- Java EE management, 296-297
- listing properties, 301
- monitoring, 296
- notifications, 297
- other types, 297
- proxies, 298
- querying, 301-302
- undeploying, 302
- using to stop a server instance, 302
- utility, 296

mdb-connection-factory element, 182, 183

message-driven beans, 71, 181

- administering, 182
- connection factory, 182
- load balancing, 282-283
- monitoring, 183
- onMessage runtime exception, 183-184
- pool monitoring, 183
- pooling, 182
- restrictions, 183-184
- using with connectors, 235-237

message security, 92-101

- application-specific, 96-99
- responsibilities, 95
- sample application, 99-101

migrate-timers command, 170

Migration Tool, 31

mime-mapping element, 155

modules

- disabling, 54-56
- lifecycle, 239

monitoring in the web container, 156

MSSQL version consistency triggers, 209

MySQL database restrictions, 129-131, 209-211

**N**

naming service, 269-274

- native library path
  - configuring for hprof, 73
  - configuring for JProbe, 75
- nested transactions, 185
- NetBeans
  - about, 31
  - profiler, 73
- nocache attribute, of cache tag, 142
- nodeagent attribute, 52, 63

**O**

Oasis Web Services Security, *See* message security

object references supported for failover, 146-147, 173

online help, 30

onMessage method, 183, 287

Open ESB Starter Kit, 111-113

Oracle automatic mapping of date and time fields, 208

Oracle Inet JDBC driver, 128, 191, 192

Oracle Thin Type 4 Driver, workaround for, 266

Oracle TopLink Essentials, 115

oracle-xa-recovery-workaround property, 267

ORDER BY validation, disabling, 206

output from servlets, 134-135

**P**

P-asserted identity authentication, 80

package-applient script, 224

package attribute, 59, 62

pass-by-reference element, 167

permissions

- changing in server.policy, 89-91
- default in server.policy, 89

persistence store

- for HTTP sessions, 145-147, 149-150, 150-151
- for stateful session bean state, 172-177

persistence.xml file, 116-117, 122

physical destinations, 279-280

ping-connection-pool command, 232, 255

pool monitoring for MDBs, 183

pooling, 179

POP3 protocol, 289-290

port attribute  
 server element, 63  
 sun-appserv-component task, 55  
 sun-appserv-deploy task, 47  
 sun-appserv-instance task, 52  
 sun-appserv-undeploy task, 50  
 portname attribute, 61  
 precompilejsp attribute, 45, 66  
 precompiling JSP files, 144  
 prefetching, 200-201  
 primary key, 187, 190  
 profilers, 72-76  
 programmatic login, 102  
 ProgrammaticLogin class, 103  
 ProgrammaticLoginPermission permission, 103  
 property attribute, 52  
 protocol attribute, 61  
 proxies, AMX, 298  
 publish-to-registry command, 109

## Q

query hints, 125-126  
 Queue interface, 280  
 QueueConnectionFactory interface, 280

## R

read-only beans, 166-167, 178-181, 201  
 deploying, 181  
 refreshing, 179-181  
 readonly.relative.refresh.mode flag, 180  
 ReadOnlyBeanNotifier, 180-181  
 READY\_EVENT, 239  
 realms  
 application-specific, 85  
 configuring, 84  
 custom, 85-87  
 supported, 84  
 redirecting a URL, 159-160  
 references supported for failover, 146-147, 173  
 refresh attribute, of cache tag, 142  
 refresh-period-in-seconds element, 179

removing servlets, 138  
 request object, 139  
 res-sharing-scope deployment descriptor  
 setting, 255-256  
 resource-adapter-mid element, 236  
 resource adapters, *See* connectors  
 resource-env-ref element, 274  
 resource managers, 261-262  
 resource-ref element, 274  
 resource references, mapping, 274  
 resourcedestdir attribute, 60  
 retrievestubs attribute, 45, 66  
 RFC 3325, 80  
 RMI/IIOP over SSL, 224-226  
 roles, 82-84

## S

sample applications, 32  
 schema capture, 198  
 schema generation  
 automatic for CMP, 192-198  
 Java Persistence options for automatic, 122-125  
 scope attribute  
 of cache tag, 142  
 of flush tag, 143  
 secondary table, 189  
 security, 79-105  
 security manager, enabling and disabling, 92  
 security map, 231-232  
 security  
 ACC, 213-214, 224-226  
 annotations, 81  
 application level, 81  
 audit modules, 87-89  
 declarative, 81  
 disabling directory listings, 155  
 EJB components, 82  
 goals, 80  
 JACC, 87  
 Java EE model, 80  
 JMS, 284  
 message security, 92-101  
 of containers, 81-82

security (*Continued*)

- programmatically, 82
- programmatically login, 102
- roles, 82-84
- server.policy file, 89-92
- Sun Java System Application Server features, 80
- web applications, 82

## sei attribute, 60

## server.policy file, 89-92

- and lifecycle modules, 242
- changing permissions, 89-91
- default permissions, 89
- ProgrammaticLoginPermission, 103

## server subelement, 63-65

## server

- administering instances using asant, 51-54
- changing the classpath of, 35
- installation, 27-28
- lib directory of, 35, 43
- life cycle events, 239

## ServerLifecycleException, 240

## server

- optimizing for development, 28
- stopping an instance using an MBean, 302
- using asant script to control, 57-58
- value-added features, 166-169

## service method, 139

## servicename attribute, 61

## ServletContext.log messages, 134

## servlets, 133-139

- caching, 135-138
- character encoding, 152-153
- destroying, 138
- engine, 138
- instantiating, 138
- invoking using a URL, 134
- output, 134-135
- removing, 138
- request handling, 139
- specification, 133
  - class loading, 155
  - mime-mapping, 155
  - object unsupported for failover, 146

## session beans, 170

session beans (*Continued*)

- container for, 171-172
  - optimizing performance, 177
  - restrictions, 177
- session cache sharing and @OrderBy, 127
- session managers, 147-151
- session persistence
- for stateful session beans, 172-177
  - for web modules, 145-147
  - object types supported, 146-147, 173
- set command
- custom MBean attributes, 248
  - custom MBean disabling, 248
  - default message security provider, 93
  - default principal settings, 83
  - java-web-start-enabled attribute, 219
  - jbi-enabled property, 112
  - JMS service settings, 278
  - JMX connector port, 248
  - transaction service settings, 264
- setCharacterEncoding method, 153
- setContentype method, 153
- setLocale method, 153
- setMonitoring method, 301
- setTransactionIsolation method, 258
- Shared Chain class loader, 35
- SHUTDOWN\_EVENT, 239
- signing client JAR files, 221-223
- Simple Object Access Protocol, *See* SOAP messages
- single sign-on, 104-105
- Sitraka web site, 74-76
- SJSXP parser, 113
- SMTP protocol, 289-290
- SOAP messages, 285-287
- SOAP with Attachments API for Java (SAAJ), 285
- solaris realm, 84
- sourcedestdir attribute, 61, 62
- specification
- application clients, 214
  - connectors, 227
  - EJB 2.1 and CMP, 187
  - EJB 2.1 and JDOQL queries, 202
  - EJB 3.0, 165
  - JAAS, 85

- specification (*Continued*)
    - Java Persistence, 115
    - JavaBeans, 139
    - JDBC, 253
    - JMX, 243, 293
    - JSP, 139
    - Liberty Alliance Project, 93
    - programmatically security, 82
    - security manager, 89
    - servlet, 133
      - class loading, 37
    - WSS, 93
  - srcdir attribute, 58
  - stack trace, generating, 71
  - STARTUP\_EVENT, 239, 241
  - stateful session beans, 171
    - object references supported for failover, 173
    - session persistence, 172-177
  - stateless session beans, 171
  - StAX API, 113
  - stubs
    - keeping, 45, 66
  - sun-appserv-admin task, 57-58
  - sun-appserv-component task, 54-56
  - sun-appserv-deploy task, 44-48
  - sun-appserv-instance task, 51-54
  - sun-appserv-jspc task, 58-59
  - sun-appserv-undeploy task, 48-51
  - sun-appserv-update task, 59-60
  - sun-cmp-mappings.xml file, 189
  - sun-ejb-jar.xml file, 175, 176
  - Sun Java Studio, 31
  - Sun Java System Message Queue, 71, 276-277
    - checking to see if running, 279
    - connector for, 277
    - varhome directory, 284
  - sun-ra.xml file, 228
  - sun-web.xml file
    - and class loaders, 37, 155
  - supportsTransactionIsolationLevel method, 258
  - Sybase
    - finder limitation, 129, 208
    - lock-when-loaded limitation, 208
  - System class loader, 35
  - System class loader (*Continued*)
    - using to circumvent isolation, 40
  - system-classpath attribute, 35
- T**
- tag libraries, 140
  - tags for JSP caching, 140-143
  - target attribute, 47, 50, 55
  - tasks, asant script, 44-62
  - TERMINATION\_EVENT, 240
  - thread pools
    - and connectors, 231
    - for bean invocation scheduling, 169
  - timeout attribute, of cache tag, 142
  - tools, for developers, 29-32
  - Topic interface, 280
  - TopicConnectionFactory interface, 280
  - toplink.application-location property, 123
  - toplink.create-ddl-jdbc-file-name property, 123
  - toplink.ddl-generation.output-mode property, 124
  - toplink.ddl-generation property, 123
  - toplink.drop-ddl-jdbc-file-name property, 123
  - TopLink Essentials, *See* Oracle TopLink Essentials
  - toplink.platform.class.name property, 116
  - transaction-support property, 234
  - transactions, 261-267
    - administration and monitoring, 186
    - and EJB components, 184
    - and non-persistent JMS messages, 284
    - and session persistence, 172, 176
    - commit options, 185
    - configuring, 264
    - flat, 185
    - global, 185
    - in the Java EE tutorial, 261-267
    - JDBC isolation levels, 258
    - local, 185
    - local or global scope of, 262-263
    - logging for recovery, 265
    - logging to a database, 265-266
    - nested, 185
    - resource managers, 261-262
    - timeouts, 168

transactions (*Continued*)

- transaction manager, 264-265
- transaction synchronization registry, 264-265
- UserTransaction, 264-265

**U**

## undeploy command

- schema generation, 125, 197
- undeployment, using asant script, 48-51
- 
- uniquetablenames attribute, 46
- 
- upload attribute, 47, 64
- 
- uribase attribute, 58
- 
- uriroot attribute, 59
- 
- URL rewriting, 144
- 
- URL, redirecting, 159-160
- 
- use-thread-pool-id element, 169
- 
- use-unique-table-names property, 196
- 
- user attribute
- server element, 63
  - sun-appserv-component task, 55
  - sun-appserv-deploy task, 47
  - sun-appserv-instance task, 52
  - sun-appserv-undeploy task, 50
- utility classes, 38-39, 39

**V**

- valves, defining custom, 157-158
- 
- varhome directory, 284
- 
- verbose attribute, 58, 61, 62
- 
- verbose mode, 72
- 
- verify attribute, 45, 66
- 
- version consistency, 200
- triggers, 209
- virtual servers, 153-154
- default, 154
- virtualservers attribute, 47, 64

**W**

- web applications, 133-164

web applications (*Continued*)

- default, 134, 154
  - distributable, 145
  - security, 82
- Web class loader, 36
- changing delegation in, 37, 155
- web container, logging and monitoring, 156
- 
- web services, 107-114
- creating portable artifacts, 108
  - debugging, 108, 110
  - deployment, 108
  - in the Java EE tutorial, 107
  - Open ESB and JBI, 111-113
  - registry, 109-110
  - security
    - See* message security
  - test page, 110
  - URL, 110
  - WSDL file, 110
- webapp attribute, 58
- 
- WebDav, 160-162
- 
- Woodstox parser, 113
- 
- wsdl attribute, 62
- 
- wsdllocation attribute, 62
- 
- WSIT, 80
- 
- WSS,
- See*
- message security

**X**

- XA resource, 262-263
- 
- XML parser, 113
- specifying alternative, 39