



Sun GlassFish Enterprise Server v3 Prelude Developer's Guide



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

Part No: 820-4496-10
October 2008

Copyright 2008 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, Enterprise JavaBeans, EJB, GlassFish, J2EE, J2SE, Java Naming and Directory Interface, JavaBeans, Javadoc, JDBC, JDK, JavaScript, JavaServer, JavaServer Pages, JMX, JSP, JVM, MySQL, NetBeans, OpenSolaris, SunSolve, Sun GlassFish, 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 2008 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, Enterprise JavaBeans, EJB, GlassFish, J2EE, J2SE, Java Naming and Directory Interface, JavaBeans, Javadoc, JDBC, JDK, JavaScript, JavaServer, JavaServer Pages, JMX, JSP, JVM, MySQL, NetBeans, OpenSolaris, SunSolve, Sun GlassFish, 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

Preface	11
Part I Development Tasks and Tools	17
1 Setting Up a Development Environment	19
Installing and Preparing the Server for Development	19
The GlassFish Project	20
Development Tools	20
The asadmin Command	20
The Administration Console	21
The Migration Tool	21
The NetBeans IDE	21
The Eclipse IDE	21
Debugging Tools	22
Profiling Tools	22
2 Class Loaders	23
The Class Loader Hierarchy	24
Delegation	25
Using the Java Optional Package Mechanism	25
Using the Endorsed Standards Override Mechanism	25
Class Loader Universes	25
Application-Specific Class Loading	26
Circumventing Class Loader Isolation	27
Using the Common Class Loader	27

3	Debugging Applications	29
	Enabling Debugging	29
	▼ To Set the Server to Automatically Start Up in Debug Mode	30
	JPDA Options	30
	Generating a Stack Trace for Debugging	31
	Enabling Verbose Mode	31
	Profiling Tools	31
	The NetBeans Profiler	32
	The HPROF Profiler	32
	The JProbe Profiler	33
Part II	Developing Applications and Application Components	37
4	Securing Applications	39
	Security Goals	40
	Enterprise Server Specific Security Features	40
	Container Security	40
	Declarative Security	41
	Programmatic Security	42
	Roles, Principals, and Principal to Role Mapping	42
	Realm Configuration	43
	Supported Realms	44
	How to Configure a Realm	44
	How to Set a Realm for a Web Application or EJB Module	44
	Creating a Custom Realm	45
	JACC Support	47
	Pluggable Audit Module Support	47
	Configuring an Audit Module	47
	The AuditModule Class	47
	The server.policy File	49
	Default Permissions	49
	Changing Permissions for an Application	49
	Enabling and Disabling the Security Manager	51
	Programmatic Login	52
	Programmatic Login Precautions	53

Granting Programmatic Login Permission	53
The ProgrammaticLogin Class	54
User Authentication for Single Sign-on	55
5 Developing Web Services	57
Deploying a Web Service	58
Web Services Registry	58
The Web Service URI, WSDL File, and Test Page	59
6 Using the Java Persistence API	61
Specifying the Database	62
Additional Database Properties	64
Configuring the Cache	64
Setting the Logging Level	64
Using Lazy Loading	64
Primary Key Generation Defaults	65
Automatic Schema Generation	65
Annotations	66
Generation Options	66
Query Hints	67
Changing the Persistence Provider	67
Restrictions and Optimizations	68
Extended Persistence Context	68
Using @OrderBy with a Shared Session Cache	68
Using BLOB or CLOB Types with the Inet Oraxo JDBC Driver	69
Database Case Sensitivity	69
Sybase Finder Limitation	70
MySQL Database Restrictions	70
7 Developing Web Applications	73
Packaging an EJB JAR File in a Web Application	73
Using Servlets	74
Invoking a Servlet With a URL	74
Servlet Output	75

Caching Servlet Results	76
About the Servlet Engine	79
Using JavaServer Pages	80
JSP Tag Libraries and Standard Portable Tags	80
Options for Compiling JSP Files	80
Creating and Managing Sessions	81
Configuring Sessions	81
Session Managers	82
Using Comet	84
Introduction to Comet	84
Grizzly Comet	86
Bayeux Protocol	96
Developing Grails Applications	99
Introduction to Groovy and Grails	99
Installing Grails	99
Creating a Simple Grails Application	100
Deploying and Running a Grails Application	101
Advanced Web Application Features	103
Internationalization Issues	103
Virtual Servers	104
Default Web Modules	105
Class Loader Delegation	106
Using the default -web.xml File	106
Configuring Logging and Monitoring in the Web Container	107
Header Management	107
Configuring Valves and Catalina Listeners	107
Alternate Document Roots	108
Redirecting URLs	110
Using a context.xml File	110
Enabling WebDav	111
Using mod_jk	113
Using SSI	114
Using CGI	116
Using PHP	117
Using Scala and Lift	117

8	Using Enterprise JavaBeans Technology	119
	Summary of EJB 3.1 Changes	119
	Value Added Features	120
	Bean-Level Container-Managed Transaction Timeouts	121
	EJB Timer Service	121
	Using Session Beans	122
	About the Session Bean Containers	122
	Session Bean Restrictions and Optimizations	123
	Handling Transactions With Enterprise Beans	123
	Flat Transactions	124
	Global and Local Transactions	124
	Administration and Monitoring	124
Part III	Using Services and APIs	125
9	Using the JDBC API for Database Access	127
	General Steps for Creating a JDBC Resource	127
	Integrating the JDBC Driver	128
	Creating a Connection Pool	128
	Testing a JDBC Connection Pool	129
	Creating a JDBC Resource	129
	Creating Web Applications That Use the JDBC API	129
	Setting a Statement Timeout	129
	Sharing Connections	130
	Wrapping Connections	130
	Obtaining a Physical Connection From a Wrapped Connection	131
	Using the <code>Connection.unwrap()</code> Method	131
	Marking Bad Connections	131
	Using Non-Transactional Connections	132
	Using JDBC Transaction Isolation Levels	133
	Allowing Non-Component Callers	134
	Restrictions and Optimizations	135
	Disabling Stored Procedure Creation on Sybase	135

10 Using the Transaction Service	137
Transaction Scope	138
Configuring the Transaction Service	139
The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction	139
11 Using the Java Naming and Directory Interface	141
Accessing the Naming Context	141
Global JNDI Names	142
Using a Custom <code>jndi.properties</code> File	142
Mapping References	143
Index	145

Tables

TABLE 2-1	Sun GlassFish Enterprise Server Class Loaders	24
TABLE 7-1	URL Fields for Servlets Within an Application	75
TABLE 7-2	SSIServlet init-param Values	115
TABLE 7-3	CGIServlet init-param Values	116
TABLE 9-1	Transaction Isolation Levels	133

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 GlassFish™ Enterprise Server environment. Topics include developer tools, security, and debugging. 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.

The following topics are addressed here:

- “Enterprise Server Documentation Set” on page 11
- “Related Documentation” on page 13
- “Typographic Conventions” on page 13
- “Symbol Conventions” on page 14
- “Default Paths and File Names” on page 15
- “Documentation, Support, and Training” on page 15
- “Searching Sun Product Documentation” on page 15
- “Third-Party Web Site References” on page 16
- “Sun Welcomes Your Comments” on page 16

Enterprise Server Documentation Set

The Enterprise Server documentation set describes deployment planning and system installation. The Uniform Resource Locator (URL) for Enterprise Server documentation is <http://docs.sun.com/coll/1343.7>. 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>Release Notes</i>	Provides 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>	Explains how to get started with the Enterprise Server product.
<i>Installation Guide</i>	Explains how to install the software and its components.
<i>Application Deployment Guide</i>	Explains how to assemble and deploy applications to the Enterprise Server and provides information about deployment descriptors.
<i>Developer's Guide</i>	Explains how to create and implement Java Platform, Enterprise Edition (Java EE platform) applications that are intended to run on the Enterprise Server. These applications follow the open Java standards model for Java EE components and APIs. This guide provides information about developer tools, security, and debugging.
<i>Add-On Component Development Guide</i>	Explains how to use published interfaces of Enterprise Server to develop add-on components for Enterprise Server. This document explains how to perform <i>only</i> those tasks that ensure that the add-on component is suitable for Enterprise Server.
<i>RESTful Web Services Developer's Guide</i>	Explains how to develop Representational State Transfer (RESTful) web services for Enterprise Server.
<i>Getting Started With JRuby on Rails for Sun GlassFish Enterprise Server</i>	Explains how to develop Ruby on Rails applications for deployment to Enterprise Server.
<i>Getting Started With Project jMaki for Sun GlassFish Enterprise Server</i>	Explains how to use the jMaki framework to develop Ajax-enabled web applications that are centered on JavaScript™ technology for deployment to Enterprise Server.
<i>Roadmap to the Java EE 5 Tutorial</i>	Explains which information in the <i>Java EE 5 Tutorial</i> is relevant to users of the v3 Prelude release of the Enterprise Server.
<i>Java EE 5 Tutorial</i>	Explains how to use Java EE 5 platform technologies and APIs to develop Java EE applications.
<i>Java WSIT Tutorial</i>	Explains how to develop web applications by using the Web Service Interoperability Technologies (WSIT). The tutorial focuses on developing web service endpoints and clients that can interoperate with Windows Communication Foundation (WCF) endpoints and clients.

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

Book Title	Description
<i>Administration Guide</i>	Explains how to configure, monitor, and manage Enterprise Server subsystems and components from the command line by using the <code>asadmin(1M)</code> utility. Instructions for performing these tasks from the Administration Console are provided in the Administration Console online help.
<i>Administration Reference</i>	Describes the format of the Enterprise Server configuration file, <code>domain.xml</code> .
<i>Troubleshooting Guide</i>	Describes common problems that you might encounter when using Enterprise Server and how to solve them.
<i>Reference Manual</i>	Provides reference information in man page format for Enterprise Server administration commands, utility commands, and related concepts.

Related Documentation

A Javadoc™ tool reference for packages that are provided with the Enterprise Server is located at <https://glassfish.dev.java.net/nonav/api/v3-prelude/index.html>. Additionally, the following resources might be useful:

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

For information about creating enterprise applications in the NetBeans™ Integrated Development Environment (IDE), see <http://www.netbeans.org/kb/60/index.html>.

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

Typographic Conventions

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

TABLE P-2 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-2 Typographic Conventions (Continued)

Typeface	Meaning	Example
AaBbCc123	What you type, contrasted with onscreen computer output	machine_name% su 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-3 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.

Default Paths and File Names

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

TABLE P-4 Default Paths and File Names

Placeholder	Description	Default Value
<i>as-install</i>	Represents the base installation directory for Enterprise Server. In configuration files, <i>as-install</i> is represented as follows: \${com.sun.aas.installRoot}	Installations on the Solaris™ operating system, Linux operating system, and Mac operating system: <i>user's-home-directory/glassfishv3-prelude/glassfish</i> Windows, all installations: <i>SystemDrive:\glassfishv3-prelude\glassfish</i>
<i>domain-root-dir</i>	Represents the directory in which a domain is created by default.	<i>as-install/domains/</i>
<i>domain-dir</i>	Represents the directory in which a domain's configuration is stored. In configuration files, <i>domain-dir</i> is represented as follows: \${com.sun.aas.instanceRoot}	<i>domain-root-dir/domain-name</i>

Documentation, Support, and Training

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

- [Documentation \(http://www.sun.com/documentation/\)](http://www.sun.com/documentation/)
- [Support \(http://www.sun.com/support/\)](http://www.sun.com/support/)
- [Training \(http://www.sun.com/training/\)](http://www.sun.com/training/)

Searching Sun Product Documentation

Besides searching Sun product documentation from the docs.sun.comSM web site, you can use a search engine by typing the following syntax in the search field:

```
search-term site:docs.sun.com
```

For example, to search for “broker,” type the following:

```
broker site:docs.sun.com
```

To include other Sun web sites in your search (for example, java.sun.com, www.sun.com, and developers.sun.com), use `sun.com` in place of `docs.sun.com` in the search field.

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 Send Comments. In the online form, provide the full document title and part number. The part number is a 7-digit or 9-digit number that can be found on the book's title page or in the document's URL. For example, the part number of this book is 820-4496.

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 GlassFish™ 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 19
- “The GlassFish Project” on page 20
- “Development Tools” on page 20

Installing and Preparing the Server for Development

For more information about stand-alone Enterprise Server installation, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

The following components are included in the full installation.

- JDK
- Enterprise Server core
 - Java Platform, Standard Edition (Java SE) 6
 - Java EE 6 compliant application server
 - Administration Console
 - `asadmin` utility
 - Other development and deployment tools
 - Java DB database, based on the [Derby database from Apache](http://db.apache.org/derby/manuals) (<http://db.apache.org/derby/manuals>)

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 Common Class Loader” on page 27.
- Set up debugging. For more information, see Chapter 3, “Debugging Applications.”
- Configure the Virtual Machine for the Java™ platform (JVM™ software). For more information, see Chapter 4, “Administering the Virtual Machine for the Java Platform,” in *Sun GlassFish Enterprise Server v3 Prelude Administration Guide*.

The GlassFish Project

Enterprise Server v3 Prelude 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.

Development Tools

The following general tools are provided with the Enterprise Server:

- “The `asadmin` Command” on page 20
- “The Administration Console” on page 21

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

- “The Migration Tool” on page 21
- “The NetBeans IDE” on page 21

The following third-party tools might also be useful:

- “The Eclipse IDE” on page 21
- “Debugging Tools” on page 22
- “Profiling Tools” on page 22

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 v3 Prelude Reference Manual*.

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

The Administration Console

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

To access the Administration Console, type `http://host:4848` 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 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>.

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 6.1 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 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.

Debugging Tools

You can use several debugging tools with the Enterprise Server. For more information, see [Chapter 3, “Debugging Applications.”](#)

Profiling Tools

You can use several profilers with the Enterprise Server. For more information, see [“Profiling Tools” on page 31.](#)

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 JVM implementation, 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 24
- “Delegation” on page 25
- “Using the Java Optional Package Mechanism” on page 25
- “Using the Endorsed Standards Override Mechanism” on page 25
- “Class Loader Universes” on page 25
- “Application-Specific Class Loading” on page 26
- “Circumventing Class Loader Isolation” on page 27

Note – For GlassFish v3 Prelude, EJB modules are not supported unless the optional EJB container add-on component is downloaded from the Update Tool. Web services are not supported unless the optional Metro (JSR 109) add-on component is downloaded from the Update Tool. For information about the Update Tool, see the [Sun GlassFish Enterprise Server v3 Prelude Installation Guide](#).

The Class Loader Hierarchy

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

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 software.
Extension	The Extension class loader loads classes from JAR files present in the system extensions directory, <i>domain-dir/lib/ext</i> . It is parent to the Public API class loader. See “Using the Java Optional Package Mechanism” on page 25.
Public API	The Public API class loader makes available all classes specifically exported by the Enterprise Server runtime for use by deployed applications. This includes, but is not limited to, Java EE APIs and other Sun GlassFish APIs. It is parent to the Common class loader.
Common	The Common class loader loads JAR files in the <i>as-install/lib</i> directory, then classes in the <i>domain-dir/lib/classes</i> directory, followed by JAR files in the <i>domain-dir/lib</i> directory. Using <i>domain-dir/lib/classes</i> or <i>domain-dir/lib</i> is recommended whenever possible, and required for custom login modules and realms. It is parent to the Connector class loader. See “Using the Common Class Loader” on page 27.
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 Applib class loader.
Applib	<p>The Applib class loader loads the library classes, specified during deployment, for a specific enabled module; see “Application-Specific Class Loading” on page 26. One instance of this class loader is present in each class loader universe; see “Class Loader Universes” on page 25. It is parent to the Archive class loader.</p> <p>When multiple deployed applications use the same library, they share the same instance of the library. One library cannot reference classes from another library.</p>
Archive	The Archive class loader loads classes from the WAR and JAR files or directories (for directory deployment) of modules deployed to the Enterprise Server. This class loader also loads any application-specific classes generated by the Enterprise Server runtime, such as stub classes or servlets generated by JSP pages.

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 class loading to its parent before attempting to load a class itself. 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.

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) (<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) (<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 modules installed on the server occurs within the context of isolated class loader universes, each of which has its own Applet and Archive class loaders.

- **Individually Deployed Module Universe** – Each individually deployed EJB JAR or web WAR 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 module's classpath; for example, a web module's classpath includes these directories:

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

Application-Specific Class Loading

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

- Use the Administration Console. Open the Applications component, then go to the page for the type of module. Select the Deploy button. Type the comma-separated paths in the Libraries field. For details, click the Help button in the Administration Console.
- Use the `asadmin deploy` command with the `--libraries` option and specify comma-separated paths. For details, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

Application libraries are included in the `Applib` 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).

If multiple 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, modules using application-specific libraries are not portable. Other ways to make libraries available are described in “[Circumventing Class Loader Isolation](#)” on page 27.

One library cannot reference classes from another library.

For general information about deployment, see the [Sun GlassFish Enterprise Server v3 Prelude 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 49.

Circumventing Class Loader Isolation

Since each module class loader universe is isolated, a module cannot load classes from another module. This prevents two similarly named classes in different modules from interfering with each other.

To circumvent this limitation for libraries, utility classes, or individually deployed modules, you can include the relevant path to the required classes. See [“Using the Common Class Loader” on page 27](#).

Using the Common Class Loader

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

Using the Common class loader makes a module accessible to all 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 v3 Prelude Release Notes](#). For configurations of supported and other drivers, see [“Configuration Specifics for JDBC Drivers” in Sun GlassFish Enterprise Server v3 Prelude Administration Guide](#).

To activate custom login modules and realms, place the JAR files in the *domain-dir/lib* directory or the class files in the *domain-dir/lib/classes* directory, then restart the server.

Debugging Applications

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

- “Enabling Debugging” on page 29
- “JPDA Options” on page 30
- “Generating a Stack Trace for Debugging” on page 31
- “Enabling Verbose Mode” on page 31
- “Profiling Tools” on page 31

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 30.

You can attach to the Enterprise Server using any JPDA compliant debugger, including that of NetBeans (<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 Administration Console. Select the Enterprise Server component and the JVM Settings tab.
- 2 Check the Debug Enabled box.
- 3 To specify a different port (from 9009, the default) to use when attaching the JVM software 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 30.

See Also For details, click the Help button in the Administration 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 software 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 software starts.

To specify a different port (from 9009, the default) to use when attaching the JVM software 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 v3 Prelude Reference Manual*.

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]
```

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`.

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 32
- “The HPROF Profiler” on page 32
- “The JProbe Profiler” on page 33

Information about comprehensive monitoring and management support in the Java™ 2 Platform, Standard Edition (JSE™ 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.

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 Administration Console. Select the Enterprise Server component and the JVM Settings tab. 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 Administration Console:

a. Select the Enterprise Server component and the JVM Settings tab. 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.

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 6, “Administering System Security,”](#) in *Sun GlassFish Enterprise Server v3 Prelude 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 40
- “Enterprise Server Specific Security Features” on page 40
- “Container Security” on page 40
- “Roles, Principals, and Principal to Role Mapping” on page 42
- “Realm Configuration” on page 43
- “JACC Support” on page 47
- “Pluggable Audit Module Support” on page 47
- “The `server.policy` File” on page 49
- “Programmatic Login” on page 52
- “User Authentication for Single Sign-on” on page 55

Note – For GlassFish v3 Prelude, EJB modules are not supported unless the optional EJB container add-on component is downloaded from the Update Tool. Web services are not supported unless the optional Metro (JSR 109) add-on component is downloaded from the Update Tool. For information about the Update Tool, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

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 several underlying authentication realms.
- 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/>).

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:

- Single sign-on across all Enterprise Server applications within a single security domain; see [“User Authentication for Single Sign-on” on page 55](#)
- Programmatic login; see [“Programmatic Login” on page 52](#)

Container Security

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

- [“Declarative Security” on page 41](#)
- [“Programmatic Security” on page 42](#)

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 v3 Prelude Application Deployment Guide*.

There are two levels of declarative security, as follows:

- [“Application Level Security” on page 41](#)
- [“Component Level Security” on page 41](#)

Application Level Security

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-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 42](#).

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 52](#)

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 v3 Prelude 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. The default principal to role mapping maps group principals to the same 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 Administration Console, select the Security component under the relevant configuration. For details, click the Help button in the Administration Console.
- Use the `asadmin set` command. For details, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#). For example, you can set the default principal as follows.

```
asadmin set --user adminuser server-config.security-service.default-principal=dsmith
asadmin set --user adminuser server-config.security-service.default-principal-password=secret
```

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

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

Realm Configuration

This section covers the following topics:

- “Supported Realms” on page 44
- “How to Configure a Realm” on page 44
- “How to Set a Realm for a Web Application or EJB Module” on page 44
- “Creating a Custom Realm” on page 45

Supported Realms

The following realms are supported in the current release of 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 Enterprise 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.

- `certificate` – Sets up the user identity in the Enterprise Server security context, and populates it with user data obtained from cryptographically verified client certificates.

Note – The `solaris` realm is not supported in GlassFish v3 Prelude.

For information about configuring realms, see [“How to Configure a Realm”](#) on page 44.

How to Configure a Realm

You can configure a realm in one of these ways:

- In the Administration Console, open the Security component under the relevant configuration and go to the Realms page. For details, click the Help button in the Administration Console.
- Use the `asadmin create-auth-realm` command to configure realms on local servers. For details, see the *Sun GlassFish Enterprise Server v3 Prelude Reference Manual*.

How to Set a Realm for a Web Application or EJB Module

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

- `web-app` element in `web.xml`
- `as-context` element in `sun-ejb-jar.xml`

For more information about the deployment descriptor files and elements, see [Appendix A, “Deployment Descriptor Files,”](#) in *Sun GlassFish Enterprise Server v3 Prelude 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.

To activate the custom login modules and realms, place the JAR files in the *domain-dir/lib* directory or the class files in the *domain-dir/lib/classes* directory, then restart the server. For more information about class loading in the Enterprise Server, see [Chapter 2, “Class Loaders.”](#)

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 <https://glassfish.dev.java.net/nonav/api/v3-prelude/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(_username, _password,
    _currentRealm, 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.

Note – The array passed to the `commitUseAuthentication` method should be newly created and otherwise unreferenced. This is because the group name array elements are set to null after authentication as part of cleanup. So the second time your custom realm executes it returns an array with null elements.

Ideally, your custom realm should not return member variables from the `authenticate` method. It should return local variables as the default `JDBCRealm` does. Your custom realm can create a local `String` array in its `authenticate` method, copy the values from the member variables, and return the `String` array. Or it can use `clone` on the member variables.

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, named `default`. An alternate provider named `simple` is also provided. To configure an alternate provider using the Administration Console, open the Security component under the relevant configuration, and select the JACC Providers component. For details, click the Help button in the Administration 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 47
- “The AuditModule Class” on page 47

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).

Configuring an Audit Module

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

- To specify an audit module using the Administration Console, open the Security component under the relevant configuration, and select the Audit Modules component. For details, click the Help button in the Administration Console.
- You can use the `asadmin create-audit-module` command to configure an audit module. For details, see the *Sun GlassFish Enterprise Server v3 Prelude Reference Manual*.

The AuditModule Class

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

For Javadoc tool pages relevant to audit modules, go to <https://glassfish.dev.java.net/nonav/api/v3-prelude/index.html> and click on the `com.sun.appserv.security.audit` 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 49
- “Changing Permissions for an Application” on page 49
- “Enabling and Disabling the Security Manager” on page 51

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/-" {  
    ...  
}
```

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/MyApp/-" {  
    ...  
}
```

For a module:

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

Note – Deployment directories may change between Enterprise Server releases.

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/cejug-clfds/ ... )  
...
```

Granting the following permission eliminates the message:

```
grant codeBase "file:${com.sun.aas.instanceRoot}/applications/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 v3 Prelude Reference Manual*.

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/>

[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 v3 Prelude 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 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 Administration 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 Administration 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 v3 Prelude Reference Manual*.

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 53
- “Granting Programmatic Login Permission” on page 53
- “The `ProgrammaticLogin` Class” on page 54

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 49](#). 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 <https://glassfish.dev.java.net/nonav/api/v3-prelude/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 8, “Administering the HTTP Service,” in *Sun GlassFish Enterprise Server v3 Prelude 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 `false`.
- `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. This chapter contains the following sections:

- “Deploying a Web Service” on page 58
- “Web Services Registry” on page 58
- “The Web Service URI, WSDL File, and Test Page” on page 59

Note – For GlassFish v3 Prelude, web services are not supported unless the optional Metro (JSR 109) add-on component is downloaded from the Update Tool. Without the Metro add-on component, a servlet cannot be a web service endpoint, and the `sun-web.xml` elements related to web services are ignored. For information about the Update Tool, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

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>).

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>)

Deploying a Web Service

You deploy a web service endpoint to the Enterprise Server just as you would any servlet. 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 58](#).

You can use the autodeployment feature to deploy a simple [JSR 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 webservices.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 the `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 59](#).

For more information about deployment, autodeployment, and deployment descriptors, see the [Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide](#). For more information about the `asadmin deploy` command, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

Web Services Registry

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

- In the Administration Console, open the Web Services component, and select the Registry tab. For details, click the Help button in the Administration 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 v3 Prelude Reference Manual*.

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

- In the Administration 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 Administration 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 v3 Prelude Reference Manual*.

For more information about registries, 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 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 SOA registry available at <http://www.sun.com/products/soa/index.jsp>.

For further information, see <https://metro.dev.java.net/>.

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 `web.xml` file, and can be overridden in the `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 Administration 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 Administration 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 v3 Prelude Application Deployment Guide](#).

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 62
- “Additional Database Properties” on page 64
- “Configuring the Cache” on page 64
- “Setting the Logging Level” on page 64
- “Using Lazy Loading” on page 64
- “Primary Key Generation Defaults” on page 65
- “Automatic Schema Generation” on page 65
- “Query Hints” on page 67
- “Changing the Persistence Provider” on page 67
- “Restrictions and Optimizations” on page 68

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

Note – For GlassFish v3 Prelude, EJB modules are not supported unless the optional EJB container add-on component is downloaded from the Update Tool. For information about the Update Tool, see the [Sun GlassFish Enterprise Server v3 Prelude Installation Guide](#).

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 type based on the connection metadata. This database type is used to issue SQL statements specific to the detected database type's dialect. You can specify the optional `eclipselink.target-database` property to guarantee that the database type 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="eclipselink.target-database"
          value="DB2"/>
      </properties>
    </persistence-unit>
  </persistence>
```

The following `eclipselink.target-database` property values are allowed. Supported platforms have been tested with the Enterprise Server and are found to be Java EE compatible.

```
//Supported platforms
JavaDB
Derby
Oracle
MySQL4
//Others available
SQLServer
```

DB2
 Sybase
 PostgreSQL
 Informix
 TimesTen
 Attunity
 HSQL
 SQLAnywhere
 DBase
 DB2Mainframe
 Cloudscape
 PointBase

For more information about the `eclipselink.target-database` property, see [Using EclipseLink JPA Extensions for Session, Target Database and Target Application Server](#).

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. 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" version="1.0">
    <persistence-unit name="em2">
      <provider>org.eclipse.persistence.jpa.PersistenceProvider</provider>
      <class>ejb3.war.servlet.JpaBean</class>
      <properties>
        <property name="eclipselink.target-database"
          value="Derby"/>
        <!-- JDBC connection properties -->
        <property name="eclipselink.jdbc.driver" value="org.apache.derby.jdbc.ClientDriver"/>
        <property name="eclipselink.jdbc.url"
value="jdbc:derby://localhost:1527/testdb;retrieveMessagesFromServerOnGetMessage=true;create=true;"/>
        <property name="eclipselink.jdbc.user" value="APP"/>
        <property name="eclipselink.jdbc.password" value="APP"/>
      </properties>
    </persistence-unit>
  </persistence>
```

For more information about `eclipselink` properties, see [“Additional Database Properties” on page 64](#).

For a list of the JDBC drivers currently supported by the Enterprise Server, see the [Sun GlassFish Enterprise Server v3 Prelude Release Notes](#). For configurations of supported and other drivers, see [“Configuration Specifics for JDBC Drivers” in Sun GlassFish Enterprise Server v3 Prelude Administration Guide](#).

To change the persistence provider, see [“Changing the Persistence Provider” on page 67](#).

Additional Database Properties

If you are using the default persistence provider, you can specify in the `persistence.xml` file the database properties listed at [How to Use EclipseLink JPA Extensions for JDBC Connection Communication](#).

For schema generation properties, see “[Generation Options](#)” on page 66. For query hints, see “[Query Hints](#)” on page 67.

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 [Using EclipseLink JPA Extensions for Entity Caching](#).

Setting the Logging Level

One of the default persistence provider's properties that you can set in the `persistence.xml` file is `eclipselink.logging.level`. For example, setting the logging level to `FINE` or higher logs all SQL statements. For details about this property, see [Using EclipseLink JPA Extensions for 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 basic information about lazy loading, see [What You May Need to Know About EclipseLink JPA Lazy Loading](#).

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 `table` 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 66
- “Generation Options” on page 66

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.

Schema generation occurs whenever an application is loaded. For example, schema generation occurs if the Enterprise Server is restarted with the application deployed.

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.

Generation Options

Optional schema generation properties control the automatic creation of database tables. You can specify them in the `persistence.xml` file. For more information, see [Using EclipseLink JPA Extensions for Schema Generation](#).

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("eclipselink.refresh", true).  
    getSingleResult();
```

For more information about the query hints available with the default provider, see [How to Use EclipseLink JPA Query Hints](#).

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.

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 27. 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>
```

To migrate from Oracle TopLink to EclipseLink, see [Migrating from Oracle TopLink to EclipseLink](http://wiki.eclipse.org/EclipseLink/Examples/MigratingFromOracleTopLink) (<http://wiki.eclipse.org/EclipseLink/Examples/MigratingFromOracleTopLink>).

Restrictions and Optimizations

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

- “Extended Persistence Context” on page 68
- “Using @OrderBy with a Shared Session Cache” on page 68
- “Using BLOB or CLOB Types with the Inet Oraxo JDBC Driver” on page 69
- “Database Case Sensitivity” on page 69
- “Sybase Finder Limitation” on page 70
- “MySQL Database Restrictions” on page 70

Extended Persistence Context

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.

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="eclipselink.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.SySQLException: 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:

- “Packaging an EJB JAR File in a Web Application” on page 73
- “Using Servlets” on page 74
- “Using JavaServer Pages” on page 80
- “Creating and Managing Sessions” on page 81
- “Using Comet” on page 84
- “Developing Grails Applications” on page 99
- “Advanced Web Application Features” on page 103

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) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

Packaging an EJB JAR File in a Web Application

The Enterprise Server supports the EJB 3.1 specification, which allows EJB JAR files to be packaged in WAR files. EJB classes must reside under `WEB-INF/classes`. For example, the structure of a `hello.war` file might look like this:

```
index.jsp
META-INF/
  MANIFEST.MF
WEB-INF/
  web.xml
  classes/
    com/
      sun/
        v3/
          demo/
            HelloEJB.class
            HelloServlet.class
```

For more information about EJB components, see [Chapter 8, “Using Enterprise JavaBeans Technology.”](#)

Note – For GlassFish v3 Prelude, EJB modules are not supported unless the optional EJB container add-on component is downloaded from the Update Tool. For information about the Update Tool, see the [Sun GlassFish Enterprise Server v3 Prelude Installation Guide](#).

For GlassFish v3 Prelude, only stateless session beans with local interfaces and entity beans that use the Java Persistence API are supported. Stateful, message-driven, and EJB 2.0 and 2.1 entity beans are not supported. Remote interfaces and remote business interfaces for any of the bean types are not supported.

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 v3 Prelude 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 74
- “Servlet Output” on page 75
- “Caching Servlet Results” on page 76
- “About the Servlet Engine” on page 79

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 7-1 URL Fields for Servlets Within an Application

URL element	Description
<code>server:port</code>	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 <code>context-root</code> or <code>servlet-name</code> unless you also wish to specify name-value parameters.
<code>context-root</code>	For an application, the context root is defined in the <code>context-root</code> element of the <code>application.xml</code> , <code>sun-application.xml</code> , or <code>sun-web.xml</code> 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 <code>.war</code> suffix.
<code>servlet-mapping</code>	The <code>servlet-mapping</code> as configured in the <code>web.xml</code> file.
<code>?name=value...</code>	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 Administration Console. Select the Enterprise Server component and the Logging tab. 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 Administration 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 <https://glassfish.dev.java.net/nonav/api/v3-prelude/index.html> and click on the `com.sun.appserv.web.cache` package.

The rest of this section covers the following topics:

- “Caching Features” on page 76
- “Default Cache Configuration” on page 77
- “Caching Example” on page 77
- “The CacheKeyGenerator Interface” on page 79

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 79](#).
- 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.
- 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 v3 Prelude Application Deployment Guide*](#) and [“dispatcher” in *Sun GlassFish Enterprise Server v3 Prelude 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 v3 Prelude 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 v3 Prelude 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 79](#)
- [“Request Handling” on page 79](#)

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
- 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 80
- “Options for Compiling JSP Files” on page 80

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.

Options for Compiling JSP Files

Enterprise Server provides the following ways of compiling JSP 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 v3 Prelude Reference Manual*.

- The `jspc` command line tool allows you to precompile JSP files at the command line. For details, see the *Sun GlassFish Enterprise Server v3 Prelude 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 81
- “Session Managers” on page 82

Configuring Sessions

This section covers the following topics:

- “HTTP Sessions, Cookies, and URL Rewriting” on page 81
- “Coordinating Session Access” on page 81
- “Saving Sessions During Redeployment” on page 81

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 v3 Prelude Application Deployment Guide* and “`cookie-properties`” in *Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide*.

Coordinating Session Access

Make sure that multiple threads don’t simultaneously modify the same session object in conflicting ways.

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.

Saving Sessions During Redeployment

Whenever a redeployment is done, the sessions at that transit time become invalid unless you use the `keepSessions=true` property of the `asadmin redeploy` command. For example:

```
asadmin redeploy --properties keepSessions=true --name hello.war
```

For details, see the *Sun GlassFish Enterprise Server v3 Prelude Reference Manual*.

The new class loader of the redeployed application is used to deserialize any sessions previously saved. The usual restrictions about serialization and deserialization apply. For example, any application-specific class referenced by a session attribute may evolve only in a backward-compatible fashion. For more information about class loaders, see [Chapter 2, “Class Loaders.”](#)

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 82, the default
- “The file Persistence Type” on page 83, which uses a file to store session data

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 v3 Prelude 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 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 v3 Prelude Application Deployment Guide*. The `sessionFileName` property specifies the name of the file where sessions are serialized and persisted if the web application or the server is stopped. To disable this behavior, specify an empty string as the value of `sessionFileName`.

For more information about the `sun-web.xml` file, see *Sun GlassFish Enterprise Server v3 Prelude 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 49.

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 v3 Prelude Application Deployment Guide* except `sessionFilename`, and supports the directory store property listed under “store-properties” in *Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide*.

For more information about the `sun-web.xml` file, see *Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide*.

Using Comet

This section explains the Comet programming technique and how to create and deploy a Comet-enabled application with the Sun GlassFish Enterprise Server.

Introduction to Comet

Comet is a programming technique that allows a web server to send updates to clients without requiring the clients to explicitly request them.

This kind of programming technique is called *server push*, which means that the server pushes data to the client. The opposite style is *client pull*, which means that the client must pull the data from the server, usually through a user-initiated event, such as a button click.

Web applications that use the Comet technique can deliver updates to clients in a more timely manner than those that use the client-pull style while avoiding the latency that results from clients frequently polling the server.

One of the many use cases for Comet is a chat room application. When the server receives a message from one of the chat clients, it needs to send the message to the other clients without requiring them to ask for it. With Comet, the server can deliver messages to the clients as they are posted rather than expecting the clients to poll the server for new messages.

To accomplish this scenario, a Comet application establishes a long-lived HTTP connection. This connection is suspended on the server side, waiting for an event to happen before being resumed. This kind of connection remains open, allowing an application that uses the Comet technique to send updates to clients when they are available rather than expecting clients to reopen the connection to poll the server for updates.

The Grizzly Implementation of Comet

One limitation of the Comet technique is that you must use it with a web server that supports non-blocking connections in order to avoid poor performance. Non-blocking connections are those that do not need to allocate one thread for each request. If the web server were to use blocking connections then it might end up holding many thousands of threads, thereby hindering its scalability.

The GlassFish server includes the Grizzly HTTP Engine, which enables asynchronous request processing (ARP) by avoiding blocking connections. Grizzly's ARP implementation accomplishes this by using the Java NIO API.

With Java NIO, Grizzly enables greater performance and scalability by avoiding the limitations experienced by traditional web servers that must run a thread for each request. Instead, Grizzly's ARP mechanism makes efficient use of a thread pool system and also keeps the state of requests so that it can keep requests alive without holding a single thread for each of them.

Grizzly supports two different implementations of Comet:

- [“Grizzly Comet” on page 86](#) — Based on ARP, this includes a set of APIs that you use from a web component to enable Comet functionality in your web application. Grizzly Comet is specific to the Sun GlassFish Enterprise Server.
- [“Bayeux Protocol” on page 96](#) — Often referred to as Cometd. This consists of the JSON-based Bayeux message protocol, a set of Dojo or Ajax libraries, and an event handler. The Bayeux protocol uses a publish/subscribe model for server/client communication. The Bayeux protocol is portable, but it is container dependent if you want to invoke it from an EJB component. The Grizzly implementation of Cometd consists of a servlet that you reference from your web application.

Client Technologies to Use With Comet

In addition to creating a web component that uses the Comet APIs, you need to enable your client to accept asynchronous updates from the web component. To accomplish this, you can use JavaScript, IFrames, or a framework, such as [Dojo](#).

An IFrame is an HTML element that allows you to include other content in an HTML page. As a result, the client can embed updated content in the IFrame without having to reload the page.

The example explained in this tutorial employs a combination of JavaScript and IFrames to allow the client to accept asynchronous updates. A servlet included in the example writes out JavaScript code to one of the IFrames. The JavaScript code contains the updated content and invokes a function in the page that updates the appropriate elements in the page with the new content.

The next section explains the two kinds of connections that you can make to the server. While you can use any of the client technologies listed in this section with either kind of connection, it is more difficult to use JavaScript with an HTTP-streaming connection.

Kinds of Comet Connections

When working with Comet, as implemented in Grizzly, you have two different ways to handle client connections to the server:

- HTTP Streaming
- long-polling

HTTP Streaming

The HTTP Streaming technique keeps a connection open indefinitely. It never closes, even after the server pushes data to the client.

In the case of HTTP streaming, the application sends a single request and receives responses as they come, reusing the same connection forever. This technique significantly reduces the network latency because the client and the server don't need to open and close the connection.

The basic life cycle of an application using HTTP-streaming is:

request --> suspend --> data available --> write response --> data available --> write response

The client makes an initial request and then suspends the request, meaning that it waits for a response. Whenever data is available, the server writes it to the response.

Long Polling

The long-polling technique is a combination of server-push and client-pull because the client needs to resume the connection after a certain amount of time or after the server pushes an update to the client.

The basic life cycle of an application using long-polling is:

request -> suspend --> data available --> write response --> resume

The client makes an initial request and then suspends the request. When an update is available, the server writes it to the response. The connection closes, and the client optionally resumes the connection.

How to Choose the Kind of Connection

If you anticipate that your web application will need to send frequent updates to the client, you should use the HTTP-streaming connection so that the client does not have to frequently reestablish a connection. If you anticipate less frequent updates, you should use the long-polling connection so that the web server does not need to keep a connection open when no updates are occurring. One caveat to using the HTTP-streaming connection is that if you are streaming through a proxy, the proxy can buffer the response from the server. So, be sure to test your application if you plan to use HTTP-streaming behind a proxy.

Grizzly Comet

The following sections describe how to use Grizzly Comet.

- [“The Grizzly Comet API” on page 87](#)
- [“The Hidden Frame Example” on page 87](#)

- “Creating a Comet-Enabled Application” on page 88
- “Developing the Web Component” on page 89
- “Creating the Client Pages” on page 92
- “Creating the Deployment Descriptor” on page 94
- “Deploying and Running a Comet-Enabled Application” on page 95

The Grizzly Comet API

Grizzly’s support for Comet includes a small set of APIs that make it easy to add Comet functionality to your web applications. The Grizzly Comet APIs that developers will use most often are the following:

- `CometContext`: A Comet context, which is a shareable space to which applications subscribe in order to receive updates.
- `CometEngine`: The entry point to any component using Comet. Components can be servlets, JavaServer Pages™ (JSP™), JavaServer Faces components, or pure Java classes.
- `CometEvent`: Contains the state of the `CometContext` object
- `CometHandler`: The interface an application implements to be part of one or more Comet contexts.

The way a developer would use this API in a web component is to perform the following tasks:

1. Register the context path of the application with the `CometContext` object:

```
CometEngine cometEngine =
    CometEngine.getEngine();
CometContext cometContext =
    cometEngine.register(contextPath)
```

2. Register the `CometHandler` implementation with the `CometContext` object:

```
cometContext.addCometHandler(handler)
```

3. Notify one or more `CometHandler` implementations when an event happens:

```
cometContext.notify((Object) handler)
```

The Hidden Frame Example

This rest of this tutorial uses the Hidden Frame example to explain how to develop Comet-enabled web applications. You can download the example from `grizzly.dev.java.net` at [Hidden example download](#). From there, you can download a prebuilt WAR file as well as a JAR file containing the servlet code.

The Hidden Frame example is so called because it uses hidden IFrames. What the example does is it allows multiple clients to increment a counter on the server. When a client increments the counter, the server broadcasts the new count to the clients using the Comet technique.

The Hidden Frame example uses the long-polling technique, but you can easily modify it to use HTTP-streaming by removing two lines. See [“Notifying the Comet Handler of an Event” on page 91](#) and [“Creating the HTML Page That Updates and Displays the Content” on page 93](#) for more information on converting the example to use the HTTP-streaming technique.

The client side of the example uses hidden IFrames with embedded JavaScript tags to connect to the server and to asynchronously post content to and accept updates from the server.

The server side of the example consists of a single servlet that listens for updates from clients, updates the counter, and writes JavaScript code to the client that allows it to update the counter on its page.

See [“Deploying and Running a Comet-Enabled Application” on page 95](#) for instructions on how to deploy and run the example.

When you run the example, the following happens:

1. The `index.html` page opens.
2. The browser loads three frames: the first one accesses the servlet using an HTTP GET; the second one loads the `count.html` page, which displays the current count; and the third one loads the `button.html` page, which is used to send the POST request.
3. After clicking the button on the `button.html` page, the page submits a POST request to the servlet.
4. The `doPost` method calls the `onEvent` method of the Comet handler and redirects the incremented count along with some JavaScript to the `count.html` page on the client.
5. The `updateCount` JavaScript function on the `count.html` page updates the counter on the page.
6. Because this example uses long-polling, the JavaScript code on `count.html` calls `doGet` again to resume the connection after the servlet pushes the update.

Creating a Comet-Enabled Application

This section uses the Hidden Frame example application to demonstrate how to develop a Comet application. The main tasks for creating a simple Comet-enabled application are the following:

- [“Developing the Web Component” on page 89](#), such as a servlet to support the Comet requests and a Comet handler to send updates to the client
- [“Creating the Client Pages” on page 92](#), one or more HTML pages that include some client-side technology to open an asynchronous connection to the server and to receive updates from the web component
- [“Creating the Deployment Descriptor” on page 94](#), which configures the web component

Developing the Web Component

This section shows you how to create a Comet-enabled web component by giving you instructions for creating the servlet in the Hidden Frame example.

Developing the web component involves performing the following steps:

1. Create a web component to support Comet requests.
2. Register the component with the Comet engine.
3. Define a Comet handler that sends updates to the client.
4. Add the Comet handler to the Comet context.
5. Notify the Comet handler of an event using the Comet context.

▼ Creating a Web Component to Support Comet

1 Create an empty servlet class, like the following:

```
import javax.servlet.*;

public class HiddenCometServlet extends HttpServlet {
    private static final long serialVersionUID = 1L;
    private String contextPath = null;
    @Override
    public void init(ServletConfig config) throws ServletException {}

    @Override
    protected void doGet(HttpServletRequest req,
        HttpServletResponse res)
        throws ServletException, IOException {}

    @Override
    protected void doPost(HttpServletRequest req,
        HttpServletResponse res)
        throws ServletException, IOException {}
}
```

2 Import the following Comet packages into the servlet class:

```
import com.sun.grizzly.comet.CometContext;
import com.sun.grizzly.comet.CometEngine;
import com.sun.grizzly.comet.CometEvent;
import com.sun.grizzly.comet.CometHandler;
```

3 Import these additional classes that you need for incrementing a counter and writing output to the clients:

```
import java.io.IOException;
import java.io.PrintWriter;
import java.util.concurrent.atomic.AtomicInteger;
```

4 Add a private variable for the counter:

```
private final AtomicInteger counter = new AtomicInteger();
```

▼ Registering the Servlet with the Comet Engine**1 In the servlet's `init` method, add the following code to get the component's context path:**

```
ServletContext context = config.getServletContext();
contextPath = context.getContextPath() + "/hidden_comet";
```

2 Get an instance of the Comet engine by adding this line after the lines from step 1:

```
CometEngine engine = CometEngine.getEngine();
```

3 Register the component with the Comet engine by adding the following lines after those from step 2:

```
CometContext cometContext = engine.register(contextPath);
cometContext.setExpirationDelay(30 * 1000);
```

▼ Defining a Comet Handler to Send Updates to the Client**1 Create a private class that implements `CometHandler` and add it to the servlet class:**

```
private class CounterHandler
    implements CometHandler<HttpServletResponse> {
    private HttpServletResponse response;
}
```

2 Add the following methods to the class:

```
public void onInitialize(CometEvent event)
    throws IOException {}

    public void onInterrupt(CometEvent event)
        throws IOException {
        removeThisFromContext();
    }

    public void onTerminate(CometEvent event)
        throws IOException {
        removeThisFromContext();
    }

    public void attach(HttpServletResponse attachment) {
        this.response = attachment;
    }

    private void removeThisFromContext() throws IOException {
```

```

        response.getWriter().close();
        CometContext context =
            CometEngine.getEngine().getCometContext(contextPath);
        context.removeCometHandler(this);
    }

```

You need to provide implementations of these methods when implementing `CometHandler`. The `onInterrupt` and `onTerminate` methods execute when certain changes occur in the status of the underlying TCP communication. The `onInterrupt` method executes when communication is resumed. The `onTerminate` method executes when communication is closed. Both methods call `removeThisFromContext`, which removes the `CometHandler` object from the `CometContext` object.

▼ Adding the Comet Handler to the Comet Context

- 1 **Get an instance of the Comet handler and attach the response to it by adding the following lines to the `doGet` method:**

```

CounterHandler handler = new CounterHandler();
handler.attach(res);

```

- 2 **Get the Comet context by adding the following lines to `doGet`:**

```

CometEngine engine = CometEngine.getEngine();
CometContext context = engine.getCometContext(contextPath);

```

- 3 **Add the Comet handler to the Comet context by adding this line to `doGet`:**

```

context.addCometHandler(handler);

```

▼ Notifying the Comet Handler of an Event

- 1 **Add an `onEvent` method to the `CometHandler` class to define what happens when an event occurs:**

```

public void onEvent(CometEvent event)
    throws IOException {
    if (CometEvent.NOTIFY == event.getType()) {
        int count = counter.get();
        PrintWriter writer = response.getWriter();
        writer.write("<script type='text/javascript'>" +
            "parent.counter.updateCount('" + count + "') +
            "</script>\n");
        writer.flush();
        event.getCometContext().resumeCometHandler(this);
    }
}

```

This method first checks if the event type is `NOTIFY`, which means that the web component is notifying the `CometHandler` object that a client has incremented the count. If the event type is `NOTIFY`, the `onEvent` method gets the updated count, and writes out JavaScript to the client. The JavaScript includes a call to the `updateCount` function, which will update the count on the clients' pages.

The last line resumes the Comet request and removes it from the list of active `CometHandler` objects. By this line, you can tell that this application uses the long-polling technique. If you were to delete this line, the application would use the HTTP-Streaming technique.

- **For HTTP-Streaming:**

Add the same code as for long-polling, except do not include the following line:

```
event.getCometContext().resumeCometHandler(this);
```

You don't include this line because you do not want to resume the request. Instead, you want the connection to remain open.

2 Increment the counter and forward the response by adding the following lines to the `doPost` method:

```
counter.incrementAndGet();
CometEngine engine = CometEngine.getEngine();
CometContext<?> context =
    engine.getCometContext(contextPath);
context.notify(null);
req.getRequestDispatcher("count.html").forward(req, res);
```

When a user clicks the button, the `doPost` method is called. The `doPost` method increments the counter. It then obtains the current `CometContext` object and calls its `notify` method. By calling `context.notify`, the `doPost` method triggers the `onEvent` method you created in the previous step. After `onEvent` executes, `doPost` forwards the response to the clients.

Creating the Client Pages

Developing the HTML pages for the client involves performing these steps:

1. Create a welcome HTML page, called `index.html`, that contains: one hidden frame for connecting to the servlet through an HTTP GET; one `IFrame` that embeds the `count.html` page, which contains the updated content; and one `IFrame` that embeds the `button.html` page, which is used for posting updates using HTTP POST.
2. Create the `count.html` page that contains an HTML element that displays the current count and the JavaScript for updating the HTML element with the new count.
3. Create the `button.html` page that contains a button for the users to submit updates.

▼ Creating a Welcome HTML Page That Contains IFrames for Receiving and Sending Updates

1 Create an HTML page called `index.html`.

2 Add the following content to the page:

```
<html>
  <head>
    <title>Comet Example: Counter with Hidden Frame</title>
  </head>
  <body>
</body>
</html>
```

3 Add IFrames for connecting to the server and receiving and sending updates to `index.html` in between the body tags:

```
<frameset>
  <iframe name="hidden" src="hidden_comet"
    frameborder="0" height="0" width="100%"></iframe>
  <iframe name="counter" src="count.html"
    frameborder="0" height="100%" width="100%"></iframe>
  <iframe name="button" src="button.html" frameborder="0" height="30%" width="100%"></iframe>
</frameset>
```

The first frame, which is hidden, points to the servlet by referencing its context path. The second frame displays the content from `count.html`, which displays the current count. The second frame displays the content from `button.html`, which contains the submit button for incrementing the counter.

▼ Creating the HTML Page That Updates and Displays the Content

1 Create an HTML page called `count.html` and add the following content to it:

```
<html>
  <head>
  </head>
  <body>
    <center>
      <h3>Comet Example: Counter with Hidden Frame</h3>
      <p>
        <b id="count">&nbsp;</b>
      <p>
    </center>
  </body>
</html>
```

This page displays the current count.

- 2 Add JavaScript code that updates the count in the page . Add the following lines in between the head tags of count .html:**

```
<script type='text/javascript'>
    function updateCount(c) {
        document.getElementById('count').innerHTML = c;
        parent.hidden.location.href = "hidden_comet";
    };
</script>
```

The JavaScript takes the updated count it receives from the servlet and updates the count element in the page. The last line in the updateCount function invokes the servlet's doGet method again to reestablish the connection.

- **For HTTP-Streaming:**

Add the same code as for long-polling, except for the following line:

```
parent.hidden.location.href = "hidden_comet"
```

This line invokes the doGet method of CometServlet again, which would reestablish the connection. In the case of HTTP-Streaming, you want the connection to remain open. Therefore, you don't include this line of code.

▼ Creating the HTML Page That Allows Submitting Updates

- **Create an HTML page called button.html and add the following content to it:**

```
<html>
    <head>
    </head>
    <body>
        <center>
            <form method="post" action="hidden_comet">
                <input type="submit" value="Click">
            </form>
        </center>
    </body>
</html>
```

This page displays a form with a button that allows a user to update the count on the server. The servlet will then broadcast the updated count to all clients.

Creating the Deployment Descriptor

This section describes how to create a deployment descriptor to specify how your Comet-enabled web application should be deployed.

▼ Creating the Deployment Descriptor

- Create a file called `web.xml` and put the following contents in it:

```
<?xml version="1.0" encoding="UTF-8"?>
  <web-app version="2.5"
    xmlns="http://java.sun.com/xml/ns/javaee"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation=
      "http://java.sun.com/xml/ns/javaee
      http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd ">

    <servlet>
      <servlet-name>HiddenCometServlet</servlet-name>
      <servlet-class>
        com.sun.grizzly.samples.comet.HiddenCometServlet
      </servlet-class>
      <load-on-startup>0</load-on-startup>
    </servlet>
    <servlet-mapping>
      <servlet-name>HiddenCometServlet</servlet-name>
      <url-pattern>/hidden_comet</url-pattern>
    </servlet-mapping>
  </web-app>
```

This deployment descriptor contains a servlet declaration and mapping for `HiddenCometServlet`. The `load-on-startup` attribute must be set to 0 so that the Comet-enabled servlet will not load until the client makes a request to it.

Deploying and Running a Comet-Enabled Application

Before running a Comet-enabled application in the Enterprise Server, you need to enable Comet in the server. Then you can deploy the application just as you would any other web application.

When running the application, you need to connect to it from at least two different browsers to experience the effect of the servlet updating all clients in response to one client posting an update to the server.

▼ Enabling Comet in the Enterprise Server

Before running a Comet-enabled application, you need to enable Comet in your application server by adding a special property to the `http-listener` element of the `domain.xml` file.

The following steps tell you how to add this property.

- 1 Open `domain-dir/config/domain.xml` in a text editor.

- 2 **Add the following property in between the `http-listener` start and end tags:**

```
<property name="cometSupport" value="true"/>
```

- 3 **Save `domain.xml` and restart the server.**

▼ Deploying the Example

These instructions tell you how to deploy the Hidden Frame example.

- 1 **Download [grizzly-comet-hidden-1.7.3.1.war](#).**

- 2 **Run the following command to deploy the example:**

```
as-install/bin/asadmin deploy grizzly-comet-hidden-1.7.3.1.war
```

▼ Running the Example

These instructions tell you how to run the Hidden Frame example.

- 1 **Open two web browsers, preferably two different brands of web browser.**

- 2 **Enter the following URL in both browsers:**

```
http://localhost:8080/grizzly-comet-hidden/index.html
```

- 3 **When the first page loads in both browsers, click the button in one of the browsers and watch the count change in the other browser window.**

Bayeux Protocol

The Bayeux protocol, often referred to as Cometd, greatly simplifies the use of Comet. No server-side coding is needed for servers such as Enterprise Server that support the Bayeux protocol. Just enable Comet and the Bayeux protocol, then write and deploy the client as described in the following tasks:

- “Enabling Comet” on page 96
- “Configuring the `web.xml` File” on page 97
- “Writing, Deploying, and Running the Client” on page 98

▼ Enabling Comet

Before running a Comet-enabled application, you need to enable Comet in your application server by adding a special property to the `http-listener` element of the `domain.xml` file.

The following steps tell you how to add this property.

- 1 **Open `domain-dir/config/domain.xml` in a text editor.**

- 2 **Add the following property in between the `http-listener` start and end tags:**

```
<property name="cometSupport" value="true"/>
```

- 3 **Save `domain.xml` and restart the server.**

▼ **Configuring the `web.xml` File**

To enable the Bayeux protocol on the Enterprise Server, you must reference the `CometdServlet` in your web application's `web.xml` file. In addition, if your web application includes a servlet, set the `load-on-startup` value for your servlet to `0` (zero) so that it will not load until the client makes a request to it.

- 1 **Open the `web.xml` file for your web application in a text editor.**
- 2 **Add the following XML code to the `web.xml` file:**

```
<servlet>
  <servlet-name>Grizzly Cometd Servlet</servlet-name>
  <servlet-class>
    com.sun.grizzly.cometd.servlet.CometdServlet
  </servlet-class>
  <init-param>
    <description>
      expirationDelay is the long delay before a request is
      resumed. -1 means never.
    </description>
    <param-name>expirationDelay</param-name>
    <param-value>-1</param-value>
  </init-param>
  <load-on-startup>1</load-on-startup>
</servlet>
<servlet-mapping>
  <servlet-name>Grizzly Cometd Servlet</servlet-name>
  <url-pattern>/cometd/*</url-pattern>
</servlet-mapping>
```

Note that the `load-on-startup` value for the `CometdServlet` is `1`.

- 3 **If your web application includes a servlet, set the `load-on-startup` value to `0` for your servlet (not the `CometdServlet`) as follows:**

```
<servlet>
  ...
  <load-on-startup>0</load-on-startup>
</servlet>
```

- 4 **Save the `web.xml` file.**

▼ Writing, Deploying, and Running the Client

The examples in this task are taken from the example chat application posted and discussed at http://weblogs.java.net/blog/jfarcand/archive/2007/02/gcometd_introdu_1.html.

1 Add script tags to the HTML page. For example:

```
<script type="text/javascript" src="chat.js"></script>
```

2 In the script, call the needed libraries. For example:

```
dojo.require("dojo.io.cometd");
```

3 In the script, use publish and subscribe methods to send and receive messages. For example:

```
cometd.subscribe("/chat/demo", false, room, "_chat");  
cometd.publish("/chat/demo", { user: room._username, chat: text});
```

4 Deploy the web application as you would any other web application. For example:

```
asadmin deploy cometd-example.war
```

5 Run the application as you would any other web application.

The context root for the example chat application is /cometd and the HTML page is index.html. So the URL might look like this:

```
http://localhost:8080/cometd/index.html
```

See Also For more information about deployment in the Enterprise Server, see the *Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide*.

For more information about the Bayeux protocol, see [Bayeux Protocol](http://svn.xantus.org/shortbus/trunk/bayeux/bayeux.html) (<http://svn.xantus.org/shortbus/trunk/bayeux/bayeux.html>).

For more information about the Dojo toolkit, see <http://dojotoolkit.org/>.

For information about pushing data from an external component such as an EJB module, see the example at http://blogs.sun.com/swchan/entry/java_api_for_cometd. Using this Grizzly Java API for Cometd makes your web application non-portable. Running your application on a server that doesn't support Grizzly Comet will not work.

For information about RESTful (REpresentational State Transfer) web services and Comet, see [RESTful Web Services and Comet](http://developers.sun.com/appserver/reference/techart/cometslideshow.html) (<http://developers.sun.com/appserver/reference/techart/cometslideshow.html>).

Developing Grails Applications

This section shows you how to get started using Groovy on Grails on the Enterprise Server by covering the following topics:

- “Introduction to Groovy and Grails” on page 99
- “Installing Grails” on page 99
- “Creating a Simple Grails Application” on page 100
- “Deploying and Running a Grails Application” on page 101

Introduction to Groovy and Grails

Groovy is a dynamic, object-oriented language for the Java Virtual Machine, which builds on the strengths of Java but has additional features inspired by languages such as Python, Ruby, and Smalltalk. For more information about Groovy, see [Groovy](http://groovy.codehaus.org) (<http://groovy.codehaus.org>).

Grails is an open-source web application framework that leverages the Groovy language and complements Java web development. Grails is a stand-alone development environment that can hide all configuration details or allow integration of Java business logic. For more information about Grails, see [Grails](http://www.grails.org) (<http://www.grails.org>).

Installing Grails

To develop and deploy Grails applications on the Enterprise Server, first install the Grails plug-in module.

▼ Installing the Grails Plug-in Module

1 Install the Grails add-on component from the Update Tool.

For information about the Update Tool, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

2 Create a `GRAILS_HOME` environment variable that points to the Grails directory, `as-install/grails`.

3 Add the `as-install/grails/bin` directory to the `PATH` environment variable.

Example 7-1 Setting UNIX Environment Variables

On Solaris, Linux, and other operating systems related to UNIX, use the following commands for steps 2 and 3:

```
set GRAILS_HOME=~/glassfish/grails
export GRAILS_HOME
cd $GRAILS_HOME
set PATH=$GRAILS_HOME/bin:$PATH
export PATH
chmod a+x $GRAILS_HOME/bin/*
```

Example 7-2 Setting Windows Environment Variables

On the Windows operating system, use the following commands for steps 2 and 3:

```
set GRAILS_HOME=C:\GlassFish\grails
set PATH=%GRAILS_HOME%\bin;%PATH%
```

Creating a Simple Grails Application

To create the `helloworld` application, perform both these tasks:

- “Creating the `helloworld` Application” on page 100
- “Creating the `hello` Controller” on page 100

For more information on creating Grails applications, see the [Grails Quick Start](http://grails.org/Quick+Start) (<http://grails.org/Quick+Start>).

▼ Creating the `helloworld` Application

1 **Go to the `as-install/grails/samples` directory.**

2 **Run the `grails create-app helloworld` command.**

The `grails create-app` command creates an application framework that you can modify.

▼ Creating the `hello` Controller

1 **Go to the `as-install/grails/samples/helloworld` directory.**

2 **Run the `grails create-controller hello` command.**

The `grails create-controller` command creates a controller file that you can modify.

3 **Edit the generated `HelloController.groovy` file so it looks like this:**

```
class HelloController {

    def world = {
        render "Hello World!"
    }
}
```

```
    }  
    //def index = { }  
}
```

Deploying and Running a Grails Application

To deploy and run your application, perform one of these tasks:

- “Running a Grails Application Using `run-app`” on page 101
- “Running a Grails Application Using Standard Deployment” on page 101

▼ Running a Grails Application Using `run-app`

1 Go to the application directory.

For example, go to the `as-install/grails/samples/helloworld` directory.

2 Run the `grails run-app` command.

The `grails run-app` command starts the Enterprise Server in the background and runs the application in one step. You don't need to create a WAR file or deploy your application.

3 To test your application, point your browser to `http://host:port/app-dir-name`.

For example, point to `http://localhost:8080/helloworld`. You should see a screen that begins, “Welcome to Grails.” Selecting the HelloController link should change the display to, “Hello World!”

See Also For details about the `grails run-app` command, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

▼ Running a Grails Application Using Standard Deployment

1 Go to the application directory.

For example, go to the `as-install/grails/samples/helloworld` directory.

2 Create the WAR file in one of the following ways:

- **Run the `grails war` command.**

This command creates a large WAR file containing all the application's dependencies.

- **Run the `grails shared-war` command.**

This command creates a small WAR file, but requires referencing of the Grails library JAR at deployment.

In the `helloworld` application, this step creates the `helloworld-0.1.war` file.

3 Deploy the WAR file in one of the following ways:

- **In the Administration Console, open the Applications component, go to the Web Applications page, select the Deploy button, and type the path to the WAR file.**

The path to the `helloworld` WAR file is

```
as-install/grails/samples/helloworld/helloworld-0.1.war.
```

If you used the `grails shared-war` command, specify the

```
as-install/grails/lib/glassfish-grails.jar
```

 file in the Libraries field.

- **On the command line, use the `asadmin deploy` command and specify the WAR file. For example:**

```
asadmin deploy helloworld-0.1.war
```

If you used the `grails shared-war` command, specify the libraries using the `--libraries` option. For example:

```
asadmin deploy --libraries $GRAILS_HOME/lib/glassfish-grails.jar helloworld-0.1.war
```

4 To test your application, point your browser to `http://host:port/war-file-name`. Do not include the `.war` extension.

For example, point to `http://localhost:8080/helloworld-0.1`. You should see a screen that begins, “Welcome to Grails.” Selecting the `HelloController` link should change the display to, “Hello World!”

See Also For details about the Administration Console, see the online help.

For details about the `asadmin deploy` command, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

For details about the `grails war` and `grails shared-war` commands, see the [Grails Quick Start \(http://grails.org/Quick+Start\)](#).

For general information about deployment, see the [Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide](#).

Advanced Web Application Features

This section includes summaries of the following topics:

- “Internationalization Issues” on page 103
- “Virtual Servers” on page 104
- “Default Web Modules” on page 105
- “Class Loader Delegation” on page 106
- “Using the default-web.xml File” on page 106
- “Configuring Logging and Monitoring in the Web Container” on page 107
- “Header Management” on page 107
- “Configuring Valves and Catalina Listeners” on page 107
- “Alternate Document Roots” on page 108
- “Redirecting URLs” on page 110
- “Using a context.xml File” on page 110
- “Enabling WebDav” on page 111
- “Using mod_jk” on page 113
- “Using SSI” on page 114
- “Using CGI” on page 116
- “Using PHP” on page 117
- “Using Scala and Lift” on page 117

Internationalization Issues

This section covers internationalization as it applies to the following:

- “The Server's Default Locale” on page 103
- “Servlet Character Encoding” on page 103

The Server's Default Locale

To set the default locale of the entire Enterprise Server, which determines the locale of the Administration Console, the logs, and so on, use the Administration Console. Select the Enterprise Server component, the Advanced tab, and the Domain Attributes tab. Then type a value in the Locale field. For details, click the Help button in the Administration 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 v3 Prelude 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 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.

If you deploy a web application and don't specify any assigned virtual servers, the web application is assigned to all currently defined virtual servers. If you then create additional virtual servers and want to assign existing web applications to them, you must redeploy the web applications. For more information about deployment, see the [Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide](#).

For more information about virtual servers, see [“virtual-server” in *Sun GlassFish Enterprise Server v3 Prelude Administration Reference*](#).

▼ To Assign a Default Virtual Server

- 1 In the Administration 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 105](#).

See Also For details, click the Help button in the Administration 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 v3 Prelude Application Deployment Guide*.
- 2 In the Administration 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 105](#).

See Also For details, click the Help button in the Administration 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 104](#). 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/docroot*. 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 a web application class loader look in the local class loader before delegating to its parent. To make the web application class loader follow the delegation model in the Servlet specification, set `delegat`=`"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 `delegat`=`"true"`, which causes the web application class loader to delegate in the same manner as the other class loaders. Use `delegat`=`"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 v3 Prelude Application Deployment Guide*.

Note – For Prelude, the `delegat` value is ignored and assumed to be set to `true`.

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 Administration Console to edit the `default-web.xml` file. For details, click the Help button in the Administration 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 Administration Console, click the Help button in the Administration Console. Logging and Monitor tabs are accessible from the Application Server page.

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. A valve class must implement the `org.apache.catalina.Valve` interface from Tomcat or previous Enterprise Server releases, or the `org.glassfish.web.valve.GlassFishValve` interface from the current Enterprise Server release. A listener class for a virtual server must implement the `org.apache.catalina.ContainerListener` or `org.apache.catalina.LifecycleListener`

interface. A listener class for a web module must implement the `org.apache.catalina.ContainerListener`, `org.apache.catalina.LifecycleListener`, or `org.apache.catalina.InstanceListener` interface.

In the `sun-web.xml` file, valve and listener properties for a web module look like this:

```
<sun-web-app ...>
  ...
  <property name="valve_1" value="org.glassfish.extension.Valve"/>
  <property name="listener_1" value="org.glassfish.extension.MyLifecycleListener"/>
</sun-web-app>
```

In the `domain.xml` file, valve and listener properties for a virtual server look like this:

```
<virtual-server ...>
  ...
  <property name="valve_1" value="org.glassfish.extension.Valve"/>
  <property name="listener_1" value="org.glassfish.extension.MyLifecycleListener"/>
</virtual-server>
```

You can define these properties for a virtual server in one of the following ways, then restart the server:

- You can define properties using the `asadmin set` command. For example:

```
asadmin set server-config.http-service.virtual-server.MyVS.property.valve_1="org.glassfish.extension.Valve"
```

- You can define virtual server properties using the Administration 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 Administration 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 v3 Prelude Application Deployment Guide* and “[virtual-server](#)” in *Sun GlassFish Enterprise Server v3 Prelude 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:

```
/srv/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:

```
/srv/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 v3 Prelude 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"/>
```

Using a context.xml File

You can define a `context.xml` file for all web applications, for web applications assigned to a specific virtual server, or for a specific web application.

To define a global context .xml file, place the file in the *domain-dir*/config directory and name it context.xml.

Use the contextXmlDefault property to specify the name and the location, relative to *domain-dir*, of the context.xml file for a specific virtual server. Specify this property in one of the following ways:

- In the Administration Console, open the HTTP Service component under the relevant configuration. Open the Virtual Servers component and scroll down to the bottom of the page. Enter contextXmlDefault as the property name and the path and file name relative to *domain-dir* as the property value.
- Use the `asadmin create-virtual-server` command. For example:

```
asadmin create-virtual-server --hosts localhost --property contextXmlDefault=config/vs1context.xml vs1
```

- Use the `asadmin set` command for an existing virtual server. For example:

```
asadmin set server-config.http-service.virtual-server.server.property.contextXmlDefault=config/mycontext.xml
```

To define a context.xml file for a specific web application, place the file in the META-INF directory and name it context.xml.

For more information about virtual servers, see “Virtual Servers” on page 104. 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>). Context parameters, environment entries, and resource definitions in context.xml are supported in the Enterprise Server.

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>
</servlet>
```

```
<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
```


Enable the security manager as described in “[Enabling and Disabling the Security Manager](#)” on [page 51](#).

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 and install Apache 2.0.x or 2.2.x HTTP Server (<http://httpd.apache.org/>).
2. 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.

3. Start httpd.
4. Enable mod_jk using the following command:

```
asadmin set server-config.http-service.http-listener.listener.property.jkEnabled=true
```

For example:

```
asadmin set server-config.http-service.http-listener.http-listener1.property.jkEnabled=true
```

Or you can use the following deprecated command, provided for backward compatibility:

```
asadmin create-jvm-options -Dcom.sun.enterprise.web.connector.enableJK=8009
```

5. 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
```

6. 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 SSI

To enable SSI (server-side includes) processing for a specific web module, add the `SSIServlet` to your `web.xml` file as follows:

```
<web-app>
  <servlet>
    <servlet-name>ssi</servlet-name>
    <servlet-class>org.apache.catalina.ssi.SSIServlet</servlet-class>
```

```

</servlet>
...
<servlet-mapping>
  <servlet-name>ssi</servlet-name>
  <url-pattern>*.shtml</url-pattern>
</servlet-mapping>
...
<mime-mapping>
  <extension>shtml</extension>
  <mime-type>text/html</mime-type>
</mime-mapping>
</web-app>

```

To enable SSI processing for all web modules, un-comment the corresponding sections in the `default-web.xml` file.

If the `mime-mapping` is not specified in `web.xml`, Enterprise Server attempts to determine the MIME type from `default-web.xml` or the operating system default.

You can configure the following `init-param` values for the `SSIServlet`.

TABLE 7-2 SSIServlet `init-param` Values

<code>init-param</code>	Type	Default	Description
<code>buffered</code>	boolean	false	Specifies whether the output should be buffered.
<code>debug</code>	int	0 (for no debugging)	Specifies the debugging level.
<code>expires</code>	Long	Expires header in HTTP response not set	Specifies the expiration time in seconds.
<code>inputEncoding</code>	String	operating system encoding	Specifies encoding for the SSI input if there is no URL content encoding specified.
<code>isVirtualWebappRelative</code>	boolean	false (relative to the given SSI file)	Specifies whether the virtual path of the <code>#include</code> directive is relative to the content-root.
<code>outputEncoding</code>	String	UTF-8	Specifies encoding for the SSI output.

For more information about SSI, see

http://httpd.apache.org/docs/2.2/mod/mod_include.html.

Using CGI

To enable CGI (common gateway interface) processing for a specific web module, add the `CGIServlet` to your `web.xml` file as follows:

```
<web-app>
  <servlet>
    <servlet-name>cgi</servlet-name>
    <servlet-class>org.apache.catalina.servlets.CGIServlet</servlet-class>
  </servlet>
  ...
  <servlet-mapping>
    <servlet-name>cgi</servlet-name>
    <url-pattern>/cgi-bin/*</url-pattern>
  </servlet-mapping>
</web-app>
```

To enable CGI processing for all web modules, un-comment the corresponding sections in the `default-web.xml` file.

Package the CGI program under the `cgiPathPrefix`. The default `cgiPathPrefix` is `WEB-INF/cgi`. For security, it is highly recommended that the contents and binaries of CGI programs be prohibited from direct viewing or download. For information about hiding directory listings, see [“Using the default-web.xml File” on page 106](#).

Invoke the CGI program using a URL of the following format:

```
http://host:8080/context-root/cgi-bin/cgi-name
```

For example:

```
http://localhost:8080/mycontext/cgi-bin/hello
```

You can configure the following `init-param` values for the `CGIServlet`.

TABLE 7-3 `CGIServlet` `init-param` Values

<code>init-param</code>	Type	Default	Description
<code>cgiPathPrefix</code>	String	<code>WEB-INF/cgi</code>	Specifies the subdirectory containing the CGI programs.
<code>debug</code>	int	0 (for no debugging)	Specifies the debugging level.
<code>executable</code>	String	<code>perl</code>	Specifies the executable for running the CGI script.

TABLE 7-3 CGIServlet init-param Values (Continued)

init-param	Type	Default	Description
parameterEncoding	String	System.getProperty("file.encoding", "UTF-8")	Specifies the parameter's encoding.
passShellEnvironment	boolean	false	Specifies whether to pass shell environment properties to the CGI program.

To work with a native executable, do the following:

1. Set the value of the `init-param` named `executable` to an empty `String` in the `web.xml` file.
2. Make sure the executable has its executable bits set correctly.
3. Use directory deployment to deploy the web module. Do not deploy it as a WAR file, because the executable bit information is lost during the process of `jar` and `unjar`. For more information about directory deployment, see the *Sun GlassFish Enterprise Server v3 Prelude Application Deployment Guide*.

Using PHP

To enable PHP, deploy the Quercus PHP interpreter to the Enterprise Server as a web module as follows:

1. Download the Quercus interpreter from <http://quercus.caucho.com/>.
2. Deploy the WAR file you downloaded to the Enterprise Server.
3. To verify that your PHP engine is working, point your browser to the default PHP script that comes with the Quercus interpreter, which is `http://localhost:8080/quercus-3.1.6/`.
4. The Quercus application directory is located at `domain-dir/applications/quercus-3.1.6/`. Place your PHP application in a subdirectory of the Quercus directory, for example `domain-dir/applications/quercus-3.1.6/myapp/`. To point your browser to the PHP application, enter `http://localhost:8080/quercus-3.1.6/myapp/`.

For more information about using the Quercus PHP interpreter, see the documentation at <http://quercus.caucho.com/quercus-3.1/doc/quercus.xtp>.

Using Scala and Lift

Scala is a general purpose programming language designed to express common programming patterns in a concise, elegant, and type-safe way. It smoothly integrates features of object-oriented and functional languages. It is also fully interoperable with Java. For details, see <http://www.scala-lang.org/>.

Lift is an expressive and elegant framework for writing web applications using Scala. Lift stresses the importance of security, maintainability, scalability and performance, while allowing for high levels of developer productivity. For details, see <http://liftweb.net/>.

It is common practice to start a Lift web application using Maven. Maven is a software project management and comprehension tool. Based on the concept of a project object model (POM), Maven can manage a project's build, reporting and documentation from a central piece of information. For details, see <http://maven.apache.org/>.

To create a new Lift project, use Maven interactively in one of these ways:

```
mvn archetype:generate -DarchetypeCatalog=http://scala-tools.org/
```

Or:

```
mvn org.apache.maven.plugins:maven-archetype-plugin:1.0-alpha-7:create \
-DarchetypeGroupId=net.liftweb \
-DarchetypeArtifactId=lift-archetype-blank \
-DarchetypeVersion=0.7.1 \
-DremoteRepositories=http://scala-tools.org/repo-releases \
-DgroupId=__my.liftapp__ -DartifactId=__liftapp__
```

Or:

```
mvn org.apache.maven.plugins:maven-archetype-plugin:1.0-alpha-7:create \
-DarchetypeGroupId=net.liftweb \
-DarchetypeArtifactId=lift-archetype-basic \
-DarchetypeVersion=0.7.1 \
-DremoteRepositories=http://scala-tools.org/repo-releases \
-DgroupId=__my.liftapp__ -DartifactId=__liftapp__
```

After coding your application, build the WAR file using the `mvn package` command. Then deploy the WAR file to the Enterprise Server as you would any other web application.

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 3.1 Changes” on page 119
- “Value Added Features” on page 120
- “EJB Timer Service” on page 121
- “Using Session Beans” on page 122
- “Handling Transactions With Enterprise Beans” on page 123

For general information about enterprise beans, see “Part Three: Enterprise Beans” in the [Java EE 5 Tutorial \(http://java.sun.com/javase/5/docs/tutorial/doc/index.html\)](http://java.sun.com/javase/5/docs/tutorial/doc/index.html).

Note – For GlassFish v3 Prelude, EJB modules are not supported unless the optional EJB container add-on component is downloaded from the Update Tool. For information about the Update Tool, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

For GlassFish v3 Prelude, only stateless session beans with local interfaces and entity beans that use the Java Persistence API are supported. Stateful, message-driven, and EJB 2.0 and 2.1 entity beans are not supported. Remote interfaces and remote business interfaces for any of the bean types are not supported.

Summary of EJB 3.1 Changes

The Enterprise Server supports and is compliant with the Sun Microsystems Enterprise JavaBeans (EJB) architecture as defined by the Enterprise JavaBeans Specification, v3.1, also known as JSR 318 (<http://jcp.org/en/jsr/detail?id=318>).

The main changes in the Enterprise JavaBeans Specification, v3.1 that impact enterprise beans in the Enterprise Server environment are as follows:

- An EJB component need not implement any interface as long as it contains one of the component defining annotations or the XML equivalent. Essentially, the local business interface is optional. For example, the following is a simple no-interface bean:

```
@Stateless
public class HelloBean {
    public String sayHello(String msg) {
        return "Hello " + msg;
    }
}
```

Even though the bean doesn't implement any interface, the client can still inject (or look up) a reference to the session bean. The client still has to perform a JNDI lookup or inject a reference of the bean. More specifically, it *cannot* use the new operator to construct the bean.

```
@EJB HelloBean h;
...
h.sayHello("bob");
```

- EJB classes can be packaged inside WAR files. These classes must reside under WEB-INF/classes. For example, the structure of a hello.war file might look like this:

```
index.jsp
META-INF/
  MANIFEST.MF
WEB-INF/
  web.xml
  classes/
    com/
      sun/
        v3/
          demo/
            HelloEJB.class
            HelloServlet.class
```

For more information about web applications, see [Chapter 7, “Developing Web Applications.”](#)

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.

- [“Bean-Level Container-Managed Transaction Timeouts” on page 121](#)

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.

EJB Timer Service

The EJB Timer Service uses a database to store persistent information about EJB timers.

The EJB Timer Service in Enterprise Server is preconfigured to use an embedded version of the Java DB database.

Note – If the optional JTS add-on component has not been downloaded from the Update Tool, you must reconfigure the JDBC resource named `jdbc/__TimerPool`. If the optional JTS add-on component has been downloaded, this is not necessary. For information about the Update Tool, see the [Sun GlassFish Enterprise Server v3 Prelude Installation Guide](#).

In the Administration Console, open the Resources component and select JDBC Resources. For details, click the Help button in the Administration Console. Change the connection pool name for the JDBC resource named `jdbc/__TimerPool` to point to the same connection pool as the one you are using for the rest of your data. Then start the database.

To enable the timer service, deploy the following application:

```
as-install/lib/install/applications/ejb-timer-service-app.war
```

You can verify that it was deployed successfully by accessing the following URL:

```
http://localhost:8080/ejb-timer-service-app/timer
```

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 v3 Prelude Release Notes](#). For configurations of supported and other drivers, see “Configuration Specifics for JDBC Drivers” in [Sun GlassFish Enterprise Server v3 Prelude 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*.

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.

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 122](#)
- [“Session Bean Restrictions and Optimizations” on page 123](#)

Information on session beans is contained in the Enterprise JavaBeans Specification, v3.1.

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.

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 v3 Prelude Application Deployment Guide*.

The Enterprise Server provides the `wscpile` and `wsdeploy` 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 v3 Prelude Reference Manual*.

Session Bean Restrictions and Optimizations

This section discusses restrictions on developing session beans and provides some optimization guidelines.

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.

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 124
- “[Global and Local Transactions](#)” on page 124
- “[Administration and Monitoring](#)” on page 124

Note – For GlassFish v3 Prelude, global (XA) transactions are not supported unless the optional JTS add-on component is downloaded from the Update Tool. Without the JTS add-on component, only local transactions are supported. For information about the Update Tool, see the [Sun GlassFish Enterprise Server v3 Prelude Installation Guide](#).

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 138](#).

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 139](#).

Administration and Monitoring

An administrator can control a number of domain-level Transaction Service settings. For details, see [“Configuring the Transaction Service” on page 139](#).

The Transaction Timeout setting can be overridden by a bean. See [“Bean-Level Container-Managed Transaction Timeouts” on page 121](#).

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.

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.

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 127
- “Creating Web Applications That Use the JDBC API” on page 129
- “Restrictions and Optimizations” on page 135

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 128
- “Creating a Connection Pool” on page 128
- “Testing a JDBC Connection Pool” on page 129
- “Creating a JDBC Resource” on page 129

For information about how to configure some specific JDBC drivers, see “[Configuration Specifics for JDBC Drivers](#)” in *Sun GlassFish Enterprise Server v3 Prelude 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 128
- “[Making the JDBC Driver JAR Files Accessible](#)” on page 128

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 v3 Prelude Release Notes*. For configurations of supported and other drivers, see “[Configuration Specifics for JDBC Drivers](#)” in *Sun GlassFish Enterprise Server v3 Prelude 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 an 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 Administration Console, open the Resources component and select Connection Pools. For details, click the Help button in the Administration Console.
- Use the `asadmin create-jdbc-connection-pool` command. For details, see the *Sun GlassFish Enterprise Server v3 Prelude Reference Manual*.

For a complete description of JDBC connection pool features, see the *Sun GlassFish Enterprise Server v3 Prelude Administration Guide*

Testing a JDBC Connection Pool

You can test a JDBC connection pool for usability in one of these ways:

- In the Administration Console, open the Resources 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 Administration Console.
- Use the `asadmin ping-connection-pool` command. For details, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

Both these commands fail and display an error message unless they successfully connect to the connection pool.

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 Administration Console, open the Resources component and select JDBC Resources. For details, click the Help button in the Administration Console.
- Use the `asadmin create-jdbc-resource` command. For details, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

Creating Web Applications That Use the JDBC API

A web application that uses the JDBC API is an application that looks up and connects to one or more databases. This section covers these topics:

- “Setting a Statement Timeout” on page 129
- “Sharing Connections” on page 130
- “Wrapping Connections” on page 130
- “Obtaining a Physical Connection From a Wrapped Connection” on page 131
- “Using the `Connection.unwrap()` Method” on page 131
- “Marking Bad Connections” on page 131
- “Using Non-Transactional Connections” on page 132
- “Using JDBC Transaction Isolation Levels” on page 133
- “Allowing Non-Component Callers” on page 134

Setting a Statement Timeout

An abnormally long running JDBC query executed by an application may leave it in a hanging state unless a timeout is explicitly set on the statement. Setting a statement timeout guarantees that all queries automatically time out if not completed within the specified period. When

statements are created, the `queryTimeout` is set according to the statement timeout setting. This works only when the underlying JDBC driver supports `queryTimeout` for `Statement`, `PreparedStatement`, `CallableStatement`, and `ResultSet`.

You can specify a statement timeout in the following ways:

- Enter a `Statement Timeout` value in the `JDBC Connection Pools` page in the Administration Console. For more information, click the `Help` button in the Administration Console.
- Specify the `--statementtimeout` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

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 5, "Administering Database Connectivity,"](#) in *Sun GlassFish Enterprise Server v3 Prelude Administration Guide*.

Wrapping Connections

If the `Wrap JDBC Objects` option is `true`, wrapped JDBC objects are returned for `Statement`, `PreparedStatement`, `CallableStatement`, `ResultSet`, and `DatabaseMetaData`. The default is `false`.

This option ensures that `Statement.getConnection()` is the same as `DataSource.getConnection()`. Therefore, this option should be `true` when both `Statement.getConnection()` and `DataSource.getConnection()` are done. The default is `false` to avoid breaking existing applications.

You can specify the `Wrap JDBC Objects` option in the following ways:

- Check or uncheck the `Wrap JDBC Objects` box on the `JDBC Connection Pools` page in the Administration Console. For more information, click the `Help` button in the Administration Console.

- Specify the `--wrapjdbcobjects` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the *Sun GlassFish Enterprise Server v3 Prelude Reference Manual*.

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); //get physical connection from wrapper
// Do db operations.
// Do not close driver connection.
con.close(); // return wrapped connection to pool.
```

Using the `Connection.unwrap()` Method

If the JDK version 1.6 is used, the Enterprise Server supports JDBC 4.0 if the JDBC driver is JDBC 4.0 compliant. Using the `Connection.unwrap()` method on a vendor-provided interface returns an object or a wrapper object implementing the vendor-provided interface, which the application can make use of to do vendor-specific database operations. Use the `Connection.isWrapperFor()` method on a vendor-provided interface to check whether the connection can provide an implementation of the vendor-provided interface. Check the JDBC driver vendor's documentation for information on these interfaces.

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){
    ds.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 Administration Console. The default is unchecked. For more information, click the Help button in the Administration Console.
- Specify the `--nonttransactionalconnections` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the *Sun GlassFish Enterprise Server v3 Prelude 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 10, “Using the Transaction Service.”](#) For information about last agent optimization, which can improve performance, see [“Transaction Scope” on page 138.](#)

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 9-1 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.

You can specify the transaction isolation level in the following ways:

- Select the value from the Transaction Isolation drop-down list on the JDBC Connection Pools page in the Administration Console. For more information, click the Help button in the Administration Console.
- Specify the `--isolationlevel` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

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 128.](#)

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 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 Administration Console. The default is `false`. For more information, click the Help button in the Administration Console.
- Specify the `--allownoncomponentcallers` option in the `asadmin create-jdbc-connection-pool` command. For more information, see the [Sun GlassFish Enterprise Server v3 Prelude 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` for the JDBC connection pool.

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 Scope” on page 138
- “Configuring the Transaction Service” on page 139
- “The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction” on page 139

For more information about the Java™ Transaction API (JTA) and Java Transaction Service (JTS), see 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) (<http://java.sun.com/javaee/5/docs/tutorial/doc/index.html>).

For information about JDBC transaction isolation levels, see “[Using JDBC Transaction Isolation Levels](#)” on page 133.

Note – For GlassFish v3 Prelude, global (XA) transactions are not supported unless the optional JTS add-on component, including the dependent Object Management Group (OMG) subcomponent, is downloaded from the Update Tool. Without this add-on component, only local transactions are supported. For information about the Update Tool, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

Transaction recovery is not implemented for GlassFish v3 Prelude, even if the JTS and OMG add-on components are installed. Therefore, all transaction service attributes and properties pertaining to transaction recovery or transaction logs are not implemented.

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 are *global* transactions. A global transaction can involve one non-XA resource if last agent optimization is enabled. Otherwise, all resourced 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 139](#).

Note – For GlassFish v3 Prelude, transaction propagation between multiple participant processes, applications, or JVM machines is not supported.

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.

Configuring the Transaction Service

You can configure the transaction service in the Enterprise Server in the following ways:

- To configure the transaction service, use the `asadmin set` command to set the following attributes and properties.

```
server-config.transaction-service.automatic-recovery = false
server-config.transaction-service.heuristic-decision = rollback
server-config.transaction-service.keypoint-interval = 2048
server-config.transaction-service.retry-timeout-in-seconds = 600
server-config.transaction-service.timeout-in-seconds = 0
server-config.transaction-service.tx-log-dir = domain-dir/logs
server-config.transaction-service.property.use-last-agent-optimization = true
```

You can use the `asadmin get` command to list all the transaction service attributes and properties. For details, see the [Sun GlassFish Enterprise Server v3 Prelude Reference Manual](#).

Changing `keypoint-interval`, `retry-timeout-in-seconds`, or `timeout-in-seconds` does not require a server restart. Changing other attributes or properties requires a server restart.

Transaction recovery is not implemented for GlassFish v3 Prelude, even if the JTS and OMG add-on components are installed. Therefore, all transaction service attributes and properties pertaining to transaction recovery or transaction logs are not implemented.

The Transaction Manager, the Transaction Synchronization Registry, and UserTransaction

To access a `UserTransaction` instance, you can either look it up using the `java:comp/UserTransaction` JNDI name or inject it using the `@Resource` annotation.

If you need to access the `javax.transaction.TransactionManager` implementation, you can look up the Enterprise Server implementation of this interface using the JNDI name `java:appserver/TransactionManager`. If possible, you should use the `javax.transaction.TransactionSynchronizationRegistry` interface instead, for portability. You can look up the implementation of this interface by using the JNDI name `java:comp/TransactionSynchronizationRegistry`. For details, see the Javadoc page for [Interface TransactionSynchronizationRegistry](http://java.sun.com/javaee/5/docs/api/javax/transaction/TransactionSynchronizationRegistry.html) (<http://java.sun.com/javaee/5/docs/api/javax/transaction/TransactionSynchronizationRegistry.html>) and [Java Specification Request \(JSR\) 907](http://www.jcp.org/en/jsr/detail?id=907) (<http://www.jcp.org/en/jsr/detail?id=907>).

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 141
- “Mapping References” on page 143

Note – For GlassFish v3 Prelude, EJB modules are not supported unless the optional EJB container add-on component is downloaded from the Update Tool. For information about the Update Tool, see the *Sun GlassFish Enterprise Server v3 Prelude Installation Guide*.

For GlassFish v3 Prelude, only stateless session beans with local interfaces and entity beans that use the Java Persistence API are supported. Stateful, message-driven, and EJB 2.0 and 2.1 entity beans are not supported. Remote interfaces and remote business interfaces for any of the bean types are not supported.

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.

Note – Each resource within the server must have a unique 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` deployment descriptor file has the highest precedence. See [“Mapping References” on page 143](#).
2. A global JNDI name assigned in a `mapped-name` element in the `ejb-jar.xml`, `web.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`, 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`.
4. A default global JNDI name is assigned in some cases if no name is assigned in deployment descriptors or annotations.
 - For 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`.

Using a Custom `jndi.properties` File

To use a custom `jndi.properties` file, place the file in the `domain-dir/lib/classes` directory or JAR it and place it in the `domain-dir/lib` directory. This adds the custom `jndi.properties` file to the Common 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 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` and `sun-ejb-ref.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 v3 Prelude Application Deployment Guide*](#).

The rest of this section uses an example of a JDBC resource lookup to describe how to reference resource factories.

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>
```


Index

Numbers and Symbols

@OrderBy and session cache sharing, 68

A

Admin Console, 21

 Audit Modules page, 47

 Debug Enabled field, 30

 Default Virtual Server field, 105

 HPROF configuration, 32

 JACC Providers page, 47

 JDBC Connection Pools page, 128

 Allow Non Component Callers field, 134

 Non-Transactional Connections field, 132

 Ping button, 129

 Statement Timeout field, 130

 Transaction Isolation field, 133

 Wrap JDBC Objects field, 130

 JDBC Resources page, 129

 JProbe configuration, 33

 Libraries field, 26

 Locale field, 103

 Logging tab, 107

 Monitor tab, 107

 online help for, 21

 Realms page, 44

 role mapping configuration, 43

 Security Manager Enabled field, 52

 Virtual Servers page, 105

 Web Services page

 Publish tab, 59

Admin Console, Web Services page (*Continued*)

 Registry tab, 58

 Test button, 60

 Write to System Log field, 75

alternate document roots, 108-110

annotation

 JNDI names, 142

 schema generation, 66

 security, 40

Applib class loader, 24

AppservPasswordLoginModule class, 45

AppservRealm class, 46

Archive class loader, 24

asadmin command, 20-21

 create-audit-module, 47

 create-auth-realm, 44

 create-jdbc-connection-pool, 128

 --allownoncomponentcallers option, 134

 --isolationlevel option, 133

 --nontransactionalconnections option, 132

 --statementtimeout option, 130

 --wrapjdbcobjects option, 131

 create-jdbc-resource, 129

 create-jvm-options

 java.security.debug option, 51

 delete-jvm-options

 java.security.manager option, 52

 deploy

 --libraries option, 26

 --precompilejsp option, 80

 generate-jvm-report, 31

 get, 139

asadmin command (*Continued*)

- ping-connection-pool, 129
 - publish-to-registry, 59
 - set
 - default principal settings, 43
 - transaction service settings, 139
- audit modules, 47-48
- AuditModule class, 47-48
- authentication
- audit modules, 48
 - JAAS, 45-46
 - programmatically login, 52
 - realms, 44
 - single sign-on, 55-56
- authorization
- audit modules, 48
 - JAAS, 45-46
 - JACC, 47
 - roles, 42-43
- automatic schema generation, Java Persistence
- options, 66

B

- Bayeux protocol, 96-98
- Bootstrap class loader, 24

C

- cache for servlets
- default configuration, 77
 - example configuration, 77
 - helper class, 77, 79
- cache sharing and @OrderBy, 68
- CacheHelper interface, 79
- cacheKeyGeneratorAttrName property, 79
- caching
- data using a non-transactional connection, 133
 - servlet results, 76-79
- Catalina listeners, defining custom, 107-108
- certificate realm, 44
- CGI, 116-117
- class-loader element, 106

- class loaders, 23-27
- application-specific, 26
 - circumventing isolation, 27
 - delegation hierarchy, 24
 - isolation, 25-26
- Comet, 84-98
- Cometd, 96-98
- command-line server configuration, *See* asadmin command
- Common class loader, 24
- using to circumvent isolation, 27
- common gateway interface, 116-117
- compiling JSP files, 80-81
- Connector class loader, 24
- context root, 75
- context.xml file, 110-111
- context, for JNDI naming, 141-142
- create-audit-module command, 47
- create-auth-realm command, 44
- create-jdbc-connection-pool command, 128
- allownoncomponentcallers option, 134
 - isolationlevel option, 133
 - nontransactionalconnections option, 132
 - statementtimeout option, 130
 - wrapjdbcobjects option, 131
- create-jdbc-resource command, 129
- create-jvm-options command, java.security.debug option, 51

D

- database properties, 64
- databases
- properties, 64
 - specifying for Java Persistence, 62-63
 - supported, 128
- debugging, 29-35
- enabling, 29-30
 - generating a stack trace, 31
 - JPDA options, 30
- DeclareRoles annotation, 42-43
- default virtual server, 105
- default web module, 75, 105-106
- default-web.xml file, 106-107

delegation, class loader, 25
 delete-jvm-options command, java.security.manager
 option, 52
 deploy command
 --libraries option, 26
 --precompilejsp option, 80
 deployment descriptor files, 143
 destroy method, 79
 development environment
 creating, 19-22
 tools for developers, 20-22
 digest authentication, 44
 directory listings, disabling, 106
 document root, 104, 106
 document roots, alternate, 108-110
 doGet method, 79, 80
 domain.xml file, configuring single sign-on, 56
 doPost method, 79, 80

E

Eclipse IDE, 21
 EclipseLink, 61
 eclipselink.target-database property, 62
 EJB 3.0, Java Persistence, 61-72
 EJB 3.1, summary of changes, 119
 EJB components
 pooling, 122
 security, 41
 ejb-ref element, 143
 EJB Timer Service, 121-122
 encoding, of servlets, 103-104
 endorsed standards override mechanism, 25
 Extension class loader, 24

F

file realm, 44
 finder limitation for Sybase, 70
 flat transactions, 124

G

generate-jvm-report command, 31
 get command, 139
 getCharacterEncoding method, 104
 getConnection method, 131
 getHeaders method, 107
 GlassFish project, 20
 Grails, 99-102
 Grizzly, Comet, 86-96
 Groovy, 99-102

H

handling requests, 79
 header management, 107
 help for Admin Console tasks, 21
 HPROF profiler, 32-33
 HTTP sessions, 81-84
 and redeployment, 81-82
 cookies, 81
 session managers, 82-84
 URL rewriting, 81
 HttpServletRequest, 77

I

Inet Oracle JDBC driver, 69
 init method, 79
 InitialContext naming service handle, 141-142
 installation, 19-20
 instantiating servlets, 79
 internationalization, 103
 isolation of class loaders, 25-26, 27

J

JACC, 47
 Java Authentication and Authorization Service
 (JAAS), 45-46
 Java Authorization Contract for Containers, *See* JACC
 Java Database Connectivity, *See* JDBC
 Java DB database, 62-63

Java Debugger (jdb), 29
Java EE tutorial, 73
Java EE, security model, 40
Java Naming and Directory Interface, *See* JNDI
Java optional package mechanism, 25
Java Persistence, 61-72
 annotation for schema generation, 66
 changing the provider, 67-68
 database for, 62-63
 deployment options for schema generation, 66
 restrictions, 68-72
Java Platform Debugger Architecture, *See* JPDA
Java Servlet API, 74
Java Transaction API (JTA), 137-139
Java Transaction Service (JTS), 137-139
JavaBeans, 80
jdbc realm, 44
JDBC
 connection pool creation, 128
 Connection wrapper, 131
 creating resources, 129
 integrating driver JAR files, 27, 128
 non-component callers, 134
 non-transactional connections, 132-133
 restrictions, 135
 sharing connections, 130
 specification, 127
 supported drivers, 128
 transaction isolation levels, 133
 tutorial, 127
JNDI
 and EJB components, 143
 defined, 141-143
 global names, 142
 mapping references, 143
 tutorial, 141
JPDA debugging options, 30
JProbe profiler, 33-35
JSP files
 command-line compiler, 80-81
 precompiling, 80-81
 specification, 80
 tag libraries, 80
jspc command, 81

JSR 109, 57
JSR 115, 40, 47, 48
JSR 181, 58
JSR 196, 40
JSR 220, 61
JSR 224, 57
JSR 318, 119
JSR 907, 139

L

last agent optimization, 138
ldap realm, 44
lib directory, and the Common class loader, 24
libraries, 26, 27
Lift, 117-118
listeners, Catalina, defining custom, 107-108
locale, setting default, 103
logging, in the web container, 107
login method, 54
LoginModule, 45
login, programmatic, 52

M

mapping resource references, 143
markConnectionAsBad method, 131-132
Maven, 117-118
Migration Tool, 21
mime-mapping element, 107
monitoring in the web container, 107
MySQL database restrictions, 70-72

N

naming service, 141-143
native library path
 configuring for hprof, 32
 configuring for JProbe, 34
nested transactions, 124
NetBeans
 about, 21

NetBeans (*Continued*)

 profiler, 32

O

online help, 21

Oracle Inet JDBC driver, 69

Oracle TopLink, 68

output from servlets, 75

P

permissions

 changing in server.policy, 49-51

 default in server.policy, 49

persistence.xml file, 62-63, 66

PHP, 117

ping-connection-pool command, 129

precompiling JSP files, 80-81

profilers, 31-35

programmatic login, 52

ProgrammaticLogin class, 54

ProgrammaticLoginPermission permission, 53

Public API class loader, 24

publish-to-registry command, 59

Q

query hints, 67

R

realms

 application-specific, 44

 configuring, 44

 custom, 45-46

 supported, 44

redirecting a URL, 110

removing servlets, 79

request object, 79

res-sharing-scope deployment descriptor setting, 130

resource-env-ref element, 143

resource-ref element, 143

resource references, mapping, 143

roles, 42-43

S

Scala, 117-118

schema generation, Java Persistence options for
 automatic, 66

security, 39-56

security manager, enabling and disabling, 51-52

security

 annotations, 40

 application level, 41

 audit modules, 47-48

 declarative, 41

 disabling directory listings, 106

 EJB components, 41

 Enterprise Server features, 40

 goals, 40

 JACC, 47

 Java EE model, 40

 of containers, 40-42

 programmatic, 42

 programmatic login, 52

 roles, 42-43

 server.policy file, 49-52

 web applications, 41

server.policy file, 49-52

 changing permissions, 49-51

 default permissions, 49

 ProgrammaticLoginPermission, 53

server-side includes, 114-115

server

 installation, 19-20

 lib directory of, 24

 optimizing for development, 20

 value-added features, 120-121

service method, 79, 80

ServletContext.log messages, 75

servlets, 74-80

 caching, 76-79

 character encoding, 103-104

servlets (*Continued*)

- destroying, 79
- engine, 79
- instantiating, 79
- invoking using a URL, 74-75
- output, 75
- removing, 79
- request handling, 79
- specification, 74
 - class loading, 106
 - mime-mapping, 107
- session beans, 122
 - container for, 122-123
 - restrictions, 123
- session cache sharing and `@OrderBy`, 68
- session managers, 82-84
- set command
 - default principal settings, 43
 - transaction service settings, 139
- `setCharacterEncoding` method, 104
- `setContentType` method, 104
- `setLocale` method, 104
- `setTransactionIsolation` method, 133
- single sign-on, 55-56
- Sitraka web site, 33-35
- specification
 - EJB 3.1, 119
 - JAAS, 45
 - Java Persistence, 61
 - JavaBeans, 80
 - JDBC, 127
 - JSP, 80
 - programmatic security, 42
 - security manager, 49
 - servlet, 74
- SSI, 114-115
- stack trace, generating, 31
- stateless session beans, 122-123
- Sun Java Studio, 21
- `sun-web.xml` file, and class loaders, 106
- `supportsTransactionIsolationLevel` method, 134
- Sybase, finder limitation, 70

T

- tag libraries, 80
- tools, for developers, 20-22
- transactions, 137-139
 - administration and monitoring, 124
 - and EJB components, 123-124
 - configuring, 139
 - flat, 124
 - global, 124
 - in the Java EE tutorial, 137-139
 - JDBC isolation levels, 133
 - local, 124
 - local or global scope of, 138
 - nested, 124
 - timeouts, 121
 - transaction manager, 139
 - transaction synchronization registry, 139
 - UserTransaction, 139

U

- unwrap method, 131
- URL rewriting, 81
- URL, redirecting, 110
- utility classes, 26, 27

V

- valves, defining custom, 107-108
- verbose mode, 31
- virtual servers, 104-105
 - default, 105

W

- web application class loader, changing delegation in, 106
- web applications, 73-118
 - default, 75, 105-106
 - security, 41
- web container, logging and monitoring, 107
- web services, 57-60

web services (*Continued*)

debugging, 58, 60

deployment, 58

registry, 58-59

test page, 59-60

URL, 59-60

WSDL file, 59-60

WebDav, 111-113

WSIT, 40

X

XA resource, 138

