



Sun Java System Application Server Enterprise Edition 8.1 2005Q2 Administration Guide



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

Part No: 819-2554-17
January 2010

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, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. 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 dérivés du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. 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	19
1 Getting Started	29
About the Sun Java System Application Server	29
What is the Application Server?	29
Application Server Architecture	30
Tools for Administration	33
Application Server Configuration	34
Configuring the Application Server	34
Configuring Domains	35
Creating a Domain	35
Deleting a Domain	36
Listing Domains	36
Starting the Domain	36
Restarting the Server or Domain	37
Stopping the Domain	37
Recreating the Domain Administration Server	38
Application Server Instances	40
About Application Server Instances	40
Defining Application Server Instances	41
About Stand-Alone Instances	43
Viewing General Server Information	43
▼ To deploy an application	43
▼ To create a new resource type	44
Administration Server Advanced Settings	45
Instance Specific Configuration Properties	47
▼ To create an Instance	48
▼ To start an instance	49

Recovering Transactions	49
▼ To stop an instance	50
▼ To shut down the Administration Server	50
Configuration Changes	50
Changing Application Server Configuration	51
Ports in the Application Server	51
▼ To view port numbers	52
▼ To change the Administrative Server port	52
▼ To change an HTTP port	52
▼ To change an IIOP port	53
▼ To configure a JMX connector using the Admin Service	53
▼ To edit the JMX connector configuration	54
Changing the J2SE Software	55
Using Online Help	55
Further Information	55
2 Deploying Applications	57
About Deployment	57
The Deployment Life Cycle	57
Types of J2EE Archive Files	59
Naming Conventions	59
Admin Console Tasks for Deploying Applications	60
▼ To deploy an enterprise application	60
▼ To edit a deployed enterprise application	62
▼ To deploy a web application	63
▼ To launch a deployed web application	65
▼ To deploy an EJB module	65
▼ To deploy a connector module	67
▼ To create a lifecycle module	68
▼ To deploy an application client module	70
▼ To specify an archive file for deployment	71
Admin Console Tasks for Listing, Undeploying, and Enabling Applications	72
▼ To list deployed applications	72
▼ To list subcomponents	73
▼ To view module descriptors of deployed applications	73

▼ To undeploy an application or module	73
▼ To enable or disable an application or module	74
▼ To manage application targets	74
▼ To deploy on additional virtual servers	75
Redeploying to Multiple Targets	75
▼ To configure dynamic reloading	76
Deployment Methods for Developers	77
▼ To use auto deploy	77
▼ To deploy an unpackaged application from a directory	78
Using the <code>deploytool</code> Utility	79
Using a Deployment Plan	79
3 JDBC Resources	81
About JDBC Resources and Connection Pools	81
JDBC Resources	81
JDBC Connection Pools	82
How JDBC Resources and Connection Pools Work Together	82
Setting Up Database Access	83
▼ To Set Up Database Access	83
▼ To integrate a JDBC driver	83
About JDBC Connection Pools	84
▼ To create a JDBC connection pool	84
▼ To edit a JDBC connection pool	85
▼ To verify connection pool settings	88
▼ To delete a JDBC connection pool	88
About JDBC Resources	88
▼ To create a JDBC resource	88
▼ To edit a JDBC resource	89
▼ To delete a JDBC resource	90
▼ To enable or disable a JDBC resource	90
About Persistence Manager Resources	91
▼ To create a persistence manager resource	91
▼ To edit a persistence manager resource	92
▼ To manage resource targets	92
▼ To delete a persistence manager resource	92

- ▼ To enable or Disable a persistence manager resource 93

- 4 Configuring Java Message Service Resources 95**
 - About JMS Resources 95
 - The JMS Provider in the Application Server 95
 - JMS Resources 95
 - The Relationship Between JMS Resources and Connector Resources 97
 - Admin Console Tasks for JMS Connection Factories 97
 - ▼ To create a JMS connection factory resource 97
 - ▼ To edit a JMS connection factory resource 100
 - ▼ To delete a JMS connection factory resource 101
 - Admin Console Tasks for JMS Destination Resources 101
 - ▼ To create a JMS destination resource 101
 - ▼ To edit a JMS destination resource 102
 - ▼ To delete a JMS destination resource 103
 - Admin Console Tasks for JMS Physical Destinations 103
 - ▼ To create a JMS physical destination 103
 - ▼ To delete a JMS physical destination 105
 - Admin Console Tasks for the JMS Provider 105
 - ▼ To configure general properties for the JMS provider 105
 - ▼ To verify JMS Service settings 109
 - ▼ To create a JMS host 109
 - ▼ To edit a JMS host 111
 - ▼ To delete a JMS host 111

- 5 Configuring JavaMail Resources 113**
 - About JavaMail 113
 - Admin Console Tasks for JavaMail 114
 - ▼ To create a JavaMail session 114
 - ▼ To edit a JavaMail session 115
 - ▼ To delete a JavaMail session 116

- 6 JNDI Resources 117**
 - About Java Naming and Directory Interface (JNDI) 117

JNDI Names and Resources	117
J2EE Naming Services	118
Naming References and Binding Information	118
About Custom Resources	119
Using Custom Resources	119
▼ To create a custom resource	120
▼ To edit a custom resource	120
▼ To delete a custom resource	121
Listing Custom Resources	121
About External JNDI Repositories and Resources	122
Using External JNDI Repositories and Resources	122
▼ To create an external resource	122
▼ To edit an external resource	123
▼ To delete an external resource	124
Listing External Resources	124
7 Connector Resources	125
About Connectors	125
Admin Console Tasks for Connector Connection Pools	126
▼ To Set Up EIS Access	126
▼ To create a connector connection pool	126
▼ To edit a connector connection pool	127
▼ To delete a connector connection pool	129
Admin Console Tasks for Connector Resources	129
▼ To create a connector resource	129
▼ To edit a connector resource	130
▼ Deleting a Connector Resource	131
▼ To configure the connector service	131
Admin Console Tasks for Administered Object Resources	132
▼ To create an administered object resource	132
▼ To edit an administered object resource	133
▼ To delete an administered object resource	133
8 J2EE Containers	135
About the J2EE Containers	135

Types of J2EE Containers	135
The Web Container	135
The EJB Container	136
Admin Console Tasks for the J2EE Containers	136
Configuring the General Web Container Settings	136
Configuring Web Container Sessions	136
Configuring the General EJB Settings	139
Configuring the Message-Driven Bean Settings	141
Configuring the EJB Timer Service Settings	142
9 Configuring Security	145
About Application Server Security	145
Overview of Security	145
About Authentication and Authorization	151
Understanding Users, Groups, Roles, and Realms	153
Introduction to Certificates and SSL	156
About Firewalls	159
Managing Security With the Admin Console	159
Admin Console Tasks for Security	162
▼ To configure security settings	162
▼ To grant access to administration tools	163
Admin Console Tasks for Realms	164
▼ To create a realm	164
▼ To edit a realm	166
▼ To delete a realm	167
▼ To set the default realm	167
Additional Information for Specific Realms	168
Admin Console Tasks for JACC Providers	176
▼ To create a JACC provider	176
▼ To edit a JACC provider	177
▼ To delete a JACC provider	178
▼ To set the active JACC provider	178
Admin Console Tasks for Audit Modules	179
▼ To create an audit module	179
▼ To edit an audit module	180

▼ To delete an audit module	181
▼ To enable or disable audit logging	181
▼ To set the active audit module	182
▼ To use the default audit module	182
Admin Console Tasks for Listeners and JMX Connectors	183
▼ To configure security for HTTP listeners	183
▼ To configure security for IIOP listeners	184
▼ To configure security for the Admin Service's JMX connector	185
▼ To set listener security properties	185
Admin Console Security Tasks for Virtual Servers	186
▼ To configure single sign-on (SSO)	186
Admin Console Tasks for Connector Connection Pools	187
About Connector Connection Pools	188
About Security Maps	188
Working with Certificates and SSL	191
About Certificate Files	191
Using Java Secure Socket Extension (JSSE) Tools	192
Using Network Security Services (NSS) Tools	196
Further Information	199
10 Configuring Message Security	201
About Message Security	201
Overview of Message Security	201
Understanding Message Security in the Application Server	202
Securing a Web Service	206
Securing the Sample Application	207
Configuring the Application Server for Message Security	207
Admin Console Tasks for Message Security	210
▼ To enable providers for message security	211
▼ To configure a message security provider	213
▼ Creating a Message Security Provider	215
▼ To delete a message security configuration	217
▼ To delete a message security provider	218
▼ To enable message security for application clients	219
Setting the Request and Response Policy for the Application Client Configuration	220

Further Information	221
11 Transactions	223
About Transactions	223
What is a Transaction?	223
Transactions in J2EE Technology	224
Admin Console Tasks for Transactions	224
Configuring Transactions	224
12 Configuring the HTTP Service	229
About the HTTP Service	229
What Is the HTTP Service?	229
Virtual Servers	229
HTTP Listeners	231
Admin Console Tasks for the HTTP Service	233
▼ To configure the HTTP Service	233
▼ To configure the HTTP Service access log	235
▼ To configure HTTP Service request processing threads	236
▼ To configure the HTTP Service keep-alive subsystem	237
▼ Configuring the HTTP Service Connection Pool	237
▼ To configure the HTTP protocol for the HTTP Service	238
▼ To configure the HTTP file cache for the HTTP Service	239
Admin Console Tasks for Virtual Servers	240
▼ To create a virtual server	240
▼ To edit a virtual server	243
▼ To delete a virtual server	243
Admin Console Tasks for HTTP Listeners	244
▼ To create an HTTP listener	244
▼ To edit an HTTP listener	246
▼ To delete an HTTP listener	247
13 Configuring the Object Request Broker	249
About the Object Request Broker	249
CORBA	249

What is the ORB?	250
IIOP Listeners	250
Admin Console Tasks for the ORB	250
▼ To configure the ORB	250
Admin Console Tasks for IIOP Listeners	251
▼ To create an IIOP listener	251
▼ To edit an IIOP listener	253
▼ To delete an IIOP listener	253
14 Thread Pools	255
About Thread Pools	255
Admin Console Tasks for Thread Pools	256
▼ To create a thread pool	256
▼ To edit a thread pool	257
▼ To delete a thread pool	258
15 Configuring Logging	259
About Logging	259
Log Records	259
The Logger Namespace Hierarchy	260
Admin Console Tasks for Logging	262
▼ To configure general logging settings	262
▼ To configure log levels	263
▼ To view the server log	264
16 Monitoring Components and Services	267
About Monitoring	267
Monitoring in the Application Server	267
Overview of Monitoring	267
About the Tree Structure of Monitorable Objects	268
About Statistics for Monitored Components and Services	271
Admin Console Tasks for Enabling and Disabling Monitoring	294
▼ To configure monitoring levels using the Admin Console	294
▼ To configure monitoring using the asadmin tool	296

Admin Console Tasks for Viewing Monitoring Data	296
▼ To view monitoring data in the Admin Console	296
Viewing Monitoring Data With the asadmin Tool	299
Using JConsole	314
▼ To disable security for the JMX Connector	314
17 Java Virtual Machine and Advanced Settings	317
Admin Console Tasks for JVM Settings	317
▼ To configure the JVM general settings	317
▼ To configure the JVM classpath settings	318
▼ To configure the JVM options	319
▼ To disable the security manager	320
▼ To configure the JVM profiler settings	321
Admin Console Tasks for Advanced Settings	321
▼ To set the advanced domain attributes	321
18 Compiling and Configuring Apache Web Server	323
Installing Apache	323
Minimum Requirements for Apache 1.3	323
Minimum Requirements for Apache 2	324
▼ To install SSL-aware Apache	325
19 Automatically Restarting a Domain or Node Agent	329
Restarting Automatically on UNIX Platforms	329
Restarting Automatically on the Microsoft Windows Platform	330
Security for Automatic Restarts	331
20 Dotted Name Attributes for domain.xml	333
Top Level Elements	333
Elements Not Aliased	335
Index	337

Figures

FIGURE 1-1	Application Server Architecture	31
FIGURE 1-2	Application Server Instance	42
FIGURE 9-1	Role Mapping	154

Tables

TABLE 1-1	Domain Attributes values	46
TABLE 1-2	Application Server Listeners that Use Ports	51
TABLE 6-1	JNDI Lookups and Their Associated References	119
TABLE 9-1	Application Server Authentication Methods	152
TABLE 9-2	Required properties for ldap realm	168
TABLE 9-3	Optional properties for ldap realm	169
TABLE 9-4	Example ldap realm values	169
TABLE 9-5	Optional properties for certificate realm	172
TABLE 9-6	Required properties for file realms	172
TABLE 10-1	Message protection policy to WS-Security SOAP message security operation mapping	208
TABLE 15-1	Application Server Logger Namespaces	260
TABLE 16-1	EJB Statistics	272
TABLE 16-2	EJB Method Statistics	272
TABLE 16-3	EJB Session Store Statistics	273
TABLE 16-4	EJB Pool Statistics	274
TABLE 16-5	EJB Cache Statistics	275
TABLE 16-6	Timer Statistics	276
TABLE 16-7	Web Container (Servlet) Statistics	276
TABLE 16-8	Web Container (Web Module) Statistics	276
TABLE 16-9	HTTP Service Statistics (applicable to Platform Edition only)	278
TABLE 16-10	JDBC Connection Pool Statistics	279
TABLE 16-11	Connector Connection Pool Statistics	280
TABLE 16-12	Connector Work Management Statistics	281
TABLE 16-13	Connection Manager (in an ORB) Statistics	281
TABLE 16-14	Thread Pool Statistics	282
TABLE 16-15	Transaction Service Statistics	282
TABLE 16-16	JVM Statistics	283
TABLE 16-17	JVM Statistics for J2SE 5.0 - Class Loading	284

TABLE 16-18	JVM Statistics for J2SE 5.0 - Compilation	284
TABLE 16-19	JVM Statistics for J2SE 5.0 - Garbage Collection	284
TABLE 16-20	JVM Statistics for J2SE 5.0 - Memory	284
TABLE 16-21	JVM Statistics for J2SE 5.0 - Operating System	285
TABLE 16-22	JVM Statistics for J2SE 5.0 - Runtime	285
TABLE 16-23	JVM Statistics for J2SE 5.0 - Thread Info	286
TABLE 16-24	JVM Statistics for J2SE 5.0 - Threads	287
TABLE 16-25	PWC Virtual Server Statistics (EE only)	288
TABLE 16-26	PWC Request Statistics (EE only)	288
TABLE 16-27	PWC File Cache Statistics (EE only)	290
TABLE 16-28	PWC Keep Alive Statistics (EE only)	291
TABLE 16-29	PWC DNS Statistics (EE only)	291
TABLE 16-30	PWC Thread Pool Statistics (EE only)	292
TABLE 16-31	PWC Connection Queue Statistics (EE only)	293
TABLE 16-32	PWC HTTP Service Statistics (EE only)	293
TABLE 16-33	Top Level	307
TABLE 16-34	Applications Level	307
TABLE 16-35	Applications - Enterprise Applications and Standalone Modules	308
TABLE 16-36	HTTP-Service Level	311
TABLE 16-37	Thread-Pools Level	312
TABLE 16-38	Resources Level	312
TABLE 16-39	Transaction-Service Level	313
TABLE 16-40	ORB Level	313
TABLE 16-41	JVM Level	314

Examples

EXAMPLE 16-1	Applications Node Tree Structure	269
EXAMPLE 16-2	HTTP Service Schematic (Platform Edition version)	269
EXAMPLE 16-3	HTTP Service Schematic (Enterprise Edition version)	270
EXAMPLE 16-4	Resources Schematic	270
EXAMPLE 16-5	Connector Service Schematic	270
EXAMPLE 16-6	JMS Service Schematic	271
EXAMPLE 16-7	ORB Schematic	271
EXAMPLE 16-8	Thread Pool Schematic	271

Preface

The *Administration Guide* describes how to configure, manage, and deploy Application Server subsystems and components from the Administration Console.

Who Should Use This Book

This *Administration Guide* is intended for information technology administrators in production environments. This guide assumes you are familiar with the following topics:

- Basic system administration tasks
- Installing software
- Using Web browsers
- Starting database servers
- Issuing commands in a terminal window

Before You Read This Book

Application Server is a component of Sun Java™ Enterprise System, a software infrastructure that supports enterprise applications distributed across a network or Internet environment. You should be familiar with the documentation provided with Sun Java Enterprise System, which can be accessed online at: <http://docs.sun.com/app/docs/prod/entsys.05q4#hic>

How This Book Is Organized

The organization of this guide corresponds to the layout of the Administration Console, the browser-based tool for administering the Application Server. Each chapter begins with conceptual information, followed by procedural sections that explain how to perform specific tasks with the Administration Console.

Conventions Used in This Book

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

Typographic Conventions

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

TABLE P-1 Typographic Conventions

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

Symbols

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

TABLE P-2 Symbol Conventions

Symbol	Description	Example	Meaning
[]	Contains optional command options.	<code>ls [-l]</code>	The <code>-l</code> option is not required.
{ }	Contains a set of choices for a required command option.	<code>-d {y n}</code>	The <code>-d</code> option requires that you use either the <code>y</code> argument or the <code>n</code> argument.
-	Joins simultaneous multiple keystrokes.	Control-A	Press the Control key while you press the A key.

TABLE P-2 Symbol Conventions (Continued)

Symbol	Description	Example	Meaning
+	Joins consecutive multiple keystrokes.	Ctrl+A+N	Press the Control key, release it, and then press the subsequent keys.
>	Indicates menu item selection in a graphical user interface.	File > New > Templates	From the File menu, choose New. From the New submenu, choose Templates.

Default Paths and File Names

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

TABLE P-3 Default Paths and File Names

Term	Description
<i>install-dir</i>	<p>By default, the Application Server installation directory is located here:</p> <ul style="list-style-type: none"> ■ Sun Java Enterprise System installations on the Solaris™ platform: /opt/SUNWappserver/appserver ■ Sun Java Enterprise System installations on the Linux platform: /opt/sun/appserver/ ■ Other Solaris and Linux installations, non-root user: <i>user's home directory</i>/SUNWappserver ■ Other Solaris and Linux installations, root user: /opt/SUNWappserver ■ Windows, all installations: <i>SystemDrive</i>: \Sun\AppServer
<i>domain-root-dir</i>	<p>By default, the directory containing all domains is located here:</p> <ul style="list-style-type: none"> ■ Sun Java Enterprise System installations on the Solaris platform: /var/opt/SUNWappserver/domains/ ■ Sun Java Enterprise System installations on the Linux platform: /var/opt/sun/appserver/domains/ ■ All other installations: <i>install-dir</i>/domains/

TABLE P-3 Default Paths and File Names (Continued)

Term	Description
<i>domain-dir</i>	By default, each domain directory is located here: <i>domain-root-dir/domain-dir</i> In configuration files, you might see <i>domain-dir</i> represented as follows: \${com.sun.aas.instanceRoot}
<i>instance-dir</i>	By default, each instance directory is located here: <i>domain-dir/instance-dir</i>

Shell Prompts

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

TABLE P-4 Shell Prompts

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

Symbols

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

TABLE P-5 Symbol Conventions

Symbol	Description	Example	Meaning
[]	Contains optional command options.	ls [-l]	The -l option is not required.
{ }	Contains a set of choices for a required command option.	-d {y n}	The -d option requires that you use either the y argument or the n argument.

TABLE P-5 Symbol Conventions (Continued)

Symbol	Description	Example	Meaning
-	Joins simultaneous multiple keystrokes.	Control-A	Press the Control key while you press the A key.
+	Joins consecutive multiple keystrokes.	Ctrl+A+N	Press the Control key, release it, and then press the subsequent keys.
>	Indicates menu item selection in a graphical user interface.	File > New > Templates	From the File menu, choose New. From the New submenu, choose Templates.

Default Paths and File Names

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

TABLE P-6 Default Paths and File Names

Term	Description
<i>install-dir</i>	<p>By default, the Application Server installation directory is located here:</p> <ul style="list-style-type: none"> ■ Sun Java Enterprise System installations on the Solaris platform: /opt/SUNWappserver/appserver ■ Sun Java Enterprise System installations on the Linux platform: /opt/sun/appserver/ ■ Other Solaris and Linux installations, non-root user: <i>user's home directory</i>/SUNWappserver ■ Other Solaris and Linux installations, root user: /opt/SUNWappserver ■ Windows, all installations: <i>SystemDrive</i>: \Sun\AppServer

TABLE P-6 Default Paths and File Names (Continued)

Term	Description
<i>domain-root-dir</i>	<p>By default, the directory containing all domains is located here:</p> <ul style="list-style-type: none"> ■ Sun Java Enterprise System installations on the Solaris platform: <i>/var/opt/SUNWappserver/domains/</i> ■ Sun Java Enterprise System installations on the Linux platform: <i>/var/opt/sun/appserver/domains/</i> ■ All other installations: <i>install-dir/domains/</i>
<i>domain-dir</i>	<p>By default, each domain directory is located here: <i>domain-root-dir/domain-dir</i></p> <p>In configuration files, you might see <i>domain-dir</i> represented as follows: \${com.sun.aas.instanceRoot}</p>
<i>instance-dir</i>	<p>By default, each instance directory is located here: <i>domain-dir/instance-dir</i></p>

Related Documentation

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

You can find a directory of URLs for the official specifications at install-dir/docs/index.htm. Additionally, the following resources might be useful.

General J2EE Information:

The J2EE 1.4 Tutorial: <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>

The J2EE Blueprints: <http://java.sun.com/reference/blueprints/index.html>

Core J2EE Patterns: Best Practices and Design Strategies by Deepak Alur, John Crupi, & Dan Malks, Prentice Hall Publishing

Java Security, by Scott Oaks, O'Reilly Publishing

Programming with Servlets and JSP files:

Java Servlet Programming, by Jason Hunter, O'Reilly Publishing

Java Threads, 2nd Edition, by Scott Oaks & Henry Wong, O'Reilly Publishing

Programming with EJB components:

Enterprise JavaBeans, by Richard Monson-Haefel, O'Reilly Publishing

Programming with JDBC:

Database Programming with JDBC and Java, by George Reese, O'Reilly Publishing

JDBC Database Access With Java: A Tutorial and Annotated Reference (Java Series), by Graham Hamilton, Rick Cattell, & Maydene Fisher

Javadocs:

Javadocs for packages provided with the Application Server are located in *install-dir/docs/api*.

Books in This Documentation Set

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

The following table summarizes the books included in the Application Server core documentation set.

TABLE P-7 Books in This Documentation Set

Book Title	Description
<i>Release Notes</i>	Late-breaking information about the software and the documentation. Includes a comprehensive, table-based summary of the supported hardware, operating system, JDBC/RDBMS.
<i>Quick Start Guide</i>	How to get started with the Application Server product.
<i>Installation Guide</i>	Installing the Application Server software and its components.
<i>Deployment Planning Guide</i>	Evaluating your system needs and enterprise to ensure that you deploy Sun Java System Application Server in a manner that best suits your site. General issues and concerns must be aware of when deploying an application server are also discussed.
<i>Developer's Guide</i>	Creating and implementing Java™ 2 Platform, Enterprise Edition (J2EE™ platform) applications intended to run on the Sun Java System Application Server that follow the Java standards model for J2EE components and APIs. Includes general information on developer tools, security, assembly, deployment, debugging, and creating lifecycle methods.
<i>J2EE 1.4 Tutorial</i>	Using J2EE 1.4 platform technologies and APIs to develop J2EE applications and deploy the applications on the Sun Java System Application Server.
<i>Administration Guide</i>	Configuring, managing, and deploying the Sun Java System Application Server subsystems and components from the Administration Console.

TABLE P-7 Books in This Documentation Set (Continued)

Book Title	Description
<i>High Availability Administration Guide</i>	Post-installation configuration and administration instructions for the high-availability database.
<i>Administration Reference</i>	Editing the Sun Java System Application Server configuration file, <code>domain.xml</code> .
<i>Upgrade and Migration Guide</i>	Migrating your applications to the new Sun Java System Application Server programming model, specifically from Application Server 6.x and 7. This guide also describes differences between adjacent product releases and configuration options that can result in incompatibility with the product specifications.
<i>Performance Tuning Guide</i>	Tuning the Sun Java System Application Server to improve performance.
<i>Troubleshooting Guide</i>	Solving Sun Java System Application Server problems.
<i>Error Message Reference</i>	Solving Sun Java System Application Server error messages.
<i>Reference Manual</i>	Utility commands available with the Sun Java System Application Server; written in manpage style. Includes the <code>asadmin</code> command line interface.

Other Server Documentation

For other server documentation, go to the following:

- Directory Server documentation http://docs.sun.com/coll/DirectoryServer_04q2
(http://docs.sun.com/coll/DirectoryServer_04q2)
- Web Server documentation http://docs.sun.com/coll/S1_websvr61_en
(http://docs.sun.com/coll/S1_websvr61_en)

Accessing Sun Resources Online

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

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

Contacting Sun Technical Support

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

Related Third-Party Web Site References

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

Sun Welcomes Your Comments

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

To share your comments, go to <http://docs.sun.com> (<http://docs.sun.com>) and click Send Comments. In the online form, provide the document title and part number. The part number is a seven-digit or nine-digit number that can be found on the title page of the book or at the top of the document. For example, the title of this book is *Sun Java System Application Server 2005Q2 Administration Guide*, and the part number is 819-2554.

Getting Started

This chapter describes the Sun Java™ System Application Server and introduces basic administration tasks. This chapter contains following sections:

- “About the Sun Java System Application Server” on page 29
- “Application Server Configuration” on page 34
- “Application Server Instances” on page 40
- “Configuration Changes” on page 50

About the Sun Java System Application Server

- “What is the Application Server?” on page 29
- “Application Server Architecture” on page 30
- “Tools for Administration” on page 33

What is the Application Server?

The Application Server provides a robust J2EE platform for the development, deployment, and management of enterprise applications. Key features include transaction management, performance, scalability, security, and integration. The Application Server supports services from Web publishing to enterprise-scale transaction processing, while enabling developers to build applications based on JavaServer Pages (JSP™), Java servlets, and Enterprise JavaBeans™ (EJB™) technology.

The Application Server Enterprise Edition provides advanced clustering and failover technologies. These features enable you to run scalable and highly available J2EE applications.

- **Clustering** - A cluster is a group of application server instances that work together as one logical entity. Each Application Server instance in the cluster has the same configuration and the same applications deployed to it.

Horizontal scaling is achieved by adding Application Server instances to a cluster, thereby increasing the capacity of the system. It is possible to add Application Server instances to a cluster without disrupting service. The HTTP, RMI/IIOP, and JMS load balancing systems distribute requests to healthy Application Server instances in the cluster.

- **High Availability** - Availability allows for failover protection of Application Server instances in a cluster. If one application server instance goes down, another Application Server instance takes over the sessions that were assigned to the unavailable server. Session information is stored in the high-availability database (HADB). HADB supports the persistence of HTTP sessions and stateful session beans.

Application Server Architecture

This section describes [Figure 1-1](#), which shows the high-level architecture of the Application Server.

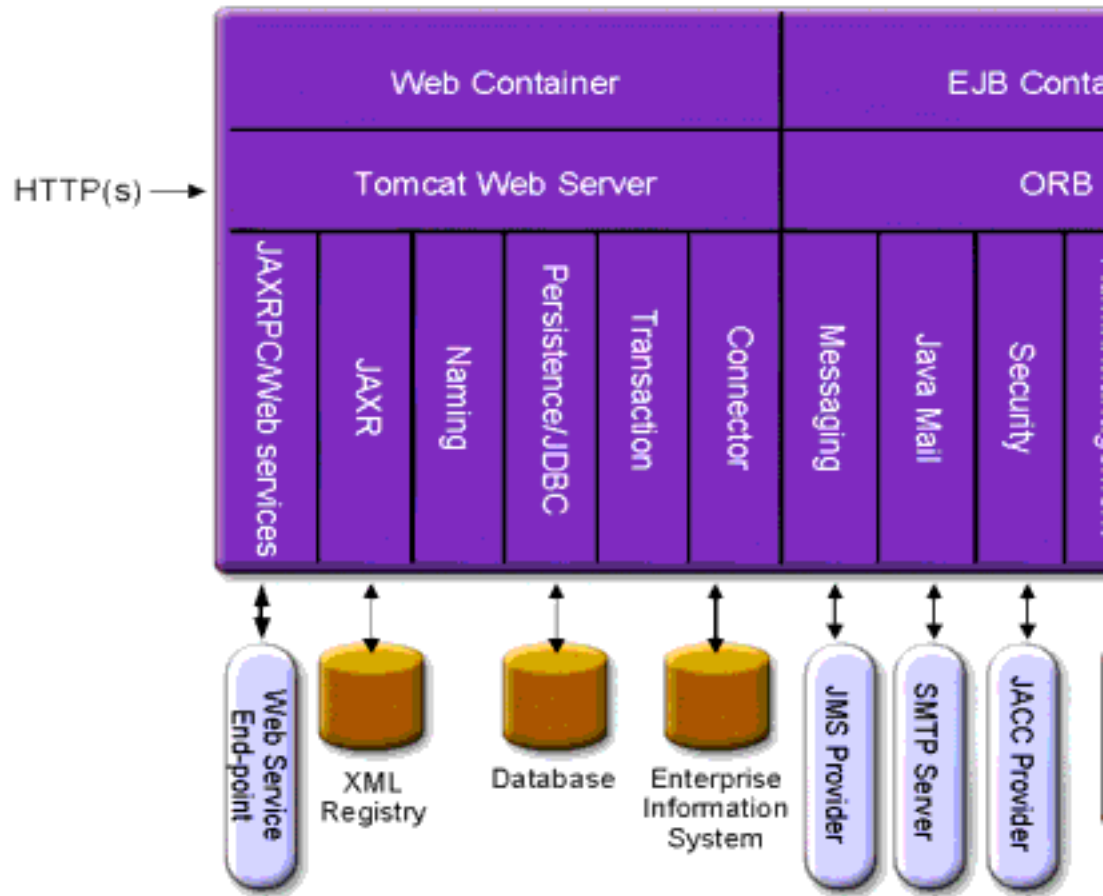


FIGURE 1-1 Application Server Architecture

- Containers** - A container is a runtime environment that provides services such as security and transaction management to J2EE components. Figure 1-1 shows the two types of J2EE containers: Web and EJB. Web components, such as JSP pages and servlets, run within the Web container. Enterprise beans, the components of EJB technology, run within the EJB container.
- Client Access** - At runtime, browser clients access Web applications by communicating with the Web server via HTTP, the protocol used throughout the internet. The HTTPS protocol is for applications that require secure communication. Enterprise bean clients communicate with the Object Request Broker (ORB) through the IIOP or IIOP/SSL (secure) protocols. The Application Server has separate listeners for the HTTP, HTTPS, IIOP, and IIOP/SSL protocols. Each listener has exclusive use of a specific port number.

- **Web Services** - On the J2EE platform, it is possible to deploy a Web application that provides a Web service implemented by Java API for XML-Based RPC (JAX-RPC). A J2EE application or component can also be a client to other Web services. Applications access XML registries through the Java API for XML Registries (JAXR).
- **Services for Applications** - The J2EE platform was designed so that the containers provide services for applications. [Figure 1–1](#) shows the following services:
 - **Naming** - A naming and directory service binds objects to names. A J2EE application locates an object by looking up its JNDI name. JNDI stands for the Java Naming and Directory Interface API.
 - **Security** - The Java Authorization Contract for Containers (JACC) is a set of security contracts defined for the J2EE containers. Based on the client's identity, the containers restrict access to the container's resources and services.
- **Transaction management** - A transaction is an indivisible unit of work. For example, transferring funds between bank accounts is a transaction. A transaction management service ensures that a transaction either completes fully or is rolled back.

Access to External Systems

The J2EE platform enables applications to access systems that are outside of the application server. Applications connect to these systems through objects called resources. One of the responsibilities of an administrator is resource configuration. The J2EE platform enables access to external systems through the following APIs and components:

- **JDBC** - A database management system (DBMS) provides facilities for storing, organizing, and retrieving data. Most business applications store data in relational databases, which applications access via the JDBC API. The information in databases is often described as persistent because it is saved on disk and exists after the application ends. The Application Server bundle includes the PointBase DBMS.
- **Messaging** - Messaging is a method of communication between software components or applications. A messaging client sends messages to, and receives messages from, any other client. Applications access the messaging provider through the Java Messaging Service (JMS) API. The Application Server includes a JMS provider.
- **Connector** - The J2EE Connector architecture enables integration between J2EE applications and existing Enterprise Information Systems (EIS). An application accesses an EIS through a portable J2EE component called a connector or resource adapter.
- **JavaMail** - Through the JavaMail API, applications connect to an SMTP server in order to send and receive email.
- **Server Administration** - The lower right-hand corner of [Figure 1-1](#) shows some of the tasks performed by the administrator of the Application Server. For example, an administrator deploys (installs) applications and monitors the server's performance. These tasks are performed with the administration tools provided by the Application Server.

Tools for Administration

- The Application Server includes three administrative tools:
 - “Admin Console” on page 33
 - “asadmin Utility” on page 33
 - “Application Server Management Extension (AMX)” on page 34

Admin Console

The Admin Console is a browser-based tool that features an easy-to-navigate interface and online help. This manual provides step-by-step instructions for using the Admin Console. The administration server must be running to use the Admin Console.

When the Application Server was installed, you chose a port number for the server, or used the default port of 4849. You also specified a user name and master password.

To start the Admin Console, in a web browser type:

```
https://hostname:port
```

For example:

```
https://kindness.sun.com:4849
```

If the Admin Console is running on the machine on which the Application Server was installed, specify `localhost` for the host name.

On Windows, start the Application Server Admin Console from the Start menu.

The installation program creates the default administrative domain (named `domain1`) with the default port number 4849, as well as an instance separate from the domain administration server (DAS). After installation, additional administration domains can be created. Each domain has its own domain administration server, which has a unique port number. When specifying the URL for the Admin Console, be sure to use the port number for the domain to be administered.

If your configuration includes remote server instances, create node agents to manage and facilitate remote server instances. It is the responsibility of the node agent to create, start, stop, and delete a server instance. Use the command line interface (CLI) commands to set up node agents.

asadmin Utility

The `asadmin` utility is a command-line tool. Use the `asadmin` utility and the commands associated with it to perform the same set of tasks that can be performed in the Admin Console. For example, start and stop domains, configure the server, and deploy applications.

Use these commands either from a command prompt in the shell, or call them from other scripts and programs. Use these commands to automate repetitive administration tasks.

To start the `asadmin` utility:

```
$ asadmin
```

To list the commands available within `asadmin`:

```
asadmin> help
```

It is also possible to issue an `asadmin` command at the shell's command prompt:

```
$ asadmin help
```

To view a command's syntax and examples, type `help` followed by the command name. For example:

```
asadmin> help create-jdbc-resource
```

The `asadmin help` information for a given command displays the UNIX man page of the command. These man pages are also available in HTML format.

Application Server Management Extension (AMX)

The Application Server Management eXtension is an API that exposes all of the Application Server configuration and monitoring JMX managed beans as easy-to-use client-side dynamic proxies implementing the AMX interfaces.

For more information on using the Application Server Management Extension, see the JMX chapter in the *Application Server Developer's Guide*.

Application Server Configuration

- [“Configuring the Application Server” on page 34](#)
- [“Configuring Domains” on page 35](#)
- [“Starting the Domain” on page 36](#)
- [“Restarting the Server or Domain” on page 37](#)
- [“Stopping the Domain” on page 37](#)
- [“Recreating the Domain Administration Server” on page 38](#)

Configuring the Application Server

Application Server domains are logical or physical units created to help the administrator manage a system configuration. A domain is broken down into smaller units including instances and node agents. A server instance is a single Java Virtual Machine (JVM) that runs

the Application Server on a single physical machine. Each domain has one or more instances. A domain must also have at least one associated node agent for the instance to function properly. Domains can be grouped together to create a cluster. Clusters allow the administrator to manage groups of hardware and software.

Configuring Domains

Administrative domains provide a basic security structure whereby different administrators can administer specific groups (domains) of application server instances. By grouping the server instances into separate domains, different organizations and administrators can share a single Application Server installation. Each domain has its own configuration, log files, and application deployment areas that are independent of other domains. If the configuration is changed for one domain, the configurations of other domains are not affected.

Each Administration console session allows you to configure and manage the domain. If you have created multiple domains, you must start an additional Administration Console session to manage each domain. Each domain has its own Domain Administration Server (DAS), with a unique port number. Each administrative domain can have multiple application server instances. However, an application server instance can belong to just one domain. When the Application Server is installed, an administrative domain named `domain1` is automatically created.

Creating a Domain

Domains are created using the `create-domain` command. The following example command creates a domain named `mydomain`. The administration server listens on port 1234 and the administrative user name is `hanan`. The command prompts for the administrative and master passwords.

```
$ asadmin create-domain --adminport 80 --adminuser hanan mydomain
```

To start the Admin Console for `mydomain` domain, in a browser, enter the following URL:

```
http://hostname:80
```

For the preceding `create-domain` example, the domain's log files, configuration files, and deployed applications now reside in the following directory:

```
domain-root-dir/mydomain
```

To create the domain's directory in another location, specify the `--domaindir` option. For the full syntax of the command, type `asadmin help create-domain`.

Deleting a Domain

Domains are deleted using the `asadmin delete-domain` command. Only the operating system user (or root) who can administer the domain can execute this command successfully. To delete a domain named `mydomain`, for example, type the following command:

```
$ asadmin delete-domain mydomain
```

Listing Domains

The domains created on a machine can be found using the `asadmin list-domains` command. To list the domains in the default `domain-root-dir` directory, type this command:

```
$ asadmin list-domains
```

To list domains that were created in other directories, specify the `--domain-dir` option.

Starting the Domain

When starting a domain, the administration server and application server instance are started. Once the application server instance is started it runs constantly, listening for and accepting requests. Each domain must be started separately.

To start a domain, type the `asadmin start-domain` command and specify the domain name. For example, to start the default domain (`domain1`), type the following:

```
$ asadmin start-domain --user admin domain1
```

If there is only one domain, omit the domain name. For the full command syntax, type `asadmin help start-domain`. If the password data is omitted, you are prompted to supply it.

The `asadmin start-domain domain1` command should be run by the user who installed the server. If it is not run by that user, the `.asadmintruststore` should be moved or copied from the home directory of installing user to the home directory of the running user.

Note that if the file is moved (not copied) from the installing user's home directory to the running user's home directory, you might experience application upgrade problems, because the upgrade/install user will no longer have the `.asadmintruststore` file in the home directory.

Restoring a Deleted `.asadmintruststore` File

If the `.asadmintruststore` file for a domain has been deleted, you can easily create a new file

▼ To create a new `.asadmintruststore` file

- 1 Use the local `asadmin start-domain` command to start a domain you want to manage.

Note that a local `asadmin` command does not require an `.asadmintruststore` file to start the domain.

- 2 Run any remote `asadmin` command.

Remote `asadmin` commands require the `--user`, `--passwordfile` (`--password`), `--host`, and `--port` options to be specified, and a target domain to be running.

- 3 Accept the new Domain Certificate by pressing “y” when prompted.

On Windows, to start the default domain:

From the Windows Start Menu, select Programs -> Sun Microsystems -> Application Server -> Start Admin Server.

Restarting the Server or Domain

Restarting the server is the same as restarting the domain. To restart the domain or server, stop and start the domain.

Stopping the Domain

Stopping a domain shuts down its administration server and application server instance. When stopping a domain, the server instance stops accepting new connections and then waits for all outstanding connections to complete. This process takes a few seconds because the server instance must complete its shutdown process. While the domain is stopped, the Admin Console or most `asadmin` commands cannot be used.

To stop a domain, type the `asadmin stop-domain` command and specify the domain name. For example, to stop the default domain (`domain1`), type the following:

```
$ asadmin stop-domain domain1
```

If there is only one domain, then the domain name is optional. For the full syntax, type `asadmin help stop-domain`.

To stop the domain using the Admin Console:

- In the tree component, select server (Admin Server) under the Standalone Instances node.
- On the General Information page, click Stop Server.

On Windows, to stop the default domain:

From the Start menu select Programs -> Sun Microsystems -> Application Server-> Stop Admin Server.

Recreating the Domain Administration Server

For mirroring purposes, and to provide a working copy of the Domain Administration Server (DAS), you must have:

- One machine (machine1) that contains the original DAS.
- A second machine (machine2) that contains a cluster with server instances running applications and catering to clients. The cluster is configured using the DAS on the first machine.
- A third backup machine (machine3) where the DAS needs to be recreated in case the first machine crashes.

Note – You must maintain a backup of the DAS from the first machine. Use `asadmin backup-domain` to backup the current domain.

▼ To migrate the DAS

The following steps are required to migrate the Domain Administration Server from the first machine (machine1) to the third machine (machine3).

1 Install the application server on the third machine just as it is installed on the first machine.

This is required so that the DAS can be properly restored on the third machine and there are no path conflicts.

a. Install the application server administration package using the command-line (interactive) mode. To activate the interactive command-line mode, invoke the installation program using the `console` option:

```
./bundle-filename -console
```

You must have root permission to install using the command-line interface.

b. Deselect the option to install default domain.

Restoration of backed up domains is only supported on two machines with same architecture and **exactly** the same installation paths (use same *install-dir* and *domain-root-dir* on both machines).

2 Copy the backup ZIP file from the first machine into the *domain-root-dir* on the third machine. You can also FTP the file.

3 Execute `asadmin restore-domain` command to restore the ZIP file onto the third machine:

```
asadmin restore-domain --filename domain-root-dir/sjsas_backup_v00001.zip domain1
```

You can backup any domain. However, while recreating the domain, the domain name should be same as the original.

4 Change `domain-root-dir/domain1/generated/tmp` directory permissions on the third machine to match the permissions of the same directory on first machine.

The default permissions of this directory are: `?drwx-----?` (or `700`).

For example:

```
chmod 700 domain-root-dir/domain1/generated/tmp
```

The example above assumes you are backing up `domain1`. If you are backing up a domain by another name, you should replace `domain1` above with the name of the domain being backed up.

5 Change the host values for the properties in the `domain.xml` file for the third machine:**6 Update the `domain-root-dir/domain1/config/domain.xml` on the third machine.**

For example, search for `machine1` and replace it with `machine3`. So, you can change:

```
<jmx-connector><property name=client-hostname value=machine1/>...
```

to:

```
<jmx-connector><property name=client-hostname value=machine3/>...
```

7 Change:

```
<jms-service... host=machine1.../>
```

to:

```
<jms-service... host=machine3.../>
```

8 Start the restored domain on machine3:

```
asadmin start-domain --user admin-user --password admin-password domain1
```

9 Change the DAS host values for properties under node agent on machine2.**10 Change `agent.das.host` property value in `install-dir/nodeagents/nodeagent/agent/config/das.properties` on machine2.****11 Restart the node agent on machine2.**

Note – Start the cluster instances using the `asadmin start-instance` command to allow them to synchronize with the restored domain.

Application Server Instances

- [“About Application Server Instances” on page 40](#)
- [“About Stand-Alone Instances” on page 43](#)
- [“Administration Server Advanced Settings” on page 45](#)
- [“Instance Specific Configuration Properties” on page 47](#)
- [“Recovering Transactions” on page 49](#)

About Application Server Instances

The Application Server creates one application server instance, called `server` at the time of installation. You can delete the `server` instance and create a new instance with a different name if you prefer.

Each Application Server instance has its own J2EE configuration, J2EE resources, application deployment areas, and server configuration settings. Changes to one application server instance have no effect on other application server instances. You can have many application server instances within one administrative domain.

For many users, one application server instance meets their needs. However, depending upon your environment, you might want to create one or more additional application server instances. For example, in a development environment you can use different application server instances to test different Application Server configurations, or to compare and test different application deployments. Because you can easily add or delete an application server instance, you can use them to create temporary “sandbox” areas to experiment with while developing.

In addition, for each application server instance you can also create virtual servers. Within a single installed application server instance you can offer companies or individuals domain names, IP Addresses, and some administration capabilities. For the users, it is almost as if they have their own web server, without the hardware and basic server maintenance. These virtual servers do not span application server instances. For more information about virtual servers, see [“To configure the JVM general settings” on page 317](#).

In operational deployments, for many purposes you can use virtual servers instead of multiple application server instances. However, if virtual servers do not meet your needs, you can also use multiple application server instances.

An Application Server instance is not started automatically. Once you start an instance, the instance runs until you stop it. When you stop an application server instance, it stops accepting

new connections, then waits for all outstanding connections to complete. If your machine crashes or is taken offline, the server quits and any requests it was servicing may be lost.

Defining Application Server Instances

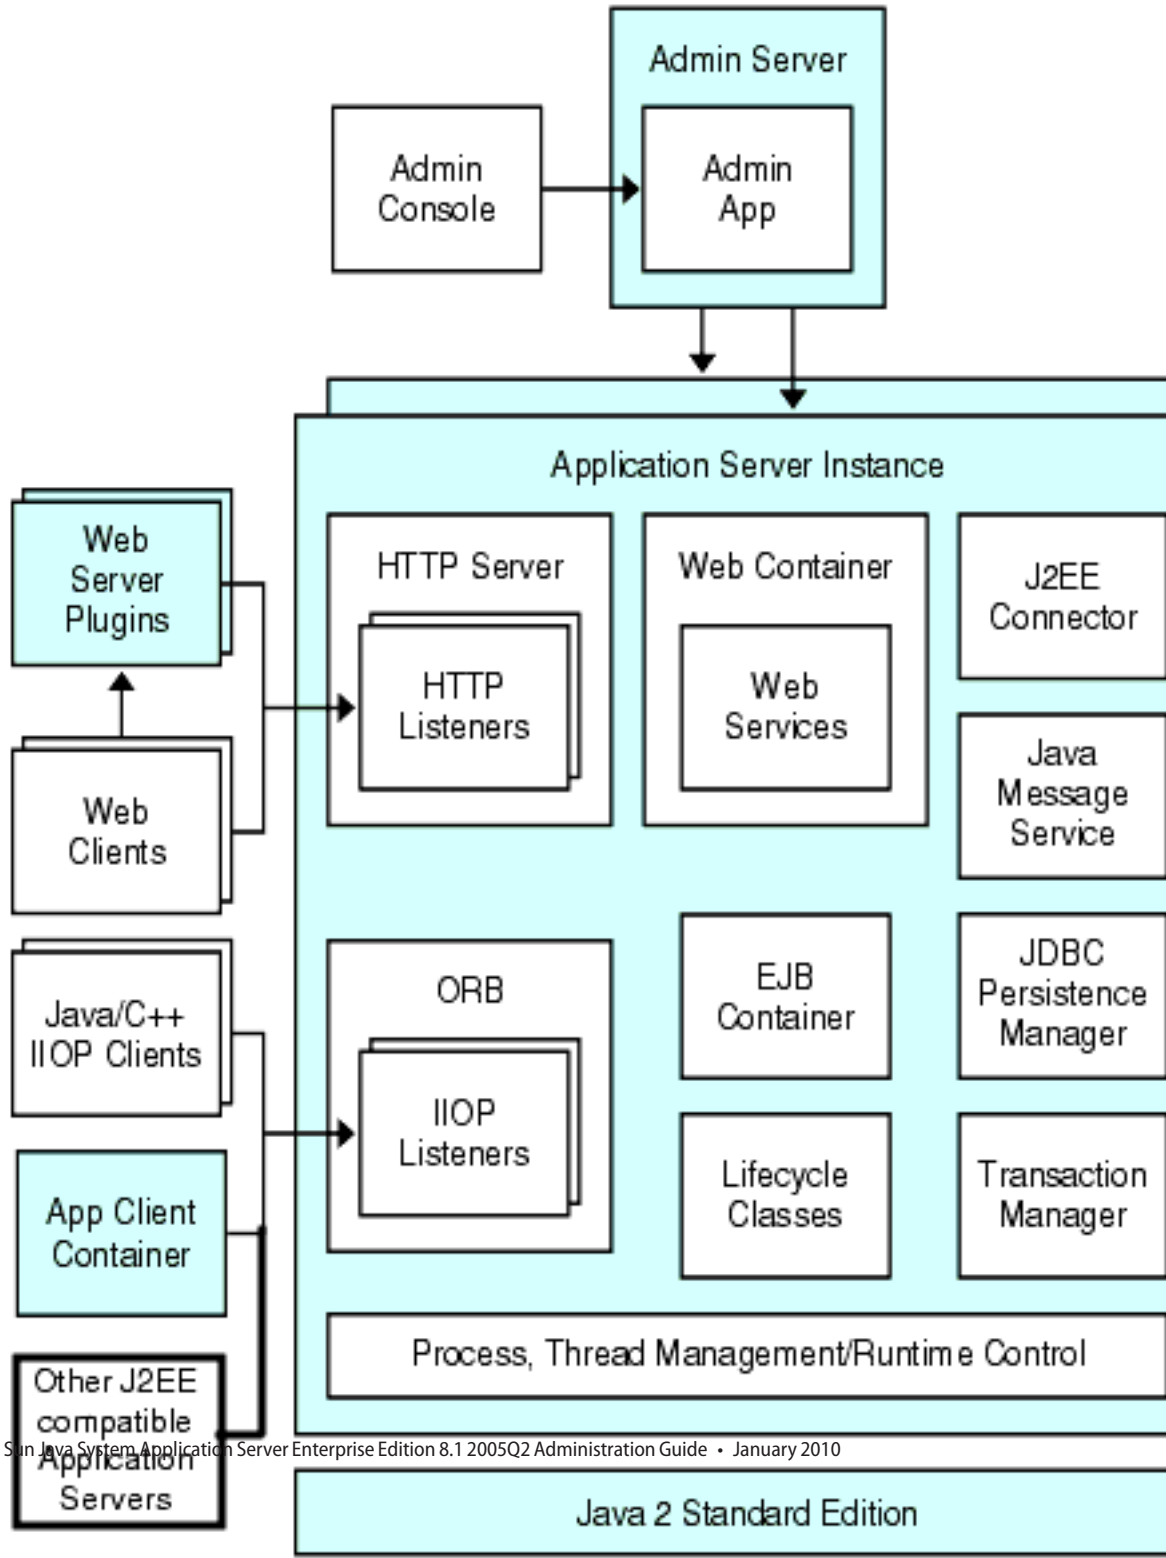
Application server instances form the basis of an application deployment. Each instance belongs to a single domain and has its own directory structure, configuration, and deployed applications. Each server instance also includes the J2EE platform web and EJB containers. Every new server instance must contain a reference to a node agent name defining the machine on which the instance will reside.

There are three types of server instances that can be created. Each server instance can only be of one type:

- In the **stand-alone server** instance the configuration is not shared by any other server instance or clusters.
- In the **shared server instance** the configuration is shared with other instances or clusters.
- In the **clustered server instance** the configuration is shared with other instances in the cluster.

A *cluster* is a group of server instances sharing the same set of applications, resources, and configuration information. A server instance can belong to only one cluster. Among other things, the cluster is used to facilitate load balancing, through distribution of a load across multiple machines, and high availability, through instance level failover.

Figure 1–2 shows an application server instance in detail. The application server instance is a building block in the clustering, load balancing, and session persistence features of the Application Server Enterprise Edition.



About Stand-Alone Instances

An Application Server instance is not started automatically. Once you start an instance, the instance runs until you stop it. When you stop an application server instance, it stops accepting new connections, then waits for all outstanding connections to complete. If your machine crashes or is taken offline, the server quits and any requests it was servicing may be lost.

See Also:

- [“To create an Instance” on page 48](#)
- [“To start an instance” on page 49](#)
- [“To stop an instance” on page 50](#)
- [“Recovering Transactions” on page 49](#)

Viewing General Server Information

From the General Tab you can perform the following tasks:

- Click Start Instance to start the instance.
- Click Stop Instance to stop the instance.
- Click View Log Files to open the server log viewer.
- Click Rotate Log File to rotate the log file for the instance.

This action schedules the log file for rotation. The actual rotation takes place the next time an entry is written to the log file. The rotation happens immediately for the default server (the DAS) but is delayed for other stand-alone server.

- Click JNDI Browsing to browse the JNDI tree for a running instance.
- Click Recover Transactions to recover incomplete transactions.

In addition, you can select the following tabs to perform these additional tasks:

- Applications Tab: deploy a selected application.
- Resources Tab: manage a selected resource.
- Properties Tab: configure instance specific properties.
- Monitor Tab: view monitoring data for JVM, Server, Thread Pools, HTTP Service, and Transaction Service.
- Advanced Tab: set general properties for deploying applications.

▼ To deploy an application

From the Applications Tab, you can enable, disable, and deploy selected applications associated with the instance.

- 1 **Select the checkbox for the desired application.**
- 2 **From the Deploy drop down menu, select the type application module you want to deploy:**
 - Enterprise Application: a J2EE application in an EAR (Enterprise Application Archive) file or directory.
 - Web Application: a collection of Web resources such as JavaServer Pages (JSPs), servlets, and HTML pages that are packaged in a WAR (Web Application Archive) file or directory.
 - EJB Module: one or more Enterprise JavaBeans (EJB components) contained in an EJB JAR (Java Archive) file or directory.
 - Connector Module: connects to an Enterprise Information System (EIS) and is packaged in a RAR (Resource Adapter Archive) file or directory.
 - Lifecycle Module: performs tasks when it is triggered by one or more events in the server's lifecycle.
 - App Client Module: also called a J2EE application client JAR file, contains the server-side routines for the client.

▼ **To create a new resource type**

From the Resources Tab, you can enable, disable, and create a new resource type to associate with the instance.

- 1 **Select the checkbox for the desired resource.**
- 2 **From the New drop down menu, select the resource type you want to create and associate with that instance:**
 - JDBC: provide applications with a means of connecting to a database.
 - Persistence Manager: required for applications with container-managed persistence beans (needed for backward compatibility).
 - JMS Connection Factory: objects that allow an application to create other JMS objects programmatically.
 - JMS Destination: represents a mail session in the JavaMail API, which provides a platform-independent and protocol-independent framework to build mail and messaging applications.
 - JavaMail: provides a platform-independent and protocol-independent framework to build mail and messaging applications.
 - Custom: represents nonstandard resources with a defined JNDI subcontext, resource type, and factory class.
 - External: enables an application to locate an external resource object in a Lightweight Directory Access Protocol (LDAP) repository.

- Connector: a program object that provides an application with a connection to an enterprise information systems (EIS).
- Admin Object: configures a JSR-160 compliant remote JMX connector.

Administration Server Advanced Settings

The Administration Server Advanced settings allow you to set general properties for deploying applications. These properties enable you to ensure and monitor that changes to deployed applications are detected and the modified classes reloaded.

Setting Applications Configurations

If dynamic reloading is enabled, the server periodically checks for changes in the files of the deployed application and automatically reloads the application with the changes. Dynamic reloading is useful in a development environment because it allows code changes to be tested quickly. In a production environment, however, dynamic reloading may degrade performance.

Dynamic reloading is intended for development environments. It is incompatible with session persistence, a production environment feature. Do not enable session persistence if dynamic deployment is enabled.

Note – Dynamic reloading is only available for the default server instance.

To configure dynamic reloading from the Applications Configuration page, configure the following:

- Reload: Enable or disable dynamic reloading with the Enabled checkbox.
- Reload Poll Interval: Specify how often the server checks for changes in the deployed applications.
- Admin Session Timeout: Specify the amount of time before the Admin Session times out and you have to log in again.

▼ To configure auto deploy settings

The auto deploy feature enables you to deploy a prepackaged application or module by copying it to the *domain-dir/autodeploy* directory.

For example, copy a file named `hello.war` to the *domain-dir/autodeploy* directory. To undeploy the application, remove the `hello.war` file from the *autodeploy* directory.

The auto deploy feature is intended for development environments. It is incompatible with session persistence, a production environment feature. Do not enable session persistence if auto deploy is enabled.

Note – Auto deploy is only available for the default server instance.

- 1 Go to the Applications Configuration page.**
- 2 Enable or disable auto deploy by selecting or deselecting the Enabled checkbox.**
- 3 In the Auto Deploy Poll Interval field, specify how often the server checks the auto deploy directory for application or module files.**
Changing the poll interval does not affect the amount of time it takes to deploy an application or module.
- 4 In the Auto Deploy directory, if you specify the directory where you build your application, then you won't have to copy the file to the default auto deploy directory.**
By default a variable is used to eliminate the need to manually change the directory for multiple server instances.
- 5 To run the verifier before deployment, select the Verifier Enabled checkbox.**
The verifier examines the structure and content of the file. Verification of large applications is often time-consuming.
- 6 To precompile JSP pages, select the JSPs checkbox.**
If you do not select this checkbox, the JSP pages are compiled at runtime when they are first accessed. Because compilation is often time-consuming, in a production environment select this checkbox.
- 7 Click the Add Property button to specify additional settings.**

Setting Domain Attributes

The following domain attributes properties are available.

TABLE 1-1 Domain Attributes values

Property	Definition
<code>com.sun.aas.installRoot</code>	Directory where the application server is installed.
<code>com.sun.aas.instanceRoot</code>	Top level directory for a server instance.
<code>com.sun.aas.hostName</code>	Name of the host (machine).
<code>com.sun.aas.javaRoot</code>	.J2SE installation directory.
<code>com.sun.aas.imqLib</code>	Library directory of the Sun Java System Message Queue software.

TABLE 1-1 Domain Attributes values (Continued)

Property	Definition
<code>com.sun.aas.configName</code>	Name of the configuration being used by a server instance.
<code>com.sun.aas.instanceName</code>	Name of the server instance. This property is not available for the default-config but can be used for customized configurations.
<code>com.sun.aas.clusterName</code>	Name of the cluster. This property is only set on the clustered server instance. This property is not available for the default-config but can be used for customized configurations.
<code>com.sun.aas.domainName</code>	Name of the domain. This property is not available for the default-config but can be used for customized configurations.

Instance Specific Configuration Properties

The instance specific Configuration Properties override the values for this instance.

Note – The default values are defined in the configuration bound to the instance.

▼ To revert the value back to the default value

- 1 Remove the override value.
- 2 Click Save.

If no override value is set, the default value is used.

▼ To add instance properties

- Click the Add Property button to specify additional settings.

The following property attribute name/value pairs for configuring the resource are available:

Property	Definition
<code>HTTP_LISTENER_PORT</code>	This port is used to listen for HTTP requests. This property specifies the port number for http-listener-1. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.
<code>HTTP_SSL_LISTENER_PORT</code>	This port is used to listen for HTTPS requests. This property specifies the port number for http-listener-2. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 requires superuser privileges.

Property	Definition
IIOB_LISTENER_PORT	This property specifies which ORB listener port for IIOB connections orb-listener listens on.
IIOB_SSL_LISTENER_PORT	This port is used for secure IIOB connections.
JMX_SYSTEM_CONNECTOR_PORT	This property specifies the port number on which the JMX connector listens. Valid values are 1-65535. On UNIX, creating sockets that listen on ports 1-1024 require superuser privileges.
IIOB_SSL_MUTUALAUTH_PORT	This property specifies which ORB listener port for IIOB connections the IIOB listener called SSL_MUTUALAUTH listens on.

▼ To delete properties

- 1 Click the for the property you wish to delete.
- 2 Click the Delete Property button.

▼ To create an Instance

- 1 In the tree component, select the Standalone Instances node.
- 2 On the Stand Alone Server Instances page, click New.
- 3 In the Name field, identify a unique name for the new instance.
- 4 Select a Node Agent.

The node agent must be started using the `asadmin start -node -agent` command on the node agent's host machine so that the server instance you are creating can be associated with that node agent.

- 5 Select a desired configuration.
 - Reference an existing configuration. No new configuration is added.
 - Make a copy of an existing configuration. A new configuration is added when the server instance or cluster is added.
- 6 To copy from a different configuration, specify it when creating the new instance.

By default, new instances are created with configurations copied from the default-config configuration.

7 For a server instance, the new configuration is named *instance-name-config*.

The configuration `default-config` is the default configuration that acts as a template for creating standalone server instance. No unclustered server instances or clusters are allowed to refer to the `default-config` configuration; it can only be copied to create new configurations. Edit the default configuration to ensure that new configurations copied from it have the correct initial settings.

More Information Equivalent `asadmin` command

`create-instance`

▼ To start an instance

1 In the tree component, expand the Stand-Alone Instances node.

2 Select the Instance you wish to start.

3 On the General tab, click Start Instance to start the instance.

The node agent associated with the instance must be started using the `asadmin start-node-agent` command before you can successfully start the instance.

Once an instance is started, it is possible to perform the following tasks from the General tab:

- **Click Stop Instance to stop the instance.**
- **Click JNDI Browsing to view the JNDI entries for that instance.**
- **Click View Log Files to view the Log Viewer and specify logging options.**
- **Click Rotate Log File.**
- **Click Recover Transactions to recover incomplete transactions.**

More Information Equivalent `asadmin` command

`start-instance`

Recovering Transactions

Transactions might be incomplete either because the server crashed or a resource manager crashed. It is essential to complete these stranded transactions and recover from the failures. Application Server is designed to recover from these failures and complete the transactions upon server startup.

If the selected server is running, then recovery will be done by the same server. If the selected server is not running, then the selected Destination Server will do the recovery.

▼ To stop an instance

- 1 In the tree component, expand the Standalone Instances node.
- 2 Select the Instance you wish to stop.
- 3 On the General tab, click Stop Instance to stop the instance.

More Information Equivalent asadmin command

```
stop-instance
```

▼ To shut down the Administration Server

- 1 In the tree component, select the Standalone Instances node.
- 2 Select the Administration Server Instance.
- 3 Click Stop.

A confirmation dialog displays to confirm that you wish to shut down the Administration Server.

Configuration Changes

- “Changing Application Server Configuration” on page 51
- “Ports in the Application Server” on page 51
- “To view port numbers” on page 52
- “To change the Administrative Server port” on page 52
- “To change an HTTP port” on page 52
- “To change an IIOP port” on page 53
- “To configure a JMX connector using the Admin Service” on page 53
- “To edit the JMX connector configuration” on page 54
- “Changing the J2SE Software” on page 55

Changing Application Server Configuration

When making any of these configuration changes, restart the server for the changes to take effect:

- Changing JVM options
- Changing port numbers
- Managing HTTP, IIOP, and JMS services
- Managing thread pools

For instructions, see [“Restarting the Server or Domain” on page 37](#).

With dynamic configuration, most changes take effect while the server is running. To make the following configuration changes, do *NOT* restart the server:

- Deploying and undeploying applications
- Adding or removing JDBC, JMS, and Connector resources and pools
- Changing logging levels
- Adding file realm users
- Changing monitoring levels
- Enabling and disabling resources and applications

Note that the `asadmin reconfig` command has been deprecated and is no longer necessary. Configuration changes are applied to the server dynamically.

Ports in the Application Server

The following table describes the port listeners of the Application Server.

TABLE 1-2 Application Server Listeners that Use Ports

Listener	Default Port Number	Description
Administrative server	4849	A domain's administrative server is accessed by the Admin Console and the <code>asadmin</code> utility. For the Admin Console, specify the port number in the URL of the browser. When executing an <code>asadmin</code> command remotely, specify the port number with the <code>--port</code> option.
HTTP	8080	The Web server listens for HTTP requests on a port. To access deployed Web applications and services, clients connect to this port.
HTTPS	8181	Web applications configured for secure communications listen on a separate port.
IIOP		Remote clients of enterprise beans (EJB components) access the beans through the IIOP listener.
IIOP, SSL		Another port is used by the IIOP listener configured for secure communications.

TABLE 1-2 Application Server Listeners that Use Ports *(Continued)*

Listener	Default Port Number	Description
IIOP, SSL and mutual authentication		Another port is used by the IIOP listener configured for mutual (client and server) authentication.

▼ To view port numbers

- 1 In the tree component, select an instance under the Standalone Instances node.
- 2 Select the Properties tab.
- 3 On the Instance Specific page, the default port numbers are identified. It is possible to set the configuration to override these values.

▼ To change the Administrative Server port

- 1 In the tree component, expand the Configurations node.
- 2 Expand the server-config (Admin Config) node.
- 3 Expand the HTTP Service node.
- 4 Expand the HTTP Listeners node.
- 5 Select the admin-listener node.
- 6 On the Edit HTTP Listener page, change the value of the Listener Port field.
- 7 Restart the server.

▼ To change an HTTP port

- 1 In the tree component, expand the HTTP Service node.
- 2 Expand the HTTP Listeners node.
- 3 Select the HTTP listener whose port number you want to change.
- 4 On the Edit HTTP Listener page, change the value of the Listener Port field.

- 5 Click Save.
- 6 Restart the server.

▼ To change an IIOP port

- 1 In the tree component, expand the Configurations node.
- 2 Expand the server-config (Admin Config) node.
- 3 Expand the ORB node.
- 4 Expand the IIOP Listeners node.
- 5 Select the listener whose port number you want to change.
- 6 On the Edit IIOP Listener page, change the value of the Listener Port field.
- 7 Click Save.
- 8 Restart the server.

▼ To configure a JMX connector using the Admin Service

Use the Admin Service to configure a JSR-160 compliant remote JMX connector, which handles communication between the domain administration server and the node agents, which manage server instances on a host machine, for remote server instances.

The Admin Service determines whether the server instance is a regular instance, a domain administration server (DAS), or a combination. A DAS is similar to a J2EE server instance, except that user applications and resources are not deployed to a DAS, though it is capable of serving user application requests. The only significant difference between a DAS and a J2EE Server Instance is that the former can not be a part of a cluster, the homogeneous unit of server instances.

- 1 Select Configurations from the tree.
- 2 Select the instance to configure:
 - a. To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.

- b. **To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.**
- 3 **Select Admin Service from the tree.**
- 4 **From the Type drop-down menu, select what you want the Admin service to configure: DAS, DAS and server, or server. Selecting DAS and server is the same as selecting DAS. The server selection selects a non-DAS server instance.**
- 5 **In the JMX Connector Name field, enter the name of the JMX connector used internally. The name of the connector is `system`.**

▼ **To edit the JMX connector configuration**

With the Edit JMX Connector screen, you can edit the configuration of the JSR-160 compliant JMX Connector.

- 1 **Select Configurations from the tree.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance’s config node. For example, for the default instance, `server`, select the `server-config` node.**
 - **To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.**
- 3 **Expand the Admin Service node and click `system`, which is the JMX connector used internally.**
- 4 **Enter the port of the JMX connector server.**

The JMX service URL is a function of the protocol, port, and address, as defined by the JSR-160 1.0 Specification.
- 5 **Enter the protocol that this JMX connector should support.**

The Application Server version 8.1 supports the `rmi_jrmp` protocol only.
- 6 **In the Realm Name field, enter the name that represents the special administrative realm.**

All authentication is handled by this realm.
- 7 **Select the Enabled checkbox to indicate that transport layer security should be used in the JMX connector.**

Changing the J2SE Software

The Application Server relies on the Java™ 2 Standard Edition (J2SE™) software. When the Application Server was installed, the directory for the J2SE software was specified. For instructions on changing the J2SE software, see “[To configure the JVM general settings](#)” on [page 317](#).

Using Online Help

The Admin Console’s online help is context-sensitive: When clicking the Help link in the upper right corner, the help browser window displays a topic related to the current Admin Console page. If the current page has no help information, the Using Online Help topic is displayed.

The online help includes conceptual topics that are not context-sensitive. To view one of these topics, select it from the table of contents in the help browser window.

▼ To go back to the previous help screen

- 1 From within the help browser window, right-click to display a selection menu.
- 2 Select Back.

Further Information

- Sun Microsystems Worldwide Training - Over 250,000 students each year are trained by Sun and its authorized centers through Web-based courses and at over 250 training sites located in more than 60 countries. For more information, see: <http://training.sun.com/>.
- *The J2EE 1.4 Tutorial* - Written for developers, the tutorial has administrative instructions for configuring JMS, setting up JavaMail resources, and managing security. To access the tutorial, go to: <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.
- *Application Server Developer’s Guide* - This guide contains development information that is specific to the Application Server.
- The `asadmin` man pages - Available in HTML format, these pages include syntax and examples for all the application server utilities including the `asadmin` utility commands.
- *Application Server Release Notes*
- docs.sun.com: Sun Product Documentation - From this site you can search for and access all of our product documentation: <http://docs.sun.com/>.
- J2EE 1.4 Documentation page - Located on our public Web site, this page has links to the technical documentation for the J2EE 1.4 platform: <http://java.sun.com/j2ee/1.4/docs/>.

- *The Quick Start Guide* - This document shows you how to deploy and run a simple Web application. The guide is in the *install-dir/docs/QuickStart.html* file.

Deploying Applications

This chapter explains how to deploy (install) J2EE applications on the Application Server. This chapter contains following sections:

- “About Deployment” on page 57
- “Admin Console Tasks for Deploying Applications” on page 60
- “Admin Console Tasks for Listing, Undeploying, and Enabling Applications” on page 72
- “Deployment Methods for Developers” on page 77

About Deployment

- “The Deployment Life Cycle” on page 57
- “Types of J2EE Archive Files” on page 59
- “Naming Conventions” on page 59

The Deployment Life Cycle

After installing the Application Server and starting a domain, you can deploy (install) J2EE applications and modules. During deployment and as the application is changed, an application or module can go through the following stages:

1. Initial Deployment

Before deploying an application or module, start the domain.

Deploy (install) an application or module to a specific stand-alone server instance or cluster. Because applications and modules are packaged in archive files, specify the archive file name during deployment. The default is to deploy to the default server instance server.

If you deploy to server instances or clusters, the application or module exists in the domain’s central repository and is referenced by any clusters or server instances you deployed to as targets.

You can also deploy to the domain using the `asadmin deploy` command, not the Admin Console. If you deploy the application or module only to the domain, it exists in the domain's central repository, but no server instances or clusters reference it until you add references, as explained in [“The Deployment Life Cycle” on page 57](#).

Deployment is dynamic: you don't need to restart the server instance after deploying application or module for applications to be available. If you do restart, all deployed applications and modules are still deployed and available.

2. Enabling or Disabling

By default, a deployed application or module is enabled, which means that it is runnable and can be accessed by clients if it has been deployed to an accessible server instance or cluster. To prevent access, disable the application or module. A disabled application or module is not uninstalled from the domain and can be easily enabled after deployment.

3. Adding or Deleting Targets for a Deployed Application or Module

Once deployed, the application or module exists in the central repository and can be referenced by a number of server instances and/or clusters. Initially, the server instances or clusters you deployed to as targets reference the application or module.

To change which server instances and clusters reference an application or module after it is deployed, change an application or module's targets using the Admin Console, or change the application references using the `asadmin` tool. Because the application itself is stored in the central repository, adding or deleting targets adds or deletes the same version of an application on different targets. However, an application deployed to more than one target can be enabled on one and disabled on another, so even if an application is referenced by a target, it is not available to users unless it is enabled on that target.

4. Redeployment

To replace a deployed application or module, redeploy it. Redeploying automatically undeploys the previously deployed application or module and replaces it with the new one.

When redeploying through the Admin Console, the redeployed application or module is deployed to the domain, and all stand-alone or clustered server instances that reference it automatically receive the new version if dynamic reconfiguration is enabled. If using the `asadmin deploy` command to redeploy, specify `domain` as the target.

For production environments, use rolling upgrade, which upgrades application without interruption in service. For more information, see [“About Rolling Upgrades.”](#)

5. Undeployment

To uninstall an application or module, undeploy it.

Types of J2EE Archive Files

A software provider packages an application or module into a archive file. To deploy the application or module, specify the archive file name. The content and structure of the archive file is defined by the specifications of the J2EE platform. Types of J2EE archive files are as follows:

- **Web Application Archive (WAR):** A WAR file consists of Web components such as servlets and JSPs, as well as static HTML pages, JAR files, tag libraries and utility classes. A WAR file name has the `.war` extension.
- **EJB JAR:** The EJB JAR file contains one or more enterprise beans, the components used for EJB technology. The EJB JAR file also includes any utility classes needed by the enterprise beans. The name of an EJB JAR file has the `.jar` extension.
- **J2EE Application Client JAR:** This JAR file contains the code for a J2EE application client, which accesses server-side components such as enterprise beans via RMI/IIOP. In the Admin Console, a J2EE application client is referred to as an “application client.” The name of the J2EE application client JAR file has the `.jar` extension.
- **Resource Adapter Archive (RAR):** A RAR file holds a resource adapter. Defined by the J2EE Connector Architecture specifications, a resource adapter is a portable component that enables enterprise beans, Web components, and application clients to access resources and foreign enterprise systems. A resource adapter is often referred to as a connector. A RAR file name has the `.rar` extension.
- **Enterprise Application Archive (EAR):** An EAR file holds one or more WAR, EJB JAR, RAR or J2EE Application Client JAR files. An EAR file name has the `.ear` extension.

The software provider might assemble an application into a single EAR file or into separate WAR, EJB JAR, and application client JAR files. In the administration tools, the deployment pages and commands are similar for all types of files.

Naming Conventions

In a given domain, the names of deployed applications and modules must be unique.

- If deploying using the Admin Console, specify the name in the Application Name field.
- If deploying using the `asadmin deploy` command, the default name of the application or module is the prefix of the JAR file that you are deploying. For example, for the `hello.war` file, the Web application name is `hello`. To override the default name, specify the `--name` option.

Modules of different types can have the same name within an application. When the application is deployed, the directories holding the individual modules are named with `_jar`, `_war` and `_rar` suffixes. Modules of the same type within an application must have unique names. In addition, database schema file names must be unique within an application.

Using a Java package-like naming scheme is recommended for module filenames, EAR filenames, module names as found in the `<module-name>` portion of the `ejb-jar.xml` files, and EJB names as found in the `<ejb-name>` portion of the `ejb-jar.xml` files. The use of this package-like naming scheme ensures that name collisions do not occur. The benefits of this naming practice apply not only to the Application Server, but to other J2EE application servers as well.

JNDI lookup names for EJB components must also be unique. Establishing a consistent naming convention might help. 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 application `pkging.ear`.

Make sure package and file names do not contain spaces or characters that are illegal for the operating system.

Admin Console Tasks for Deploying Applications

- [“To deploy an enterprise application” on page 60](#)
- [“To deploy a web application” on page 63](#)
- [“To launch a deployed web application” on page 65](#)
- [“To deploy an EJB module” on page 65](#)
- [“To deploy an application client module” on page 70](#)
- [“To deploy a connector module” on page 67](#)
- [“To create a lifecycle module” on page 68](#)
- [“To deploy an application client module” on page 70](#)

▼ To deploy an enterprise application

An enterprise application is packaged in an EAR file, a type of archive file that contains any type of J2EE stand-alone modules, such as WAR and EJB JAR files.

- 1 In the tree component, expand the Applications node.
- 2 Select the Enterprise Applications node.
- 3 On the Enterprise Applications page, click Deploy.

4 On the Deployment page, specify the location of the EAR file to deploy.

The server machine is the host that is running the application server domain administration server. The client machine is the host on which you are viewing the Admin Console through a browser.

- **If the file resides on or is accessible from the client machine, click the radio button to specify a package file to upload to the Application Server.**

Click Browse to browse to the file, or type the full path to the file.

- **If the file resides on the server machine, or to deploy an unpackaged application from an exploded directory, click the radio button to specify a package file or a directory path that must be accessible from the server.**

Type the full path name to the file or directory. Deploying from an exploded directory is for advanced developers and is not recommended for production environments.

5 Click Next to display the Deploy Enterprise Application page.**6 On the Deploy Enterprise Application page, specify the settings for the application.**

- a. **In the Application Name field, either retain the default name, which is the prefix of the file name, or type another name.**

The default name appears if you chose to upload a file. The application name must be unique.

- b. **To disable the application so that is unavailable after deployment, select the Disabled radio button.**

By default, an application is available as soon as it is deployed.

- c. **If the application has already been deployed, select the Redeploy checkbox to redeploy it; otherwise you see an error.**

You can also choose a different application name and deploy it under a new name.

- d. **To verify the structure and contents of the file before deployment, select the Verifier checkbox.**

Verification of large applications can be time-consuming. Verify the file if you suspect it is corrupt or non-portable.

- e. **To precompile JSP pages, select the JSPs checkbox.**

If you do not select this checkbox, the JSP pages are compiled at runtime when they are first accessed. Because compilation is often time-consuming, in a production environment select this checkbox.

f. Choose a high availability setting.

To enable high availability for the application, select the Availability checkbox. If availability is enabled for an application, it must also be enabled at all higher levels (named configuration and web container or EJB container) as well.

g. Choose the targets to which to deploy the application.

From the list of available targets, choose the target or targets and click Add. Targets can be clusters or stand-alone server instances. If you do not select a target, the application is deployed to the default server instance server.

If you are redeploying, don't select targets. Anything you select here is ignored. Any target clustered or stand-alone server instance that references the deployed application automatically references the new, redeployed application if dynamic reconfiguration is enabled for the cluster or stand-alone instance. For more information about how to redeploy applications without interruption of service, see "Upgrading an Application."

h. Choose whether to generate RMI stubs.

If you choose to generate RMI stubs, static RMI-IIOP stubs are generated and put into the client JAR file.

7 Click OK to deploy the application.

More Information Equivalent asadmin command
deploy

▼ To edit a deployed enterprise application

- 1 In the tree component, expand the Applications node.**
- 2 Expand the Enterprise Applications node.**
- 3 Select the node for the deployed application.**
- 4 On the Enterprise Application page, change the description.**
- 5 In Enterprise Edition, enable or disable high-availability.**

If availability is enabled for an application, it must also be enabled at all higher levels (named configuration and web container or EJB container) as well.

▼ To deploy a web application

A web application is packaged in a WAR file, a type of archive file that contains components such as servlets and JSP files.

1 In the tree component, expand the Applications node.

2 Select the Web Applications node.

3 On the Web Applications page, click Deploy.

4 On the Deployment page, specify the location of the WAR file to deploy.

The server machine is the host that is running the application server domain administration server. The client machine is the host on which you are viewing the Admin Console through a browser.

- If the file resides on or is accessible from the client machine, click the radio button to specify a package file to upload to the Application Server.

Click Browse to browse to the file, or type the full path to the file.

- If the file resides on the server machine, or to deploy an unpackaged application from an exploded directory, click the radio button to specify a package file or a directory path that must be accessible from the server.

Type the full path name to the file or directory. Deploying from an exploded directory is for advanced developers and is not recommended for production environments.

5 Click Next to display the Deploy Web Application page.

6 On the Deploy Web Application page, specify the settings for the application.

- a. In the Application Name field, either retain the default name, which is the prefix of the file name, or type another name.**

The default name appears if you chose to upload a file. The application name must be unique.

- b. In the Context Root field, enter a string that identifies the Web application.**

In the URL of the Web application, the context root immediately follows the port number (`http://host:port/context-root/...`). Make sure that the context root starts with a forward slash, for example: `/hello`.

- c. To disable the application so that is unavailable after deployment, select the Disabled radio button.**

By default, an application is available as soon as it is deployed.

- d. If the application has already been deployed, select the Redeploy checkbox to redeploy it; otherwise you see an error.**

You can also choose a different application name and deploy it under a new name.

- e. To verify the structure and contents of the file before deployment, select the Verifier checkbox.**

Verification of large applications is often time-consuming. Verify the file if you suspect it is corrupt or non-portable.

- f. To precompile JSP pages, select the JSPs checkbox.**

If you do not select this checkbox, the JSP pages are compiled at runtime when they are first accessed. Because compilation is often time-consuming, in a production environment select this checkbox.

- g. Choose a high availability setting.**

To enable high availability for the application, select the Availability checkbox. If availability is enabled for an application, it must also be enabled at all higher levels (named configuration and web container or EJB container) as well.

- h. Choose the targets to which to deploy the application.**

From the list of available targets, choose the target or targets and click Add. Targets can be clusters or stand-alone server instances. If you do not select a target, the application is deployed to the default server instance server.

If you are redeploying, don't select targets. Anything you select here is ignored. Any target clustered or stand-alone server instance that references the deployed application automatically references the new, redeployed application if dynamic reconfiguration is enabled for the cluster or stand-alone instance. For more information about how to redeploy applications without interruption of service, see "About Rolling Upgrades."

- i. Choose whether to generate RMI stubs.**

If you choose to generate RMI stubs, static RMI-IIOP stubs are generated and put into the client JAR file.

7 Click OK to deploy the application.

More Information Equivalent asadmin command

deploy

▼ To launch a deployed web application

After deploying an application, you can launch it from the Admin Console. The server and HTTP listener must be running for the application to launch.

- 1 In the tree component, expand the Applications node.
- 2 Click Web Applications.
- 3 Click the Launch link for the web application.
- 4 Click a link on the Web Application Links page to launch the application.

▼ To deploy an EJB module

An EJB module, also called an EJB JAR file, contains enterprise beans.

- 1 In the tree component, expand the Applications node.
- 2 Select the EJB Modules node.
- 3 On the EJB Module page, click Deploy.
- 4 On the Deployment page, specify the location of the JAR file to deploy.

The server machine is the host that is running the application server domain administration server. The client machine is the host on which you are viewing the Admin Console through a browser.

- If the file resides on or is accessible from the client machine, click the radio button to specify a package file to upload to the Application Server.

Click Browse to browse to the file, or type the full path to the file.

- If the file resides on the server machine, or to deploy an unpackaged application from an exploded directory, click the radio button to specify a package file or a directory path that must be accessible from the server.

Type the full path name to the file or directory. Deploying from an exploded directory is for advanced developers and is not recommended for production environments.

- 5 Click Next to display the Deploy EJB Module page.

6 On the Deploy EJB Module page, specify the settings for the module.

- a. In the Application Name field, either retain the default name, which is the prefix of the file name, or type another name.**

The default name appears if you chose to upload a file. The application name must be unique.

- b. To disable the module so that is unavailable after deployment, select the Disabled radio button.**

By default, an module is available as soon as it is deployed.

- c. If the module has already been deployed, select the Redeploy checkbox to redeploy it; otherwise you see an error.**

You can also choose a different application name and deploy it under a new name.

- d. To verify the structure and contents of the file before deployment, select the Verifier checkbox.**

Verification of large applications can be time-consuming. Verify the file if you suspect it is corrupt or non-portable.

- e. Choose a high availability setting.**

To enable high availability for the module, select the Availability checkbox. If availability is enabled for an module, it must also be enabled at all higher levels (named configuration and web container or EJB container) as well.

- f. Choose the targets to which to deploy the module.**

From the list of available targets, choose the target or targets and click Add. Targets can be clusters or stand-alone server instances. If you do not select a target, the module is deployed to the default server instance server.

If you are redeploying, don't select targets. Anything you select here is ignored. Any target clustered or stand-alone server instance that references the deployed module automatically references the new, redeployed module if dynamic reconfiguration is enabled for the cluster or stand-alone instance. For more information about how to redeploy modules without interruption of service, see "About Rolling Upgrades."

- g. Choose whether to generate RMI stubs.**

If you choose to generate RMI stubs, static RMI-IIOP stubs are generated and put into the client JAR file.

7 Click OK to deploy the module.

More Information Equivalent asadmin command
deploy

▼ To deploy a connector module

A connector, also known as a resource adapter, is packaged in a type of archive file called a RAR file.

1 In the tree component, expand the Applications node.

2 Select the Connector Modules node.

3 On the Connector Modules page, click Deploy.

4 On the Deployment page, specify the location of the RAR file to deploy.

The server machine is the host that is running the application server domain administration server. The client machine is the host on which you are viewing the Admin Console through a browser.

- If the file resides on or is accessible from the client machine, click the radio button to specify a package file to upload to the Application Server.

Click Browse to browse to the file, or type the full path to the file.

- If the file resides on the server machine, or to deploy an unpackaged module from an exploded directory, click the radio button to specify a package file or a directory path that must be accessible from the server.

Type the full path name to the file or directory. Deploying from an exploded directory is for advanced developers and is not recommended for production environments.

5 Click Next to display the Deploy Connector Module page.

6 On the Deploy Connector Module page, specify the settings for the module.

a. In the Application Name field, either retain the default name, which is the prefix of the file name, or type another name.

The default name appears if you chose to upload a file. The application name must be unique.

b. In the Thread Pool Id field, specify the thread pool for the resource adapter you are deploying.

By default, the Application Server services work requests from all resource adapters from its default thread pool. Use this field to associate a specific user-created thread pool to service work requests from a resource adapter.

- c. **To disable the module so that is unavailable after deployment, select the Disabled radio button.**

By default, a module is available as soon as it is deployed.

When you enable or disable a connector module, you also enable or disable the connector resources and connection pools that point to the module.

- d. **If the module has already been deployed, select the Redeploy checkbox to redeploy it; otherwise you see an error.**

You can also choose a different application name and deploy it under a new name.

- e. **To verify the structure and contents of the file before deployment, select the Verifier checkbox.**

Verification of large applications is often time-consuming. Verify the file if you suspect it is corrupt or non-portable.

- f. **If the resource adapter has additional properties specified, they are displayed.**

Use the table to modify the default values of these properties.

- g. **Choose the targets to which to deploy the module.**

From the list of available targets, choose the target or targets and click Add. Targets can be clusters or stand-alone server instances. If you do not select a target, the module is deployed to the default server instance server.

If you are redeploying, don't select targets. Anything you select here is ignored. Any target clustered or stand-alone server instance that references the deployed module automatically references the new, redeployed module if dynamic reconfiguration is enabled for the cluster or stand-alone instance. For more information about how to redeploy modules without interruption of service, see "About Rolling Upgrades."

7 Click OK to deploy the module.

More Information Equivalent asadmin command

deploy

▼ To create a lifecycle module

A lifecycle module performs tasks when it is triggered by one or more events in the server lifecycle. These server events are:

- Initialization
- Startup

- Ready to service requests
- Shutdown

Lifecycle modules are not part of the J2EE specification, but are an enhancement to the Application Server.

- 1 In the tree component, expand the Applications node.**
- 2 Select the Lifecycle Modules node.**
- 3 On the Lifecycle Modules page, click New.**
- 4 On the Create Lifecycle Module page, specify these settings:**
 - a. In the Name field, type a name that denotes the function of the module.**
 - b. In the Class Name field, type the fully qualified name of the lifecycle module's class file.**
 - c. If the JAR file containing the lifecycle is in the server's classpath, leave the Classpath field blank. Otherwise, type the fully qualified path.**

If you don't specify the classpath, you must unpack the classes in *domain-dir/applications/lifecycle-module/module-name*. If you specify a classpath, nothing else is required.
 - d. In the Load Order field, type an integer greater than 100 and less than the operating system's MAXINT value.**

The integer determines the order in which lifecycle modules are loaded when the server starts up. Modules with smaller integers are loaded sooner.
 - e. When you start the server, it loads lifecycle modules that are already deployed.**

By default, if a load fails, the server continues the startup operation. To prevent the server from starting up when a load fails, select the On Load Failure checkbox.
 - f. To disable the module so that is unavailable after deployment, select the Disabled radio button.**

Because a lifecycle module is invoked at server startup, a disabled lifecycle module won't cease to have an effect until the server instance is restarted.
 - g. Choose the targets to which to deploy the module.**

From the list of available targets, choose the target or targets and click Add. Targets can be clusters or stand-alone server instances. If you do not select a target, the module is deployed to the default server instance server.

- 5 Click OK.

More Information Equivalent `asadmin` command

```
create-lifecycle-module
```

▼ To deploy an application client module

An application client module, also called a J2EE application client JAR file, contains the server-side routines for the client.

- 1 In the tree component, expand the **Applications** node.

- 2 Select the **App Client Modules** node.

- 3 On the **Application Client Modules** page, click **Deploy**.

- 4 On the **Deployment** page, specify the location of the JAR file to deploy.

The server machine is the host that is running the application server domain administration server. The client machine is the host on which you are viewing the Admin Console through a browser.

- If the file resides on or is accessible from the client machine, click the radio button to specify a package file to upload to the Application Server.

Click **Browse** to browse to the file, or type the full path to the file.

- If the file resides on the server machine, or to deploy an unpackaged module from an exploded directory, click the radio button to specify a package file or a directory path that must be accessible from the server.

Type the full path name to the file or directory. Deploying from an exploded directory is for advanced developers and is not recommended for production environments.

- 5 Click **Next** to display the **Deploy Application Client Module** page.

- 6 On the **Deploy Application Client Module** page, specify the settings for the module.

- a. In the **Application Name** field, either retain the default name, which is the prefix of the file name, or type another name.

The default name appears if you chose to upload a file. The application name must be unique.

- b. If the module has already been deployed, select the Redeploy checkbox to redeploy it; otherwise you see an error.**

You can also choose a different application name and deploy it under a new name.

- c. To verify the structure and contents of the file before deployment, select the Verifier checkbox.**

Verification of large applications can be time-consuming. Verify the file if you suspect it is corrupt or non-portable.

- d. Choose the targets to which to deploy the module.**

From the list of available targets, choose the target or targets and click Add. Targets can be clusters or stand-alone server instances. If you do not select a target, the module is deployed to the default server instance server.

If you are redeploying, don't select targets. Anything you select here is ignored. Any target clustered or stand-alone server instance that references the deployed module automatically references the new, redeployed module if dynamic reconfiguration is enabled for the cluster or stand-alone instance. For more information about how to redeploy modules without interruption of service, see "About Rolling Upgrades."

- e. Choose whether to generate RMI stubs.**

If you choose to generate RMI stubs, static RMI-IIOP stubs are generated and put into the client JAR file.

For the client-side routines:

- Typically, the application provider ships a JAR file containing the client-side routines.
- The application provider gets the client-side stubs by specifying the `--retrieve` option of the `asadmin deploy` command.

- 7 Click OK to deploy the module.**

More Information Equivalent `asadmin` command

`deploy`

▼ To specify an archive file for deployment

Click Deploy on an application or module page, to access the Deployment page. On the Deployment page, specify the location of the archive file in which the application or module is packaged.

The server machine is the host that is running the application server domain administration server. The client machine is the host on which you are viewing the Admin Console through a browser.

- 1 **If the file resides on or is accessible from the client machine, click the radio button to specify a package file to upload to the Application Server.**

Click Browse to browse to the file, or type the full path to the file.

- 2 **If the file resides on the server machine, or to deploy an unpackaged application from an exploded directory, click the radio button to specify a package file or a directory path that must be accessible from the server.**

Type the full path name to the file or directory. Deploying from an exploded directory is for advanced developers and is not recommended for production environments.

Admin Console Tasks for Listing, Undeploying, and Enabling Applications

- [“To list deployed applications” on page 72](#)
- [“To list subcomponents” on page 73](#)
- [“To view module descriptors of deployed applications” on page 73](#)
- [“To undeploy an application or module” on page 73](#)
- [“To enable or disable an application or module” on page 74](#)
- [“To configure dynamic reloading” on page 76](#)

▼ To list deployed applications

- 1 **In the tree component, expand the Applications node.**
- 2 **Expand the node for the application or module type.**

To view the details of a deployed application or module either:

- In the tree component, select the node of the application or module.
- On the page, select an entry in the Application Name column.

More Information Equivalent asadmin command

`list-components`

▼ To list subcomponents

Enterprise and Web applications, EJB Modules and Connector Modules contain subcomponents. For example, a web application might contain one or more servlets.

- 1 In the tree component, expand the Applications node.
- 2 Expand the node for the type of application or module for which to view descriptors.
- 3 Select the node for the deployed application or module.
- 4 On the Application or Module page, note the contents of the Sub Components table.

More Information Equivalent asadmin command

```
list-sub-components
```

▼ To view module descriptors of deployed applications

For Enterprise Applications, Web Applications, EJB Modules, Connector Modules, and App Client Modules, you can view the module deployment descriptors.

- 1 In the tree component, expand the Applications node.
- 2 Select the node for the type of application or module for which to view descriptors.
- 3 Select the node for the deployed application or module.
- 4 Select the Descriptor tab.
- 5 To see the text of the descriptor file, click the file name.

The page displays the file contents. This information is read-only.

▼ To undeploy an application or module

Undeploying an application or module uninstalls it from the domain and removes references to it from all instances.

- 1 In the tree component, expand the Applications node.
- 2 Select the node for the type of application or module want to undeploy.

- 3 In the table listing the deployed applications, select the checkbox for the application or module you want to undeploy.
- 4 Click **Undeploy**.

More Information Equivalent `asadmin` command
`undeploy`

▼ To enable or disable an application or module

If a deployed application or module is enabled, it is accessible by clients. If it is disabled, it is still deployed but is not accessible by clients. By default, when you deploy an application or module, it is enabled because the **Enable on All Targets** radio button is selected by default.

- 1 In the tree component, expand the **Applications** node.
- 2 Expand the node for the application type.
- 3 To enable a deployed application or module, select the checkbox next to the deployed application or module.
To enable an application on a single target:
 - a. Select the node for the application.
 - b. Click the **Targets** tab.
 - c. Select the checkbox next to the deployed application or module.
- 4 Click **Enable** or **Disable**.
These buttons enable or disable the application on all targets.

More Information Equivalent `asadmin` commands
`enable` and `disable`

▼ To manage application targets

After deploying an application or module, manage the server instances and clusters that reference it by managing targets.

- 1 In the tree component, expand the **Applications** node.

- 2 Expand the node for the application type.
- 3 Select the node for the deployed application.
- 4 Select the Targets tab.
- 5 To enable or disable an application on a specific target instance or cluster, click the checkbox next to the target and click Enable or Disable.
- 6 To add or delete targets for the application, choose Manage Targets.
- 7 Add or remove targets and click OK.
The application is now available on the revised list of targets.

More Information Equivalent asadmin commands

create-application-ref and delete-application-ref

▼ To deploy on additional virtual servers

After deploying an application or module to a target server instances or clusters, you can associate it with additional virtual servers.

- 1 From the deployed application or module's Target page, click the Manage Virtual Servers link next to the target.
- 2 Add or remove virtual server targets from the list of available virtual servers.
- 3 Click OK.

Redeploying to Multiple Targets

If an application is deployed to multiple targets (stand-alone server instances or clusters), you have two ways of redeploying to multiple targets. Use one of the following methods to make sure all server instances that reference an application receive the most recent version.

Development Environment

In a development environment, simply redeploy the application. The application is redeployed to the domain, and all targets that reference it automatically receive the new version if dynamic reconfiguration is enabled for the target server instances. Dynamic reconfiguration is enabled by default. If dynamic reconfiguration is not enabled for a server instance, it continues to use the old version until the server instance is restarted.

Production Environment

In a production environment, follow the steps detailed in “About Rolling Upgrades.”

▼ To configure dynamic reloading

If dynamic reloading is enabled, the server periodically checks for changes in a deployed application and automatically reloads the application with the changes. Changes are signaled by a date change to a file called `.reload` that you create manually. The applications must be installed in `domain-dir/applications/j2ee-modules/module-name` or `domain-dir/applications/j2ee-apps/app-name`

For example:

```
/opt/SUNWappserver/domain/domain1/applications/j2ee-modules/webapps-simple
```

Dynamic reloading is useful in a development environment because it allows code changes to be tested quickly. In a production environment, however, dynamic reloading may degrade performance.

Note – Dynamic reloading is only available for the default server instance.

Dynamic reloading is intended for development environments. It is incompatible with session persistence, a production environment feature. Do not enable session persistence if dynamic reloading is enabled.

- 1 In the tree component, expand the **Stand-Alone Instances** node.
- 2 Click **server (Admin Server)**.
- 3 Click **Advanced**.
- 4 On the **Applications Configuration** page, configure the following:
 - **Reload:** Enable or disable dynamic reloading with the **Enabled** checkbox.
 - **Reload Poll Interval:** Specify how often the server checks for changes in the deployed applications.
 - **Admin Session Timeout:** Specify the amount of time before the Admin Session times out and you have to log in again.

Next Steps After configuring the system to use dynamic reloading, for every application to be reloaded dynamically, create a file called `.reload` and put it in the application’s directory. The file does not have any content. When you change the application, change the date of the file (for example, using the UNIX `touch` command), and the changes are reloaded automatically.

Deployment Methods for Developers

See Also:

- “To use auto deploy” on page 77
- “To use auto deploy” on page 77
- “To deploy an unpackaged application from a directory” on page 78
- “Using the deploytool Utility” on page 79
- “Using a Deployment Plan” on page 79

▼ To use auto deploy

The auto deploy feature enables you to deploy a prepackaged application or module by copying it to the *domain-dir/autodeploy* directory.

For example, copy a file named `hello.war` to the *domain-dir/autodeploy* directory. To undeploy the application, remove the `hello.war` file from the *autodeploy* directory.

You can also use the Admin Console or `asadmin` tool to undeploy the application. In that case, the archive file remains.

Note – Auto deploy is only available for the default server instance.

The auto deploy feature is intended for development environments. It is incompatible with session persistence, a production environment feature. Do not enable session persistence if auto deployment is enabled

- 1 In the tree component, expand the Stand-Alone Instances node.
- 2 Click server (Admin Server).
- 3 Click Advanced.
- 4 On the Applications Configuration page, configure the following:
 - a. Enable or disable auto deploy by selecting or deselecting the Enabled checkbox.
 - b. In the Auto Deploy Poll Interval field, specify how often the server checks the auto deploy directory for application or module files.

Changing the poll interval does not affect the amount of time it takes to deploy an application or module.

- c. **In the Auto Deploy Directory, if you specify the directory where you build your application, then you won't have to copy the file to the default auto deploy directory.**

The default is a directory called `autodeploy` in the server instance's root directory.

By default a variable is used to eliminate the need to manually change the directory for multiple server instances. For more information on these variables, see [“To set the advanced domain attributes” on page 321](#).

- d. **To run the verifier before deployment, select the Verifier.**

The verifier examines the structure and content of the file. Verification of large applications is often time-consuming.

- e. **To precompile JSP pages, select Precompile.**

If you do not select this checkbox, the JSP pages are compiled at runtime when they are first accessed. Because compilation is often time-consuming, in a production environment select this checkbox.

▼ **To deploy an unpackaged application from a directory**

This feature is for advanced developers.

Use directory deployment only to deploy to the default server instance (server). You cannot use it to deploy to clusters or stand-alone server instances.

A directory containing an unpackaged application or module is sometimes called an exploded directory. The contents of the directory must match the contents of a corresponding J2EE archive file. For example, if you deploy a Web application from a directory, the contents of the directory must be the same as a corresponding WAR file. For information about the required directory contents, see the appropriate specifications.

You can change the deployment descriptor files directly in the exploded directory.

If your environment is configured to use dynamic reloading, you can also dynamically reload applications deployed from the directory. For more information, see [“To configure dynamic reloading” on page 76](#).

- 1 **In the Admin Console, begin the deployment process. See [“To deploy a web application” on page 63](#).**
- 2 **On the Deployment page, specify the following:**
 - a. **Click the radio button to specify a package file or a directory path that must be accessible from the server.**
 - b. **In the File Or Directory field, enter the name of the exploded directory.**

More Information Equivalent `asadmin` command

```
deploydir
```

Using the `deploytool` Utility

Designed for software developers, the `deploytool` utility packages and deploys J2EE applications and modules. For instructions on how to use `deploytool`, see *The J2EE 1.4 Tutorial*.

Using a Deployment Plan

This feature is for advanced developers.

A deployment plan is an JAR file that contains only the deployment descriptors that are specific to the Application Server. These deployment descriptors, for example `sun-application.xml`, are described in the *Application Server Developer's Guide*. The deployment plan is part of the implementation of *JSR 88: J2EE Application Deployment*. Use a deployment plan to deploy an application or module that does not contain the deployment descriptors that are specific to the Application Server.

To deploy using a deployment plan, specify the `--deploymentplan` option of the `asadmin` `deploy` command. The following command, for example, deploys the enterprise application in the `myrosterapp.ear` file according to the plan specified by the `mydeployplan.jar` file.

```
$ asadmin deploy --user admin ---deploymentplan mydeployplan.jar myrosterapp.ear
```

In the deployment plan file for an enterprise application (EAR), the `sun-application.xml` file is located at the root. The deployment descriptor for each module is stored according to this syntax: `module-name.sun-dd-name`, where the `sun-dd-name` depends on the module type. If a module contains a CMP mappings file, the file is named `module-name.sun-cmp-mappings.xml`. A `.dbschema` file is stored at the root level with each forward slash character (`/`) replaced by a pound sign (`#`). The following listing shows the structure of the deployment plan file for an enterprise application (EAR).

```
$ jar -tvf mydeployplan.jar
420 Thu Mar 13 15:37:48 PST 2003 sun-application.xml
370 Thu Mar 13 15:37:48 PST 2003 RosterClient.war.sun-web.xml
418 Thu Mar 13 15:37:48 PST 2003 roster-ac.jar.sun-application-client.xml
1281 Thu Mar 13 15:37:48 PST 2003 roster-ejb.jar.sun-ejb-jar.xml
2317 Thu Mar 13 15:37:48 PST 2003 team-ejb.jar.sun-ejb-jar.xml
3432 Thu Mar 13 15:37:48 PST 2003 team-ejb.jar.sun-cmp-mappings.xml
84805 Thu Mar 13 15:37:48 PST 2003 team-ejb.jar.RosterSchema.dbschema
```

In the deployment plan for a web application or a module file, the deployment descriptor that is specific to the Application Server is at the root level. If a stand-alone EJB module contains a CMP bean, the deployment plan includes the `sun-cmp-mappings.xml` and `.dbschema` files at the root level. In the following listing, the deployment plan describes a CMP bean.

```
$ jar r -tvf myotherplan.jar
3603 Thu Mar 13 15:24:20 PST 2003 sun-ejb-jar.xml
3432 Thu Mar 13 15:24:20 PST 2003 sun-cmp-mappings.xml
84805 Thu Mar 13 15:24:20 PST 2003 RosterSchema.dbschema
```


JDBC Resources

This chapter explains how to configure JDBC resources, which are required by applications that access databases. This chapter contains the following sections:

- “About JDBC Resources and Connection Pools” on page 81
- “Setting Up Database Access” on page 83
- “About JDBC Connection Pools” on page 84
- “About JDBC Resources” on page 88
- “About Persistence Manager Resources” on page 91

About JDBC Resources and Connection Pools

- “JDBC Resources” on page 81
- “JDBC Connection Pools” on page 82
- “How JDBC Resources and Connection Pools Work Together” on page 82

JDBC Resources

To store, organize, and retrieve data, most applications use relational databases. J2EE applications access relational databases through the JDBC API.

A JDBC resource (data source) provides applications with a means of connecting to a database. Typically, the administrator creates a JDBC resource for each database accessed by the applications deployed in a domain. (However, more than one JDBC resource can be created for a database.)

To create a JDBC resource, specify a unique JNDI name that identifies the resource. (See the section JNDI Names and Resources.) Expect to find the JNDI name of a JDBC resource in `java:comp/env/jdbc` subcontext. For example, the JNDI name for the resource of a payroll database could be `java:comp/env/jdbc/payrolldb`. Because all resource JNDI names are in

the `java:comp/env` subcontext, when specifying the JNDI name of a JDBC resource in the Admin Console, enter only `jdbc/name`. For example, for a payroll database specify `jdbc/payrolldb`.

JDBC Connection Pools

To create a JDBC resource, specify the connection pool with which it is associated. Multiple JDBC resources can specify a single connection pool.

A JDBC connection pool is a group of reusable connections for a particular database. Because creating each new physical connection is time consuming, the server maintains a pool of available connections to increase performance. When an application requests a connection, it obtains one from the pool. When an application closes a connection, the connection is returned to the pool.

The properties of connection pools can vary with different database vendors. Some common properties are the database's name (URL), user name, and password.

How JDBC Resources and Connection Pools Work Together

To store, organize, and retrieve data, most applications use relational databases. J2EE applications access relational databases through the JDBC API. Before an application can access a database, it must get a connection.

At runtime, here's what happens when an application connects to a database:

1. The application gets the JDBC resource (data source) associated with the database by making a call through the JNDI API.

Given the resource's JNDI name, the naming and directory service locates the JDBC resource. Each JDBC resource specifies a connection pool.

2. Via the JDBC resource, the application gets a database connection.

Behind the scenes, the application server retrieves a physical connection from the connection pool that corresponds to the database. The pool defines connection attributes such as the database name (URL), user name, and password.

3. Now that it's connected to the database, the application can read, modify, and add data to the database.

The applications access the database by making calls to the JDBC API. The JDBC driver translates the application's JDBC calls into the protocol of the database server.

4. When it's finished accessing the database, the application closes the connection.

The application server returns the connection to the connection pool. Once it's back in the pool, the connection is available for the next application.

Setting Up Database Access

- [“To Set Up Database Access” on page 83](#)
- [“To integrate a JDBC driver” on page 83](#)

▼ To Set Up Database Access

1 Install a supported database product.

For a list of database products supported by the Application Server, see the *Release Notes*.

2 Install a JDBC driver for the database product.

3 Make the driver’s JAR file accessible to the domain’s server instance. See [“To integrate a JDBC driver” on page 83](#).

4 Create the database.

Usually, the application provider delivers scripts for creating and populating the database.

5 Create a connection pool for the database. See [“To create a JDBC connection pool” on page 84](#).

6 Create a JDBC resource that points to the connection pool. See [“To create a JDBC resource” on page 88](#).

▼ To integrate a JDBC driver

A JDBC driver translates an application’s JDBC calls into the protocol of the database server.

- To integrate the JDBC driver into an administrative domain, do either of the following:
 - Make the driver accessible to the common class loader.
 - a. Copy the driver’s JAR and ZIP files into the `domain-dir/lib` directory or copy its class files into the `domain-dir/lib/ext` directory.
 - b. Restart the domain.
 - Make the driver accessible to the system class loader.
 - a. In the Admin Console’s tree view (left pane), select Configurations.
 - b. Select the desired configuration (for example, default-config).

- c. Select JVM Settings.
- d. On the JVM Settings page, click the Path Settings tab.
- e. In the Classpath Suffix field, enter the fully-qualified path name for the driver's JAR file.
- f. Click Save.
- g. Restart the server.

About JDBC Connection Pools

- “To create a JDBC connection pool” on page 84
- “To edit a JDBC connection pool” on page 85
- “To verify connection pool settings” on page 88
- “To delete a JDBC connection pool” on page 88

▼ To create a JDBC connection pool

A JDBC connection pool is a group of reusable connections for a particular database. When creating the pool with the Admin Console, the Administrator is actually defining the aspects of a connection to a specific database.

Before You Begin Before creating the pool, you must first install and integrate the JDBC driver. See [“To integrate a JDBC driver” on page 83](#).

When building the Create Connection Pool pages, certain data specific to the JDBC driver and the database vendor must be entered. Before proceeding, gather the following information:

- Database vendor name
 - Resource type, such as `javax.sql.DataSource` (local transactions only)
`javax.sql.XADataSource` (global transactions)
 - Data source class name
 - Required properties, such as the database name (URL), user name, and password
- 1 In the Admin Console's tree view (left pane), expand the Resources node.
 - 2 Under Resources, expand the JDBC node.
 - 3 Under JDBC, select the Connection Pools node.
 - 4 On the Connection Pools page, click New.

- 5 On the first Create Connection Pool page, specify the following general settings:
 - a. In the Name field, enter a logical name for the pool.
Specify this name when creating a JDBC resource.
 - b. Select an entry from the Resource Type combo box.
 - c. Select an entry from the Database Vendor combo box.
- 6 Click Next.
- 7 On the second Create Connection Pool page, specify the value for the DataSource Class Name field.
If the JDBC driver has a DataSource class for the resource type and database vendor specified in the previous page, then the value of the DataSource Class Name field is provided.
- 8 Click Next.
- 9 On the third and last Create Connection Pool page, perform these tasks:
 - a. In the General Settings section, verify that the values are correct.
 - b. For the fields in the Pool Settings, Connection Validation, and Transaction Isolation sections, retain the default values.
It is most convenient to change these settings at a later time. See [“To edit a JDBC connection pool” on page 85](#).
 - c. In the Additional Properties table, add the required properties, such as database name (URL), user name, and password.
- 10 Click Finish.

More Information Equivalent asadmin command

```
create-jdbc-connection-pool
```

▼ To edit a JDBC connection pool

The Edit JDBC Connection Pool page provides the means to change all of the settings for an existing pool except its name.

- 1 In the tree component, expand the Resources node.

- 2 Under Resources, expand the JDBC node.
- 3 Under JDBC, expand the Connection Pools node.
- 4 Select the node for the pool you want to edit.
- 5 On the Edit JDBC Connection Pool page, make any necessary changes.

a. Change general settings.

The values of the general settings depend on the specific JDBC driver that is installed. These settings are the names of classes or interfaces in the Java programming language.

Parameter	Description
DataSource Class Name	The vendor-specific class name that implements the DataSource and / or XADataSource APIs. This class must be in the classpath.
Resource Type	Choices include javax.sql.DataSource (local transactions only), javax.sql.XADataSource (global transactions), and javax.sql.ConnectionPoolDataSource (local transactions, possible performance improvements).

b. Change pool settings.

A set of physical database connections reside in the pool. When an application requests a connection, the connection is removed from the pool, and when the application releases the connection, it is returned to the pool.

Parameter	Description
Initial and Minimum Pool Size	The minimum number of connections in the pool. This value also determines the number of connections that are created when the pool is first created or when application server starts.
Maximum Pool Size	The maximum number of connections in the pool.
Pool Resize Quantity	When the pool shrinks toward the minimum pool size it is resized in batches. This value determines the number of connections in the batch. Making this value too large delays connection recycling; making it too small will be less efficient.
Idle Timeout	The maximum time in seconds that a connection can remain idle in the pool. After this time expires, the connection is removed from the pool.
Max Wait Time	The amount of time the application requesting a connection will wait before getting a connection timeout. If the wait time is long, the application might appear to hang indefinitely.

c. Change connection validation settings.

Optionally, the application server can validate connections before they are passed to applications. This validation allows the application server to automatically reestablish database connections if the database becomes unavailable due to network failure or database server crash. Validation of connections incurs additional overhead and slightly reduces performance.

Parameter	Description
Connection Validation	Select the Required checkbox to enable connection validation.
Validation Method	<p>The application server can validate database connections in three ways: auto-commit, metadata, and table-based. auto-commit and metadata - The application server validates a connection by calling the <code>con.getAutoCommit()</code> and <code>con.getMetaData()</code> methods. However, because many JDBC drivers cache the results of these calls, you should check with the driver vendor to determine whether these calls are cached or not.</p> <p>auto-commit validation makes use of two methods for validating the connection: <code>con.getAutoCommit()</code> and <code>con.setAutoCommit()</code> to change the state of auto-commit. This allows actual connection validation. <code>con.getAutoCommit()</code> might or might not contact the database, depending on the implementation. <code>con.setAutoCommit()</code> is wrapped for different purposes, such as a connection pool.</p> <p>Note – Databases such as Oracle do some caching for <code>con.setAutoCommit()</code>. For such databases, actual validation may not happen, so table-based validation is recommended.</p> <p>table - The application queries a database table that are specified. The table must exist and be accessible. Do not use an existing table that has a large number of rows or a table that is already frequently accessed.</p>
Table Name	If you selected table from the Validation Method combo box, then specify the name of the database table.
On Any Failure	If you select the checkbox labelled Close All Connections, if a single connection fails, then the application closes all connections in the pool and then reestablish them. If you do not select the checkbox, then individual connections are used.

d. Change transaction isolation settings.

Because a database is usually accessed by many users concurrently, one transaction might update data while another attempts to read the same data. The isolation level of a transaction defines the degree to which the data being updated is visible to other transactions. For details on isolation levels, refer to the documentation of the database vendor.

Parameter	Description
Transaction Isolation	Makes it possible to select the transaction isolation level for the connections of this pool. If left blank, the application uses the default isolation levels provided by the JDBC driver.
Guaranteed Isolation Level	Only applicable if the isolation level has been specified. If you select the Guaranteed checkbox, all connections in the pool have the same isolation level. For example, if the isolation level for the connection is changed (using <code>con.setTransactionIsolation()</code>) when last used, this mechanism changes the status back to the Guaranteed level.

e. Change properties.

In the Additional Properties table, it is possible to specify properties, such as the database name (URL), user name, and password. Because the properties vary with database vendor, consult the vendor's documentation for details.

6 Click Save.

- See Also**
- [“To create a JDBC connection pool” on page 84](#)
 - [“To verify connection pool settings” on page 88](#)
 - [“To delete a JDBC connection pool” on page 88](#)

▼ To verify connection pool settings

1 Start the database server.

2 Click Ping.

The Admin Console attempts to connect to the database. If an error message displays, check to see if the database server was restarted.

▼ To delete a JDBC connection pool

1 In the tree component, expand the Resources node.

2 Under Resources, expand the JDBC node.

3 Under JDBC, select the Connection Pools node.

4 On the Connection Pools page, select the checkbox for the pool to be deleted.

5 Click Delete.

More Information Equivalent asadmin command

```
delete-jdbc-connection-pool
```

About JDBC Resources

- [“To create a JDBC resource” on page 88](#)
- [“To edit a JDBC resource” on page 89](#)
- [“To delete a JDBC resource” on page 90](#)
- [“To enable or disable a JDBC resource” on page 90](#)

▼ To create a JDBC resource

A JDBC resource (data source) provides applications with a means of connecting to a database.

Before You Begin Before creating a JDBC resource, first create a JDBC connection pool. See [“To create a JDBC connection pool”](#) on page 84.

- 1 In the tree component, expand the Resources node.
- 2 Under Resources, expand the JDBC node.
- 3 Under JDBC, select the JDBC Resources node.
- 4 On the JDBC Resources page, click New.
- 5 On the Create JDBC Resource page, specify the resource’s settings:
 - a. In the JNDI Name field, type a unique name.
By convention, the name begins with the jdbc/ string. For example: jdbc/payrolldb. Don’t forget the forward slash.
 - b. From the Pool Name combo box, choose the connection pool to be associated with the new JDBC resource.
 - c. If you want the resource to be unavailable, deselect the Enabled checkbox.
By default, the resource is available (enabled) as soon as it is created.
 - d. In the Description field, type a short description of the resource.
 - e. In the Targets section, specify the targets (clusters and standalone server instances) on which the resource is available.
Select the desired target on the left, and click Add to add it to the list of selected targets.
- 6 Click OK.

More Information Equivalent asadmin command
`create-jdbc-resource`

▼ To edit a JDBC resource

- 1 In the tree component, expand the Resources node.
- 2 Under Resources, expand the JDBC node.
- 3 Under JDBC, expand the JDBC Resources node.

- 4 Select the node for the JDBC resource to be edited.
- 5 On the Edit JDBC Resource page, it is possible to perform these tasks:
 - a. From the Pool Name combo box, select a different connection pool.
 - b. In the Description field, change the short description of the resource.
 - c. Select or deselect the checkbox to enable or disable the resource.
 - d. Select the Targets tab to change the targets (clusters and standalone server instances) on which the resource is available.

Select the checkbox for an existing target in the list, then click Enable to enable the resource for that target or Disable to disable the resource for that target.

Click Manage Targets to add or remove targets to the list. In the Manage Targets page, select the desired target in the Available list on the left, and click Add to add it to the list of selected targets. Click Remove to remove a target from the Selected list.

Click OK to save the changes to the available targets.
- 6 Click Save to apply the edits.

▼ To delete a JDBC resource

- 1 In the tree component, expand the Resources node.
- 2 Under Resources, expand the JDBC node.
- 3 Under JDBC, select the Connection Pools node.
- 4 On the Connection Pools page, select the checkbox for the pool to be deleted.
- 5 Click Delete.

▼ To enable or disable a JDBC resource

- 1 In the tree component, expand the JDBC Resources node or expand the Standalone Instances to select the Server Instance node Resource tab.
- 2 On the Resources page, select the checkbox for the resource to be enabled or disabled.

- 3 Click Enable or Disable.

About Persistence Manager Resources

- “To create a persistence manager resource” on page 91
- “To edit a persistence manager resource” on page 92
- “To manage resource targets” on page 92
- “To delete a persistence manager resource” on page 92
- “To enable or Disable a persistence manager resource” on page 93

▼ To create a persistence manager resource

This feature is needed for backward compatibility. To run on version 7 of the Application Server, a persistent manager resource was required for applications with container-managed persistence beans (a type of EJB component). Using a JDBC resource instead is recommended.

- 1 In the tree component, expand the Resources node.
- 2 Under Resources, select the Persistence Managers node.
- 3 On the Persistence Managers page, click New.
- 4 On the Create Persistence Manager page, specify these settings:
 - a. In the JNDI Name field, type a unique name.
For example: jdo/mypm. Don't forget the forward slash.
 - b. In the Factory Class field, retain the default class provided with this release, or type in the class of another implementation.
 - c. From the Connection Pool combo box, choose the connection pool that the new persistence manager resource will belong to.
 - d. To disable the new persistence manager resource, deselect the Enabled check box.
By default, the new persistence manager resource is enabled.
 - e. In the Targets section, specify the targets (clusters and standalone server instances) on which the resource is available.
Select the desired target on the left, and click Add to add it to the list of selected targets.
- 5 Click OK.

More Information Equivalent asadmin command

`create-persistence-resource`

▼ To edit a persistence manager resource

- 1 From the **Edit Persistence Manager Properties** tab, select the **Add Property** button.
A new row is added to the **Additional Properties** table.
- 2 Add the desired property and value.

▼ To manage resource targets

- 1 Select the **Targets** tab to change the targets (clusters and standalone server instances) where the resource resides.
- 2 Select the checkbox for an existing target in the list, then click **Enable** to enable the resource for that target or **Disable** to disable the resource for that target.
- 3 Click **Manage Targets** to add or remove targets to the list.
In the **Manage Targets** page, select the desired target in the **Available** list on the left, and click **Add** to add it to the list of selected targets. Click **Remove** to remove a target from the **Selected** list.
- 4 Click **OK** to save the changes to the available targets.
- 5 Click **Save**.

▼ To delete a persistence manager resource

- 1 In the tree component, expand the **Persistence Managers** node.
- 2 Select the **Persistence Managers** node.
- 3 On the **Persistence Managers** page, select the checkbox for the persistence manager that you want to delete.
- 4 Click **Delete**.

More Information Equivalent asadmin command
`delete-persistence-resource`

▼ **To enable or Disable a persistence manager resource**

- 1 In the tree component, expand the Persistence Managers node.
- 2 Select the checkbox for the resource to be enabled or disabled.
- 3 Click Enable or Disable.

Configuring Java Message Service Resources

This chapter describes how to configure resources for applications that use the Java Message Service (JMS) API. It contains the following sections:

- “About JMS Resources” on page 95
- “Admin Console Tasks for JMS Connection Factories” on page 97
- “Admin Console Tasks for JMS Destination Resources” on page 101
- “Admin Console Tasks for JMS Physical Destinations” on page 103
- “Admin Console Tasks for the JMS Provider” on page 105

About JMS Resources

- “The JMS Provider in the Application Server” on page 95
- “JMS Resources” on page 95
- “The Relationship Between JMS Resources and Connector Resources” on page 97

The JMS Provider in the Application Server

The Application Server implements the Java Message Service (JMS) API by integrating the Sun Java System Message Queue (formerly Sun ONE Message Queue) software into the Application Server. For basic JMS API administration tasks, use the Application Server Admin Console. For advanced tasks, including administering a Message Queue cluster, use the tools provided in the *MQ-install-dir/imq/bin* directory.

For details about administering Message Queue, see the *Message Queue Administration Guide*.

JMS Resources

The Java Message Service (JMS) API uses two kinds of administered objects:

- Connection factories, objects that allow an application to create other JMS objects programmatically
- Destinations, which serve as the repositories for messages

These objects are created administratively, and how they are created is specific to each implementation of JMS. In the Application Server, perform the following tasks:

- Create a connection factory by creating a connection factory resource
- Create a destination by creating two objects:
 - A physical destination
 - A destination resource that refers to the physical destination

JMS applications use the JNDI API to access the connection factory and destination resources. A JMS application normally uses at least one connection factory and at least one destination. To learn what resources to create, study the application or consult with the application developer.

There are three types of connection factories:

- `QueueConnectionFactory` objects, used for point-to-point communication
- `TopicConnectionFactory` objects, used for publish-subscribe communication
- `ConnectionFactory` objects, which can be used for both point-to-point and publish-subscribe communications; these are recommended for new applications

There are two kinds of destinations:

- Queue objects, used for point-to-point communication
- Topic objects, used for publish-subscribe communication

The chapters on JMS in the *J2EE 1.4 Tutorial* provide details on these two types of communication and other aspects of JMS (see <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>).

The order in which the resources are created does not matter.

For a J2EE application, specify connection factory and destination resources in the Application Server deployment descriptors as follows:

- Specify a connection factory JNDI name in a `resource-ref` or an `mdb-connection-factory` element.
- Specify a destination resource JNDI name in the `ejb` element for a message-driven bean and in the `message-destination` element.
- Specify a physical destination name in a `message-destination-link` element, within either a `message-driven` element of an enterprise bean deployment descriptor or a `message-destination-ref` element. In addition, specify it in the `message-destination` element. (The `message-destination-ref` element replaces the `resource-env-ref`

element, which is deprecated in new applications.) In the `message-destination` element of an Application Server deployment descriptor, link the physical destination name with the destination resource name.

The Relationship Between JMS Resources and Connector Resources

The Application Server implements JMS by using a system resource adapter named `jmsra`. When a user creates JMS resources, the Application Server automatically creates connector resources that appear under the Connectors node in the Admin Console's tree view.

For each JMS connection factory that a user creates, the Application Server creates a connector connection pool and connector resource. For each JMS destination a user creates, the Application Server creates an admin object resource. When the user deletes the JMS resources, the Application Server automatically deletes the connector resources.

It is possible to create connector resources for the JMS system resource adapter by using the Connectors node of the Admin Console instead of the JMS Resources node. See [Chapter 7, "Connector Resources,"](#) for details.

Admin Console Tasks for JMS Connection Factories

- ["To create a JMS connection factory resource" on page 97](#)
- ["To edit a JMS connection factory resource" on page 100](#)
- ["To delete a JMS connection factory resource" on page 101](#)

▼ To create a JMS connection factory resource

- 1 In the tree component, expand the Resources node, then expand the JMS Resources node.
- 2 Select the Connection Factories node.
- 3 On the JMS Connection Factories page, click New.
The Create JMS Connection Factory page appears.
- 4 In the JNDI Name field, type the name of the connection factory.
It is a recommended practice to use the naming subcontext prefix `jms/` for JMS resources. For example: `jms/ConnectionFactory1`.
- 5 From the Type drop-down list, choose either `javax.jms.ConnectionFactory`, `javax.jms.QueueConnectionFactory`, or `javax.jms.TopicConnectionFactory`.

- 6 **Select the Enabled checkbox to enable the resource at run time.**
- 7 **In the Advanced area, change values as needed for the connection factory attributes.**
For details about these attributes, see “[To edit a JMS connection factory resource](#)” on page 100. The Application Server applies these attributes to the connector connection pool created for the connection factory.
- 8 **For a JMS connection factory resource, specify the Transaction Support value as follows:**
 - **Specify `XATransaction` (the default value) for a resource that can be used for transactions that involve the use of more than one resource within a transaction scope.**
For example, this resource plus a JDBC resource, a connector resource, or another JMS connection factory resource. This value offers the most flexibility. A resource that is configured as `XATransaction` will participate in two-phase commit operations.
 - **Specify `LocalTransaction` for a resource that can be used either for transactions that involve only one resource within the transaction scope or as the last agent in a distributed transaction that involves more than one XA resource.**
This value offers significantly better performance. A resource that is configured as `LocalTransaction` will not be used in two-phase commit operations.
 - **Specify `NoTransaction` for a resource that can never participate in transactions.**
This setting is of limited use in JMS applications.
- 9 **In the Additional Properties area, provide values for properties required by applications. The following table lists the available properties.**

Property Name	Description
<code>ClientId</code>	Specifies a client ID for a connection factory that will be used by a durable subscriber.
<code>AddressList</code>	<p>Specifies the names (and, optionally, port numbers) of a message broker instance or instances with which applications will communicate. Each address in the list specifies the host name (and, optionally, host port and connection service) for the connection. For example, the value might be <code>earth</code> or <code>earth:7677</code>. Specify the port number if the message broker is running on a port other than the default (7676). If the property setting specifies multiple hosts and ports in a clustered environment, the first available host on the list is used unless the <code>AddressListBehavior</code> property is set to <code>RANDOM</code>.</p> <p>For details, see the <i>Message Queue Developer's Guide for Java Clients</i>.</p> <p>Default: The local host and default port number (7676). The client will attempt a connection to a broker on port 7676 of the local host.</p>

Property Name	Description
MessageServiceAddressList	Specifies the list of addresses. This property name is deprecated. Use <code>AddressList</code> instead.
UserName	The user name for the connection factory. Default: guest
Password	The password for the connection factory. Default: guest
ReconnectEnabled	If enabled (value = true), specifies that the client runtime attempts to reconnect to a message server (or the list of addresses in the <code>AddressList</code>) when a connection is lost. Default: true
ReconnectAttempts	Specifies the number of attempts to connect (or reconnect) for each address in the <code>AddressList</code> before the client runtime tries the next address in the list. A value of -1 indicates that the number of reconnect attempts is unlimited (the client runtime attempts to connect to the first address until it succeeds). Default: 3
ReconnectInterval	Specifies the interval in milliseconds between reconnect attempts. This applies for attempts on each address in the <code>AddressList</code> and for successive addresses in the list. If the interval is too short, the broker does not have time to recover. If it is too long, the reconnect might represent an unacceptable delay. Default: 30000
AddressListBehavior	Specifies whether connection attempts are in the order of addresses in the <code>AddressList</code> attribute (<code>PRIORITY</code>) or in a random order (<code>RANDOM</code>). <code>RANDOM</code> means that the reconnect chooses a random address from the <code>AddressList</code> . If many clients are likely to attempt a connection using the same connection factory, this value prevents them from all being connected to the same address. <code>PRIORITY</code> means that the reconnect always tries to connect to the first server address in the <code>AddressList</code> and uses another address only if the first broker is not available. Default: <code>RANDOM</code>
AddressListIterations	Specifies the number of times the client runtime iterates through the <code>AddressList</code> in an effort to establish (or reestablish) a connection). A value of -1 indicates that the number of attempts is unlimited. Default: 3

10 In the Targets area, do the following:

- a. **From the Available column, select the target or targets to which applications that use the resource will be deployed.**

The available targets include the available clusters and server instances as well as the default server instance, server.

- b. **Click Add to move the target to the Selected column.**

11 Click OK to save the connection factory.**More Information** Equivalent asadmin command

```
create-jms-resource
```

▼ To edit a JMS connection factory resource

- 1 **In the tree component, expand the Resources node, then expand the JMS Resources node.**

- 2 **Expand the Connection Factories node.**

- 3 **Select the connection factory to be edited.**

- 4 **On the Edit JMS Connection Factory page, you can perform these tasks:**

- Modify the text in the Description field.
- Select or deselect the Enabled checkbox to enable or disable the resource.
- Change the values of the attributes in the Advanced area.
- Add, remove, or modify properties.

- 5 **Optionally, click the Targets tab to go to the JMS Connection Factory Resource Targets page. On this page, do the following:**

- a. **Click Manage Targets to open the Manage Resource Targets page.**

On this page, move targets between the Available column and the Selected column. Make sure to place in the Selected column the target or targets to which applications that use the resource will be deployed. The available targets include the available clusters and server instances, as well as the default server instance, server. Click OK to save the changes.

- b. **Select the checkbox for a target, then click Enable or Disable to enable or disable the resource for that target.**

- 6 **Click Save to save the changes.**

▼ To delete a JMS connection factory resource

- 1 In the tree component, expand the Resources node, then expand the JMS Resources node.
- 2 Select the Connection Factories node.
- 3 On the JMS Connection Factories page, select the checkbox next to the name of the connection factory to be deleted.
- 4 Click Delete.

More Information Equivalent `asadmin` command

```
delete-jms-resource
```

Admin Console Tasks for JMS Destination Resources

- [“To create a JMS destination resource” on page 101](#)
- [“To edit a JMS destination resource” on page 102](#)
- [“To delete a JMS destination resource” on page 103](#)

▼ To create a JMS destination resource

- 1 In the tree component, expand the Resources node, then expand the JMS Resources node.
- 2 Select the Destination Resources node.
- 3 On the JMS Destination Resources page, click New. The Create JMS Destination Resource page appears.
- 4 In the JNDI Name field, type the name of the resource.
It is a recommended practice to use the naming subcontext prefix `jms/` for JMS resources. For example: `jms/Queue`.
- 5 From the Type drop-down list, choose either `javax.jms.Topic` or `javax.jms.Queue`.
- 6 Select the Enabled checkbox to enable the resource at run time.
- 7 In the Additional Properties area, provide values for properties.
The following table lists the available properties.

Property Name	Description
Name	(Required) The name of the physical destination to which the resource refers.
Description	A description of the physical destination.

8 In the Targets area, do the following:

- a. From the Available column, select the target or targets to which applications that use the resource will be deployed.**

The available targets include the available clusters and server instances, as well as the default server instance, server.

- b. Click Add to move the target to the Selected column.**

9 Click OK.

More Information Equivalent asadmin command

```
create-jms-resource
```

▼ To edit a JMS destination resource

- 1 In the tree component, expand the Resources node, then expand the JMS Resources node.
- 2 Expand the Destination Resources node.
- 3 Select the destination resource to be edited.
- 4 On the Edit JMS Destination Resource page, you can perform these tasks:
 - Change the type of the resource.
 - Modify the text in the Description field.
 - Select or deselect the Enabled checkbox to enable or disable the resource.
 - Add, remove, or modify the Name or Description property.
- 5 Click Save to save the changes.

- 6 **Optionally, click the Targets tab to go to the JMS Destination Resource Targets page. On this page, do the following:**
 - a. **Click Manage Targets to open the Manage Resource Targets page.**

On this page, move targets between the Available column and the Selected column. Make sure to place in the Selected column the target or targets to which applications that use the resource will be deployed. The available targets include the available clusters and server instances, as well as the default server instance, server. Click OK to save the changes.
 - b. **Select the checkbox for a target, then click Enable or Disable to enable or disable the resource for that target.**

▼ To delete a JMS destination resource

- 1 In the tree component, expand the Resources node, then expand the JMS Resources node.
- 2 Select the Destination Resources node.
- 3 On the JMS Destination Resources page, select the checkbox next to the name of the destination resource to be deleted.
- 4 Click Delete.

More Information Equivalent asadmin command

```
delete-jms-resource
```

Admin Console Tasks for JMS Physical Destinations

- [“To create a JMS physical destination” on page 103](#)
- [“To delete a JMS physical destination” on page 105](#)

▼ To create a JMS physical destination

For production purposes, always create physical destinations. During the development and testing phase, however, this step is not required. The first time an application accesses a destination resource, Message Queue automatically creates the physical destination specified by the Name property of the destination resource. The physical destination is temporary and expires after a period specified by a Message Queue configuration property.

- 1 In the tree component, expand the Configurations node, then expand the Java Message Service node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.
 - To configure the default settings for future instances that use a copy of default-config, select the default-config node.
- 3 Select the Physical Destinations node.
- 4 On the Physical Destinations page, click New.
The Create Physical Destination page appears.
- 5 In the Physical Destination Name field, type the name of the destination (for example, PhysicalQueue).
- 6 From the Type drop-down list, choose either topic or queue.
- 7 In the Additional Properties area, click Add Property to add a property.
The following table lists the one property currently available.

Property Name	Description
maxNumActiveConsumers	The maximum number of consumers that can be active in load-balanced delivery from a queue destination. A value of -1 means an unlimited number. The default is 1 if the destination is created for a standalone server instance and -1 if it is created for a cluster.

To modify the value of this property or to specify other physical destination properties, use the *MQ-install-dir/imq/bin/imqcmd* command. See the *Message Queue Administration Guide* for more information.

- 8 Click OK.

More Information Equivalent asadmin command

```
create-jmsdest
```


System Destination

The Physical Destinations page shows the system destination, a queue named `mq.sys.dmq`, to which expired and undeliverable messages are redirected. You can create destination resources, consumers, and browsers for this destination. You cannot delete it or send messages to it.

▼ To delete a JMS physical destination

- 1 In the tree component, expand the Configurations node, then expand the Java Message Service node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Select the Physical Destinations node.
- 4 On the Physical Destinations page, select the checkbox next to the name of the destination to be deleted.
- 5 Click Delete.

If you try to delete the system destination `mq.sys.dmq`, an error message appears.

More Information Equivalent `asadmin` command

```
delete-jmsdest
```

Admin Console Tasks for the JMS Provider

- [“To configure general properties for the JMS provider” on page 105](#)
- [“To create a JMS host” on page 109](#)
- [“To edit a JMS host” on page 111](#)
- [“To delete a JMS host” on page 111](#)

▼ To configure general properties for the JMS provider

Use the JMS Service page to configure properties to be used by all JMS connections.

For more information about configuring the JMS service, see the *Application Server Developer's Guide*.

1 In the tree component, select the Configurations node.

2 Select the instance to configure:

- **To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.**
- **To configure the default settings for future instances that use a copy of default-config, select the default-config node.**

3 Select the Java Message Service node to open the JMS Service page.

4 Edit the value in the Startup Timeout field to change the time the Application Server waits for the JMS service to start before aborting the startup.

On a slow or overloaded system, increase the value from the default (60).

5 From the Type drop-down list:

- Choose LOCAL (the default for the server-config configuration) to access the JMS service on the local host. The JMS service is started and managed by the Application Server.
- Choose REMOTE (the default for the default-config configuration) to access the JMS service on another system or on a cluster. If you choose REMOTE, the JMS service is not started by the Application Server the next time the server starts. Instead, the JMS service is started and managed via Message Queue, so you must start the Message Queue broker separately. For information about starting the broker, see the *Message Queue Administration Guide*. If you choose this value and are using a remote host, follow the instructions in [“To edit a JMS host” on page 111](#) to specify the name of the remote host.

6 In the Start Arguments field, type arguments to customize the JMS service startup.

Use any arguments available through the `MQ-install-dir/imq/bin/imqbrokerd` command.

7 Use the Reconnect checkbox to specify whether the JMS service attempts to reconnect to a message server (or the list of addresses in the AddressList) when a connection is lost.

By default, reconnection is enabled.

8 In the Reconnect Interval field, type the number of seconds between reconnect attempts.

This applies for attempts on each address in the AddressList and for successive addresses in the list. If it is too short, this time interval does not give a broker time to recover. If it is too long, the reconnect might represent an unacceptable delay.

The default value is 60 seconds.

- 9 In the Reconnect Attempts field, type the number of attempts to connect (or reconnect) for each address in the AddressList before the client runtime tries the next address in the list.**

A value of -1 indicates that the number of reconnect attempts is unlimited (the client runtime attempts to connect to the first address until it succeeds). The default value is 3.

- 10 Choose a host from the Default JMS Host drop-down list. The default is `default_JMS_host`.**

- 11 In the Address List Behavior drop-down list, choose whether connection attempts are in the order of addresses in the AddressList (priority) or in a random order (random).**

`priority` means that the reconnect always tries to connect to the first server address in the AddressList and uses another one only if the first broker is not available.

If there are many clients attempting a connection using the same connection factory, specify `random` to prevent them from all being connected to the same address.

The default is `random`.

- 12 In the Address List Iterations field, type the number of times the JMS service iterates through the AddressList in an effort to establish (or reestablish) a connection).**

A value of -1 indicates that the number of attempts is unlimited.

The default value is 3.

- 13 In the MQ Scheme and MQ Service fields, type the Message Queue address scheme name and the Message Queue connection service name if a non-default scheme or service is to be used.**

The full syntax for a message service address is

scheme://address-syntax

where the *scheme* and *address-syntax* are described in the table below.

The MQ Scheme and MQ Service are the values shown in the first two columns of the following table.

Scheme Name	Connection Service	Description	Address Syntax
mq	jms and ssljms	Message Queue client runtime makes a connection to the MQ Port Mapper at the specified host and port. The Port Mapper returns a list of the dynamically established connection service ports, and the Message Queue client runtime then makes a connection to the port hosting the specified connection service.	<p><code>[hostName][:port][serviceName]</code></p> <p>Defaults: <code>hostName = localhost</code>, <code>port = 7676</code>, <code>serviceName = jms</code></p> <p>Defaults only apply to the <code>jms</code> connection service. For the <code>ssljms</code> connection service, all variables need to be specified.</p> <p>Example: <code>mq:MyHost:7677/ssljms</code></p>

Scheme Name	Connection Service	Description	Address Syntax
mqtcp	jms	Message Queue client runtime makes a TCP connection to the specified host and port (bypassing the MQ Port Mapper) to establish a connection.	<i>hostName:port/jms</i> Example: <code>mqtcp:localhost:7676/jms</code>
mqssl	ssljms	Message Queue client runtime makes a secure SSL connection to the specified host and port (bypassing the MQ Port Mapper) to establish a connection.	<i>hostName:port/ssljms</i> Example: <code>mqssl:localhost:7676/ssljms</code>
http	httpjms	Message Queue client runtime makes an HTTP connection to an Message Queue tunnel servlet at the specified URL. (The broker must be configured to access the HTTP tunnel servlet, as described in the Message Queue <i>Administrator's Guide</i> .)	<i>hostName:port/contextRoot/tunnel</i> If multiple broker instances are using the same tunnel servlet, then the syntax for connecting to a specific broker instance (rather than a randomly selected one) is: <code>http://hostName:port/contextRoot/tunnel?serverName=hostName:instanceName</code>
https	httpsjms	Message Queue client runtime makes a secure HTTPS connection to the specified Message Queue tunnel servlet URL. (The broker must be configured to access the HTTPS tunnel servlet, as described in the Message Queue <i>Administrator's Guide</i> .)	<i>hostName:port/contextRoot/tunnel</i> If multiple broker instances are using the same tunnel servlet, then the syntax for connecting to a specific broker instance (rather than a randomly selected one) is: <code>https://hostName:port/contextRoot/tunnel?serverName=hostName:instanceName</code>

14 In the Additional Properties area, click Add Property to add a property.

The following table lists the available Sun Java System Message Queue broker configuration properties.

Property Name	Description
<code>instance-name</code>	Specifies the full Message Queue broker instance name. The default is <code>imqbroker</code> .
<code>instance-name-suffix</code>	Specifies a suffix to add to the full Message Queue broker instance name. The suffix is separated from the instance name by an underscore character (<code>_</code>). For example, if the instance name is <code>imqbroker</code> , appending the suffix <code>xyz</code> changes the instance name to <code>imqbroker_xyz</code> .
<code>append-version</code>	If <code>true</code> , appends the major and minor version numbers, preceded by underscore characters (<code>_</code>), to the full Message Queue broker instance name. For example, if the instance name is <code>imqbroker</code> , appending the version numbers changes the instance name to <code>imqbroker_8_0</code> . The default is <code>false</code> .

- 15 Click **Save** to save the changes, or click **Load Defaults** to restore the default values for the service.

More Information Accessing Remote Servers

Changing the provider and host to a remote system causes all JMS applications to run on the remote server. To use both the local server and one or more remote servers, create a connection factory resource with the `AddressList` property to create connections that access remote servers. See [“To create a JMS connection factory resource” on page 97](#).

▼ To verify JMS Service settings

- 1 In the tree component, select the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s **config** node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Select the **Java Message Service** node to open the **JMS Service** page.
- 4 Click **Ping**.

If the JMS service is up and running, the message “Ping succeeded: JMS service is running” appears.

More Information Equivalent asadmin command

```
jms-ping
```

▼ To create a JMS host

The Application Server Platform Edition expects there to be exactly one JMS host, whose default name is `default_JMS_host`. It is possible to create additional hosts, but the Application Server has no way of knowing about them.

- 1 In the tree component, expand the **Configurations** node.

2 Select the instance to configure:

- **To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.**
- **To configure the default settings for future instances that use a copy of default-config, select the default-config node.**

3 Expand the Java Message Service node.

4 Select the JMS Hosts node.

5 On the JMS Hosts page, click New. The Create JMS Host page appears.

6 In the Name field, type the name of the host. For example:

`NewJmsHost`

7 In the Host field, type the name or Internet Protocol (IP) address of the system where the JMS host will run (`localhost` or the name of the local or remote system).

8 In the Port field, type the port number of the JMS service.

Change this field only if the JMS service to be used is running on a non-default port. (The default port is 7676.)

9 In the Admin Username and Admin Password fields, type the Message Queue broker user name and password.

These are different from the Application Server user name and password. Edit these fields only if the Message Queue broker values have been changed using the `MQ-install-dir/imq/bin/imqusermgr` command. The default values are `admin` and `admin`.

10 Click OK.

More Information Equivalent `asadmin` command

`create-jms-host`

▼ To edit a JMS host

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.
 - To configure the default settings for future instances that use a copy of default-config, select the default-config node.
- 3 Expand the Java Message Service node.
- 4 Select the JMS Hosts node.
- 5 On the JMS Hosts page, select the host to be edited.
- 6 On the Edit JMS Host page, it is possible to perform these tasks:
 - Change the host name or Internet Protocol (IP) address in the Host field.
 - Change the port number of the JMS service in the Port field.
 - Change the values in the Admin Username and Admin Password fields.
- 7 Click Save to save the changes, or click Load Defaults to restore the default values for the host.

▼ To delete a JMS host

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.
 - To configure the default settings for future instances that use a copy of default-config, select the default-config node.
- 3 Expand the Java Message Service node.
- 4 Select the JMS Hosts node.

- 5 On the *JMS Hosts* page, select the checkbox next to the name of the host to be deleted.
- 6 Click **Delete**.



Caution – Do not delete all JMS hosts. Doing so will prevent Application Server from restarting. You must leave at least one JMS host.

More Information Equivalent `asadmin` command

`delete-jms-host`

Configuring JavaMail Resources

This chapter describes how to configure resources for applications that use the JavaMail API. It contains the following sections:

- “About JavaMail” on page 113
- “Admin Console Tasks for JavaMail” on page 114

About JavaMail

The JavaMail API is a set of abstract APIs that model a mail system. The API provides a platform-independent and protocol-independent framework to build mail and messaging applications. The JavaMail API provides facilities for reading and sending email. Service providers implement particular protocols.

The JavaMail API is implemented as a Java platform optional package and is also available as part of the J2EE platform.

The Application Server includes the JavaMail API along with JavaMail service providers that allow an application component to send email notifications over the Internet and to read email from IMAP and POP3 mail servers.

For more information about the JavaMail API, go to the JavaMail web site (<http://java.sun.com/products/javamail/>).

Admin Console Tasks for JavaMail

- [“To create a JavaMail session” on page 114](#)
- [“To edit a JavaMail session” on page 115](#)
- [“To delete a JavaMail session” on page 116](#)

▼ To create a JavaMail session

- 1 In the tree component, expand the Resources node, then select the JavaMail Sessions node.

- 2 On the JavaMail Sessions page, click New.

The Create JavaMail Session page appears.

- 3 In the JNDI Name field, type the name of the session.

It is a recommended practice to use the naming subcontext prefix `mail/` for JavaMail resources. For example: `mail/MySession`.

- 4 In the Mail Host field, type the host name of the default mail server.

The connect methods of the Store and Transport objects use this value if a protocol-specific host property is not supplied. The name must be resolvable to an actual host name.

- 5 In the Default User field, type the user name to provide when connecting to a mail server.

The connect methods of the Store and Transport objects use this value if a protocol-specific username property is not supplied.

- 6 In the Default Return Address field, type the email address of the default user, in the form *username@host . domain*.

- 7 Next to the Session label, select the Disable on all targets radio button if you do not want to enable the mail session at this time.

- 8 In the Advanced area, change the field values only if the Application Server's mail provider has been reconfigured to use a non-default store or transport protocol.

By default, the Store Protocol is `imap`; the Store Protocol Class is `com.sun.mail.imap.IMAPStore`; the Transport Protocol is `smtp`; and the Transport Protocol Class is `com.sun.mail.smtp.SMTPTransport`.

- 9 **Select the Debug checkbox to enable extra debugging output, including a protocol trace, for this mail session.**
If the JavaMail log level is set to FINE or finer, the debugging output is generated and is included in the system log file. See [“To configure log levels” on page 263](#) for information about setting the log level.
- 10 **In the Additional Properties area, click Add Property to add properties required by applications, such as a protocol-specific host or username property.**
The JavaMail API documentation lists the available properties <http://java.sun.com/products/javamail/javadocs/index.html>).
- 11 **In the Targets area, do the following:**
 - a. **From the Available column, select the target or targets to which applications that use the resource will be deployed.**
The available targets include the available clusters and server instances as well as the default server instance, server.
 - b. **Click Add to move the target to the Selected column.**
- 12 **Click OK to save the session.**

More Information Equivalent asadmin command
create-javamail-resource

▼ To edit a JavaMail session

- 1 In the tree component, expand the Resources node, then select the JavaMail Sessions node.
- 2 On the JavaMail Sessions page, select the session to be edited.
- 3 On the Edit JavaMail Session page, you can perform these tasks:
 - Modify the values in the Mail Host, Default User, Default Return Address, and Description fields.
 - Select or deselect the Enabled checkbox to enable or disable the resource.
 - Modify the values in the Advanced fields.
 - Add, remove, or modify properties.

- 4 **Click the Targets tab to go to the JavaMail Session Targets page. On this page, do the following:**
 - a. **Click Manage Targets to open the Manage Resource Targets page.**

On this page, move targets between the Available column and the Selected column. Make sure to place in the Selected column the target or targets to which applications that use the resource will be deployed. The available targets include the available clusters and server instances as well as the default server instance, `server`. Click OK to save the changes.
 - b. **Select the checkbox for a target, then click Enable or Disable to enable or disable the resource for that target.**
- 5 **Click Save to save the changes, or click Load Defaults to restore the default values for a mail session.**

▼ To delete a JavaMail session

- 1 **In the tree component, expand the Resources node, then select the JavaMail Sessions node.**
- 2 **On the JavaMail Sessions page, select the checkbox next to the name of the session to be deleted.**
- 3 **Click Delete.**

More Information Equivalent `asadmin` command

```
delete-javamail-resource
```

JNDI Resources

- [“About Java Naming and Directory Interface \(JNDI\)” on page 117](#)
- [“About Custom Resources” on page 119](#)
- [“About External JNDI Repositories and Resources” on page 122](#)

About Java Naming and Directory Interface (JNDI)

This section discusses the Java Naming and Directory Interface (JNDI). JNDI is an application programming interface (API) for accessing different kinds of naming and directory services. J2EE components locate objects by invoking the JNDI lookup method.

This section covers the following topics:

- [“JNDI Names and Resources” on page 117](#)
- [“J2EE Naming Services” on page 118](#)
- [“Naming References and Binding Information” on page 118](#)

JNDI Names and Resources

JNDI is the acronym for the Java Naming and Directory Interface API. By making calls to this API, applications locate resources and other program objects. A resource is a program object that provides connections to systems, such as database servers and messaging systems. (A JDBC resource is sometimes referred to as a data source.) Each resource object is identified by a unique, people-friendly name, called the JNDI name. A resource object and its JNDI name are bound together by the naming and directory service, which is included with the Application Server. To create a new resource, a new name-object binding is entered into the JNDI.

J2EE Naming Services

A JNDI name is a people-friendly name for an object. These names are bound to their objects by the naming and directory service that is provided by a J2EE server. Because J2EE components access this service through the JNDI API, the object usually uses its JNDI name. For example, the JNDI name of the PointBase database is `jdbc/Pointbase`. When it starts up, the Application Server reads information from the configuration file and automatically adds JNDI database names to the name space.

J2EE application clients, enterprise beans, and web components are required to have access to a JNDI naming environment.

The application component's naming environment is a mechanism that allows customization of the application component's business logic during deployment or assembly. Use of the application component's environment allows the application component to be customized without the need to access or change the application component's source code.

A J2EE container implements the application component's environment, and provides it to the application component instance as a JNDI naming context. The application component's environment is used as follows:

- The application component's business methods access the environment using the JNDI interfaces. The application component provider declares in the deployment descriptor all the environment entries that the application component expects to be provided in its environment at runtime.
- The container provides an implementation of the JNDI naming context that stores the application component environment. The container also provides the tools that allow the deployer to create and manage the environment of each application component.
- A deployer uses the tools provided by the container to initialize the environment entries that are declared in the application component's deployment descriptor. The deployer sets and modifies the values of the environment entries.
- The container makes the environment naming context available to the application component instances at runtime. The application component's instances use the JNDI interfaces to obtain the values of the environment entries.

Each application component defines its own set of environment entries. All instances of an application component within the same container share the same environment entries. Application component instances are not allowed to modify the environment at runtime.

Naming References and Binding Information

A resource reference is an element in a deployment descriptor that identifies the component's coded name for the resource. More specifically, the coded name references a connection factory for the resource. In the example given in the following section, the resource reference name is `jdbc/SavingsAccountDB`.

The JNDI name of a resource and the name of the resource reference are not the same. This approach to naming requires that you map the two names before deployment, but it also decouples components from resources. Because of this de-coupling, if at a later time the component needs to access a different resource, the name does not need to change. This flexibility also makes it easier for you to assemble J2EE applications from preexisting components.

The following table lists JNDI lookups and their associated references for the J2EE resources used by the Application Server.

TABLE 6-1 JNDI Lookups and Their Associated References

JNDI Lookup Name	Associated Reference
<code>java:comp/env</code>	Application environment entries
<code>java:comp/env/jdbc</code>	JDBC DataSource resource manager connection factories
<code>java:comp/env/ejb</code>	EJB References
<code>java:comp/UserTransaction</code>	UserTransaction references
<code>java:comp/env/mail</code>	JavaMail Session Connection Factories
<code>java:comp/env/url</code>	URL Connection Factories
<code>java:comp/env/jms</code>	JMS Connection Factories and Destinations
<code>java:comp/ORB</code>	ORB instance shared across application components

About Custom Resources

- [“Using Custom Resources” on page 119](#)
- [“To create a custom resource” on page 120](#)
- [“To edit a custom resource” on page 120](#)
- [“To delete a custom resource” on page 121](#)
- [“Listing Custom Resources” on page 121](#)

Using Custom Resources

A custom resource accesses a local JNDI repository and an external resource accesses an external JNDI repository. Both types of resources need user-specified factory class elements, JNDI name attributes, etc. In this section, we will discuss how to configure JNDI connection factory resources, for J2EE resources, and how to access these resources.

Within the Application Server, you can create, delete, and list resources, as well as `list-jndi-entities`.

▼ To create a custom resource

- 1 In the left pane of the Admin Console, open the Application Server instance for the JNDI configuration to be modified.

- 2 Open the JNDI tab and click Custom Resources.

If any custom resources have been created already, they are listed in the right pane. To create a new custom resource, click New. Open the JNDI tab and click New. A page for adding a new custom resource appears.

- 3 In the JNDI Name field, enter the name to use to access the resource.

This name will be registered in the JNDI naming service.

- 4 In the Resource Type field, enter a fully qualified type definition, as shown in the example above.

The Resource Type definition follows the format, xxx.xxx.

- 5 In the Factory Class field, enter a factory class name for the custom resource to be created.

The Factory Class is the user-specified name for the factory class. This class implements the `javax.naming.spi.ObjectFactory` interface.

- 6 In the Description field, enter a description for the resource to be creating.

This description is a string value and can include a maximum of 250 characters.

- 7 In the Additional Properties section, add the property name and value.

- 8 Mark the Custom Resource Enabled checkbox, to enable the custom resource.

- 9 Click OK to save your custom resource.

If the custom resource is deployed on a cluster or a stand-alone instance, you can manage targets using the Targets tab. The tab appears after the custom resource has been created. Set the target by entering the target name and clicking OK.

More Information Equivalent `asadmin` command

```
create-custom-resource
```

▼ To edit a custom resource

- 1 In the left pane of the Admin Console, open the Application Server instance for the JNDI configuration to be modified.

- 2 **Open JNDI and select Custom Resources.**
If any custom resources have been created already, they are listed in the right pane.
- 3 **Click on the file name in the right pane.**
- 4 **Edit the Resource Type field, the Factory Class field, or the Description field.**
- 5 **Mark the Custom Resource Enabled checkbox, to enable the custom resource.**
- 6 **Click Save to save the changes to the custom resource.**

▼ To delete a custom resource

- 1 **In the left pane of the Admin Console, open the JNDI tab.**
- 2 **Click Custom Resources.**
If any custom resources have been created already, they are listed in the right pane.
- 3 **Click in the box next to the name of the resource to be deleted.**
- 4 **Click Delete. The custom resource is deleted.**

More Information Equivalent `asadmin` command

```
delete-custom-resource
```

Listing Custom Resources

To list custom resources, type the `asadmin list-custom-resources` command. For example, to list custom resources on the host `plum`, type the following:

```
$asadmin list-custom-resources --host plum target6
```

For the full context, type `asadmin help list-custom-resources`.

About External JNDI Repositories and Resources

- “Using External JNDI Repositories and Resources” on page 122
- “To create an external resource” on page 122
- “To edit an external resource” on page 123
- “To delete an external resource” on page 124
- “Listing External Resources” on page 124

Using External JNDI Repositories and Resources

Often applications running on the Application Server require access to resources stored in an external JNDI repository. For example, generic Java objects could be stored in an LDAP server as per the Java schema. External JNDI resource elements let users configure such external resource repositories. The external JNDI factory must implement `javax.naming.spi.InitialContextFactory` interface.

An example of the use of an external JNDI resource is:

```
<resources>
<!-- external-jndi-resource element specifies how to access J2EE resources
-- stored in an external JNDI repository. The following example
-- illustrates how to access a java object stored in LDAP.
-- factory-class element specifies the JNDI InitialContext factory that
-- needs to be used to access the resource factory. property element
-- corresponds to the environment applicable to the external JNDI context
-- and jndi-lookup-name refers to the JNDI name to lookup to fetch the
-- designated (in this case the java) object.
-->
<external-jndi-resource jndi-name="test/myBean"
  jndi-lookup-name="cn=myBean"
  res-type="test.myBean"
  factory-class="com.sun.jndi.ldap.LdapCtxFactory">
  <property name="PROVIDER-URL" value="ldap://ldapserver:389/o=myObjects" />
  <property name="SECURITY_AUTHENTICATION" value="simple" />
  <property name="SECURITY_PRINCIPAL", value="cn=joeSmith, o=Engineering" />
  <property name="SECURITY_CREDENTIALS" value="changeit" />
</external-jndi-resource>
</resources>
```

▼ To create an external resource

- 1 In the left pane of the Admin Console, open the Application Server instance for the JNDI configuration to be modified.

2 Open JNDI and select External Resources.

If any external resources have been created already, they are listed in the right pane.

3 To create a new external resource, click New.**4 In the JNDI Name field, enter the name that is to be used to access the resource.**

This name is registered in the JNDI naming service.

5 In the Resource Type field, enter a fully qualified type definition, as shown in the example above.

The Resource Type definition follows the format, xxx.xxx.

6 In the JNDI Lookup field, enter the JNDI value to look up in the external repository.

For example, when creating an external resource to connect to an external repository, to test a bean class, the JNDI Lookup can look like this: `cn=testmybean`.

7 In the Factory Class field, enter a JNDI factory class external repository, for example,

`com.sun.jndi.ldap`.

This class implements the `javax.naming.spi.ObjectFactory` interface.

8 In the Description field, enter a description for the resource to be created.

This description is a string value and can include a maximum of 250 characters.

9 In the Additional Properties section, add the property name and value.**10 Mark the External Resource Enabled checkbox, to enable the external resource.****11 Click OK to save the external resource.**

If the external resource is deployed on a cluster or a stand-alone instance, you can manage targets using the Targets tab. The tab appears after the external resource has been created. Set the target by entering the target name and clicking OK.

More Information Equivalent `asadmin` command

```
create-jndi-resource
```

▼ To edit an external resource**1 In the left pane of the Admin Console, open the Application Server instance for the JNDI configuration to be modified.**

- 2 **Open JNDI and select External Resources.**
If any external resources have been created already, they are listed in the right pane.
- 3 **To edit an external resource, click on the file name in the right pane.**
- 4 **Edit the Resource Type field, the JNDI Lookup field, the Factory Class field, or the Description field.**
- 5 **Mark the External Resource Enabled checkbox, to enable the external resource.**
- 6 **Click Save to save the changes to the external resource.**

▼ **To delete an external resource**

- 1 **In the left pane of the Admin Console, open the JNDI tab.**
- 2 **Click External Resources.**
If any external resources have been created already, they are listed in the right pane.
- 3 **Click the box next to the name of the resource to be deleted.**
- 4 **Click Delete. The external resource is deleted.**

More Information Equivalent asadmin command

```
delete-jndi-resource
```

Listing External Resources

To list external resources, type the `asadmin list-jndi-resources` command and specify the JNDI name. For example, to list an external resource, type the following:

```
$asadmin list-jndi-resources --user adminuser --host plum jndi_name_test
```

For the full context, type `asadmin help list-jndi-resources`.

Connector Resources

This chapter explains how to configure connectors, which are used to access enterprise information systems (EISs). This chapter contains the following sections:

- [“About Connectors” on page 125](#)
- [“Admin Console Tasks for Connector Connection Pools” on page 126](#)
- [“Admin Console Tasks for Connector Resources” on page 129](#)
- [“Admin Console Tasks for Administered Object Resources” on page 132](#)

About Connectors

Also called a resource adapter, a connector module is a J2EE component that enables applications to interact with enterprise information systems (EISs). EIS software includes various types of systems: enterprise resource planning (ERP), mainframe transaction processing, and non-relational databases, among others. Like other J2EE modules, to install a connector module you deploy it.

A connector connection pool is a group of reusable connections for a particular EIS. To create a connector connection pool, specify the connector module (resource adapter) that is associated with the pool.

A connector resource is a program object that provides an application with a connection to an EIS. To create a connector resource, specify its JNDI name and its associated connection pool. Multiple connector resources can specify a single connection pool. The application locates the resource by looking up its JNDI name. (For more information on JNDI, see the section JNDI Names and Resources.) The JNDI name of a connector resource for an EIS is usually in the `java:comp/env/eis-specific` subcontext.

The Application Server implements JMS by using a connector module (resource adapter). See the section, [The Relationship Between JMS Resources and Connector Resources](#).

Admin Console Tasks for Connector Connection Pools

- “To Set Up EIS Access” on page 126
- “To create a connector connection pool” on page 126
- “To edit a connector connection pool” on page 127
- “To delete a connector connection pool” on page 129

▼ To Set Up EIS Access

- 1 Deploy (install) a connector. See [“To deploy a connector module” on page 67.](#)
- 2 Create a connection pool for the connector. See [“To create a connector connection pool” on page 126.](#)
- 3 Create a connector resource that is associated with the connection pool. See [“To create a connector resource” on page 129.](#)

▼ To create a connector connection pool

Before You Begin Before creating the pool, deploy the connector module (resource adapter) associated with the pool. The values that are specified for the new pool depend on the connector module that is deployed.

- 1 In the tree component, expand the Resource node and then the Connectors node.
- 2 Select the Connector Connection Pools node.
- 3 On the Connector Connection Pools page, click New.
- 4 On the first Create Connector Connection Pool page, specify the following settings:
 - a. In the Name field, enter a logical name for the pool.
Specify this name when creating a connector resource.
 - b. Select an entry from the Resource Adapter combo box.
The combo box displays a list of deployed resource adapters (connector modules).
- 5 Click Next.

- 6 **On the second Create Connector Connection Pool page, select a value from the Connection Definition combo box.**

The choices in the combo box depend on the resource adapter. Typically, a type of `ConnectionFactory` is specified, a factory instance to get a connection to the EIS.

- 7 **Click Next.**

- 8 **On the third and last Create Connector Connection Pool page, perform these tasks:**

- a. **In the General Settings section verify that the values are correct.**

- b. **For the fields in the Pool Settings section, the default values can be retained.**

These settings can be changed at a later time. See [Editing a Connector Connection Pool](#).

- c. **In the Additional Properties table, add any required properties.**

In the previous Create Connector Connection Pool page, you selected a class in the Connection Definition combo box. If this class is in the server's classpath, then the Additional Properties table displays default properties.

- 9 **Click Finish.**

More Information Equivalent `asadmin` command

```
create-connector-connection-pool
```

▼ To edit a connector connection pool

- 1 **In the tree component, expand the Resources node and then the Connectors node.**
- 2 **Expand the Connector Connection Pools node.**
- 3 **Select the node for the pool you want to edit.**
- 4 **On the Edit Connector Connection Pool page, you can change settings that control the number of connections in the pool. See the following table.**

Parameter	Description
Initial and Minimum Pool Size	The minimum number of connections in the pool. This value also determines the number of connections placed in the pool when the pool is first created or when application server starts.

Parameter	Description
Maximum Pool Size	The maximum number of connections in the pool.
Pool Resize Quantity	When the pool shrinks toward the minimum pool size it is resized in batches. This value determines the number of connections in the batch. Making this value too large will delay connection recycling; making it too small will be less efficient.
Idle Timeout	The maximum time in seconds that a connection can remain idle in the pool. After this time expires, the connection will be removed from the pool.
Max Wait Time	The amount of time the application that has requested a connection will wait before getting a connection timeout. Because the default wait time is long, the application might appear to hang indefinitely.
On Any Failure	If you select the checkbox labelled Close All Connections, if a single connection fails, then the application server will close all connections in the pool and then reestablish them. If you do not select the checkbox, then individual connections will be reestablished only when they are used.
Transaction Support	<p>Use the Transaction Support list to select the type of transaction support for the connection pool. The chosen transaction support overrides the transaction support attribute in the resource adapter associated with this connection pool in a downward compatible way. In other words, it can support a lower transaction level than that specified in the resource adapter or the same transaction level as that specified in resource adapter, but it cannot specify a higher level.</p> <p>The transaction support options include the following.</p> <p>The None selection from the Transaction Support menu indicates that the resource adapter does not support resource manager local or JTA transactions and does not implement XAResource or LocalTransaction interfaces.</p> <p>Local transaction support means that the resource adapter supports local transactions by implementing the LocalTransaction interface. Local transactions are managed internal to a resource manager and involve no external transaction managers.</p> <p>XA transaction support means that the resource adapter supports resource manager local and JTA transactions by implementing the LocalTransaction and XAResource interfaces. XA transactions are controlled and coordinated by a transaction manager external to a resource manager. Local transactions are managed internal to a resource manager and involve no external transaction managers.</p>

5 In the Additional Properties table, specify name-value pairs.

The properties specified depend on the resource adapter used by this pool. The name-value pairs specified by the deployer using this table can be used to override the default values for the properties defined by the resource-adapter vendor.

- 6 On the **Security Maps** tabbed pane, create or modify a security map for the connection pool. See [“About Security Maps” on page 188](#) for information on how to create a security map.
- 7 Click **Save**.

▼ To delete a connector connection pool

- 1 In the tree component, expand the **Resources** node and then the **Connectors** node.
- 2 Select the **Connector Connection Pools** node.
- 3 On the **Connector Connection Pools** page, select the checkbox for the pool to be deleted.
- 4 Click **Delete**.

More Information Equivalent `asadmin` command
`delete-connector-connection-pool`

Admin Console Tasks for Connector Resources

- [“To create a connector resource” on page 129](#)
- [“To edit a connector resource” on page 130](#)
- [“Deleting a Connector Resource” on page 131](#)
- [“To configure the connector service” on page 131](#)

▼ To create a connector resource

A connector resource (data source) provides applications with a connection to an EIS.

Before You Begin Before creating a connector resource, first create a connector connection pool.

- 1 In the tree component, expand the **Resources** node and then the **Connectors** node.
- 2 Expand the **Connector Resources** node.
- 3 On the **Connector Resources** page, click **New**.

- 4 **On the Create Connector Resources page, specify the resource's settings:**
 - a. **In the JNDI Name field, type a unique name, for example: `eis/myERP`.**
Don't forget the forward slash.
 - b. **From the Pool Name combo box, choose the connection pool to which the new connector resource belongs.**
 - c. **To change the resource to be unavailable, select the Disable on All Targets radio button.**
By default, the resource is available (enabled) as soon as it is created.
 - d. **In the Targets section of the page, select the domain, cluster, or server instances where the connector resource will reside, from the Available field and click Add.**
If you do not want to deploy the connector resource to one of the domains, clusters, or server instances listed in the Selected field, select it from the field and click Remove.
- 5 **Click OK.**

More Information Equivalent `asadmin` command

`create-connector-resource`

▼ **To edit a connector resource**

- 1 **In the tree component, expand the Resources node and then the Connectors node.**
- 2 **Expand the Connector Resources node.**
- 3 **Select the node for the connector resource that you want to edit.**
- 4 **On the Edit Connector Resources page, you can select a different connection pool from the Pool Name menu.**
- 5 **On the Targets tabbed pane, you can edit the targets on which the connector resource is deployed by clicking Manage Targets.**
See [“To create a connector resource” on page 129](#) for more information on targets.
- 6 **Click Save to apply the edits.**

▼ Deleting a Connector Resource

- 1 In the tree component, expand the Resources node and then the Connectors node.
- 2 Select the Connector Resources node.
- 3 On the Connector Resources page, select the checkbox for the resource to be deleted.
- 4 Click Delete.

More Information Equivalent `asadmin` command

```
delete-connector-resource
```

▼ To configure the connector service

Use the Connector Service screen to configure the connector container for all resource adapters deployed to this cluster or server instance.

- 1 Select Configurations from the tree.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.
 - To configure the default settings for future instances that use a copy of default-config, select the default-config node.
- 3 Select the Connector Service node.
- 4 Specify the shutdown timeout in seconds in the Shutdown Timeout field.

Enter an integer representing the number of seconds that the application server waits to allow the `ResourceAdapter.stop` method of the connector module's instance to complete. Resource adapters that take longer than the specified shutdown timeout are ignored by the application server and the shutdown procedure continues. The default shutdown timeout is 30 seconds. Click Load Defaults to select the default shutdown timeout for the resource adapters deployed to this cluster or server instance.

Admin Console Tasks for Administered Object Resources

- [“To create an administered object resource” on page 132](#)
- [“To edit an administered object resource” on page 133](#)
- [“To delete an administered object resource” on page 133](#)

▼ To create an administered object resource

Packaged within a resource adapter (connector module), an administered object provides specialized functionality for an application. For example, an administered object might provide access to a parser that is specific to the resource adapter and its associated EIS. The object can be administered; that is, it can be configured by an administrator. To configure the object, add name-value property pairs in the Create or Edit Admin Object Resource pages. When creating an administered object resource, associate the administered object to a JNDI name.

The Application Server implements JMS by using resource adapter. For each JMS destination created, the Application Server automatically creates an administered object resource.

- 1 In the tree component, expand the Resources node and then the Connectors node.**
- 2 Expand the Admin Object Resources node.**
- 3 On the Admin Object Resources page, click New.**
- 4 On the Admin Object Resources page, specify the following settings:**
 - a. In the JNDI Name field, type a unique name that identifies the resource.**
 - b. In the Resource Type field, enter the Java type for the resource.**
 - c. From the Resource Adapter combo box, select the resource adapter that contains the administered object.**
 - d. Select or deselect the checkbox to enable or disable the resource.**
 - e. Click Next.**
- 5 On the second Create Admin Object Resource page, the following tasks can be performed.**
 - a. To configure the administered object with name-value property pairs, click Add Property.**

- b. In the **Targets** section of the page, select the domain, cluster, or server instances where the administered object will reside, from the **Available** field and click **Add**.

To undeploy the administered object to one of the domains, clusters, or server instances listed in the **Selected** field, select it from the field and click **Remove**.

- 6 Click **Finish**.

More Information Equivalent `asadmin` command

```
create-admin-object
```

▼ To edit an administered object resource

- 1 In the tree component, expand the **Resource** node and then the **Connectors** node.
- 2 Expand the **Administered Object Resources** node.
- 3 Select the node for the administered object resource to be edited.
- 4 On the **Edit Administered Object Resources** page, modify values specified in **Creating an Administered Object Resource**.
- 5 On the **Targets** tabbed pane, edit the targets on which the administered object is deployed by clicking **Manage Targets**.
See [“To create an administered object resource” on page 132](#) for more information on targets.
- 6 Click **Save** to apply the edits.

▼ To delete an administered object resource

- 1 In the tree component, expand the **Resources** node and then the **Connectors** node.
- 2 Select the **Administered Object Resources** node.
- 3 On the **Administered Object Resources** page, select the checkbox for the resource to be deleted.
- 4 Click **Delete**.

More Information Equivalent `asadmin` command

```
delete-admin-object
```


J2EE Containers

This chapter explains how to configure the J2EE containers included in the server. This chapter contains following sections:

- “About the J2EE Containers” on page 135
- “Admin Console Tasks for the J2EE Containers” on page 136

About the J2EE Containers

This section describes the J2EE containers included with the Application Server.

- “Types of J2EE Containers” on page 135
- “The Web Container” on page 135
- “The EJB Container” on page 136

Types of J2EE Containers

J2EE containers provide runtime support for J2EE application components. J2EE application components use the protocols and methods of the container to access other application components and services provided by the server. The Application Server provides an application client container, an applet container, a Web container, and an EJB container. For a diagram that shows the containers, see the section “[Application Server Architecture](#)” on page 30.

The Web Container

The Web Container is a J2EE container that hosts web applications. The web container extends the web server functionality by providing developers the environment to run servlets and JavaServer Pages (JSP files).

The EJB Container

Enterprise beans (EJB components) are Java programming language server components that contain business logic. The EJB container provides local and remote access to enterprise beans.

There are three types of enterprise beans: session beans, entity beans, and message-driven beans. Session beans represent transient objects and processes and typically are used by a single client. Entity beans represent persistent data, typically maintained in a database. Message-driven beans are used to pass messages asynchronously to application modules and services.

The container is responsible for creating the enterprise bean, binding the enterprise bean to the naming service so other application components can access the enterprise bean, ensuring only authorized clients have access to the enterprise bean's methods, saving the bean's state to persistent storage, caching the state of the bean, and activating or passivating the bean when necessary.

Admin Console Tasks for the J2EE Containers

- [“Configuring the General Web Container Settings” on page 136](#)
- [“Configuring the General EJB Settings” on page 139](#)
- [“Configuring the Message-Driven Bean Settings” on page 141](#)
- [“Configuring the EJB Timer Service Settings” on page 142](#)

Configuring the General Web Container Settings

In this release, there are no container-wide settings for the Web container in the Admin Console.

Configuring Web Container Sessions

This section describes the HTTP session settings in the Web container. HTTP sessions are unique web sessions that have their state data written to a persistent store.

- [“To set the session timeout value” on page 137](#)
- [“To configure the manager properties” on page 137](#)
- [“To configure the store properties” on page 138](#)

▼ To set the session timeout value

- 1 In the tree component, select the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the Web Container node.
- 4 Click the Session Properties tab.
- 5 In the Session Timeout field, enter the number of seconds that a session is valid.
- 6 Click Save.

▼ To configure the manager properties

The session manager provides the means to configure how sessions are created and destroyed, where session state is stored, and the maximum number of sessions.

To change the session manager settings:

- 1 In the tree component select the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the Web Container node.
- 4 Click the Manager Properties tab.
- 5 Set the Reap Interval value.

The Reap Interval field is the number of seconds before the inactive session data is deleted from the store.

6 Set the Max Sessions value.

The Max Sessions field is the maximum number of sessions allowed.

7 Set the Session Filename value.

The Session Filename field is the file that contains the session data.

8 Set the Session ID Generator Classname value.

The Session ID Generator Classname field allows you to specify a custom class for generating unique session IDs. Only one session ID generator class per server instance is permitted, and all instances in a cluster must use the same session ID generator to prevent session key collision.

Custom session ID generator classes must implement the `com.sun.enterprise.util.uuid.UuidGenerator` interface:

```
package com.sun.enterprise.util.uuid;

public interface UuidGenerator {

    public String generateUuid();
    public String generateUuid(Object obj); //obj is the session object
}
```

The class must be in the Application Server classpath.

9 Click Save.**▼ To configure the store properties****1 In the tree component, select the Configurations node.****2 Select the instance to configure:**

- **To configure a particular instance, select the instance's config node. For example, the default instance, server, select the `server-config` node.**
- **To configure the default settings for all instances, select the `default-config` node.**

3 Select the Web Container node.**4 Click the Store Properties tab.****5 Set the Reap Interval.**

The Reap Interval field is the number of seconds before the inactive session data is deleted from the store.

6 Click Save.

Configuring the General EJB Settings

This section describes the following settings, which apply to all enterprise bean containers on the server:

- “Session Store Location” on page 139
- “To configure EJB pool settings” on page 139
- “To configure EJB cache settings” on page 140

To override the defaults on a per-container basis, adjust the values in the enterprise bean’s `sun-ejb-jar.xml` file. For details, see the *Application Server Developer’s Guide*.

Session Store Location

The Session Store Location field specifies the directory where passivated beans and persisted HTTP sessions are stored on the file system.

Passivated beans are enterprise beans that have had their state written to a file on the file system. Passivated beans typically have been idle for a certain period of time, and are not currently being accessed by clients.

Similar to passivated beans, persisted HTTP sessions are individual web sessions that have had their state written to a file on the file system.

The Commit Option field specifies how the container caches passivated entity bean instances between transactions.

Option B caches entity bean instances between transactions, and is selected by default. Option C disables caching.

▼ To configure EJB pool settings

The container maintains a pool of enterprise beans in order to respond to client requests without the performance hit that results from creating the beans. These settings only apply to stateless session beans and entity beans.

If you experience performance problems in an application that uses deployed enterprise beans, creating a pool, or increasing the number of beans maintained by an existing pool, can help increase the application’s performance.

By default, the container maintains a pool of enterprise beans.

1 In the tree component, select the Configurations node.

- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance's config node. For example, the default instance, server, select the server-config node.**
 - **To configure the default settings for all instances, select the default-config node.**
- 3 **Select the EJB Container node.**
- 4 **Under Pool Settings in the Initial and Minimum Pool Size field enter the minimum number of beans the container creates in the pool.**
- 5 **In the Maximum Pool Size field enter the maximum number of beans the container maintains in the pool, at any time.**
- 6 **In the Pool Resize Quantity field enter the number of beans that will be removed from the pool if they are idle for more than the time specified in Pool Idle Timeout.**
- 7 **In the Pool Idle Timeout field enter the time, in seconds, that a bean in the pool can remain idle before it will be removed from the pool.**
- 8 **Click Save.**
- 9 **Restart the Application Server.**

▼ **To configure EJB cache settings**

The container maintains a cache of enterprise bean data for the most used enterprise beans. This allows the container to respond more quickly to requests from other application modules for data from the enterprise beans. This section applies only to stateful session beans and entity beans.

Cached enterprise beans are in one of three states: active, idle, or passivated. An active enterprise bean is currently being accessed by clients. An idle enterprise bean's data is currently in the cache, but no clients are accessing the bean. A passivated bean's data is temporarily stored, and read back into the cache if a client requests the bean.

- 1 **In the tree component, select the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance's config node. For example, the default instance, server, select the server-config node.**
 - **To configure the default settings for all instances, select the default-config node.**

- 3 Select the EJB Container node.**
- 4 Adjust the maximum cache size in the Max Cache Size field.**

Increase the maximum number of beans to cache to eliminate the overhead of bean creation and destruction. However, if the cache is increased, the server consumes more memory and resources. Be sure your operating environment is sufficient for your cache settings.
- 5 Adjust the cache resize quantity in the Cache Resize Quantity field.**

When the maximum number of cached beans is reached, the container removes a number of passivated beans from the backup store, set to 32 by default.
- 6 Adjust the rate, in seconds, at which the cache cleanup is scheduled for entity beans in the Cache Idle Timeout field.**

If a cached entity bean has been idle a certain amount of time, it is passivated. That is, the bean's state is written to a backup store.
- 7 Adjust the time, in seconds, after which stateful session beans are removed from the cache or passivated store in the Removal Timeout field.**
- 8 Configure the policy the container uses to remove stateful session beans in the Removal Selection Policy field.**

The container decides which stateful session beans to remove based on the policy set in the Removal Selection Policy field. There are three possible policies the container uses to remove beans from the cache:

 - Not recently used (NRU)
 - First in, first out (FIFO)
 - Least recently used (LRU)

The NRU policy removes a bean that hasn't been used recently. The FIFO policy removes the oldest bean in the cache. The LRU policy removes the least recently accessed bean. By default, the NRU policy is used by the container.

Entity beans are always removed using the FIFO policy.
- 9 Click Save.**
- 10 Restart the Application Server.**

Configuring the Message-Driven Bean Settings

The pool for message-driven beans is similar to the pool for session beans described in [“To configure EJB pool settings” on page 139](#). By default, the container maintains a pool of message—driven beans.

To adjust the configuration of this pool:

▼ To configure MDB pool settings

- 1 In the tree component, select the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the EJB Container node.
- 4 Click the MDB Settings tab.
- 5 Under Pool Settings in the Initial and Minimum Pool Size field, enter the minimum number of message beans the container creates in the pool.
- 6 In the Maximum Pool Size field, enter the maximum number of beans the container maintains in the pool, at any time.
- 7 In the Pool Resize Quantity field, enter the number of beans that are removed from the pool if they are idle for more than the time specified in Pool Idle Timeout.
- 8 In the Pool Idle Timeout field, enter the time, in seconds, that a bean in the pool can remain idle before it is removed from the pool.
- 9 Click Save.
- 10 Restart the Application Server.

Configuring the EJB Timer Service Settings

The timer service is a persistent and transactional notification service provided by the enterprise bean container used to schedule notifications or events used by enterprise beans. All enterprise beans except stateful session beans can receive notifications from the timer service. Timers set by the service are not destroyed when the server is shut down or restarted.

▼ To configure the Timer Service

- 1 In the tree component, select the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the EJB Container node.
- 4 Click the EJB Timer Service tab.
- 5 Set the minimum delivery interval in milliseconds in the Minimum Delivery Interval field.

Minimum Delivery Interval is the minimum number of milliseconds allowed before the next timer expiration for a particular timer can occur. Setting this interval too low can cause server overload.
- 6 Set the maximum number of attempts the timer service makes to deliver the notification in the Maximum Redeliveries field.
- 7 Set the interval, in milliseconds, between redelivery attempts in the Redelivery Interval field.
- 8 Click Save.
- 9 Restart the Application Server.

▼ To use an external database with the Timer Service

The timer service by default uses an embedded database to store timers.

- 1 Set up a JDBC resource for the database, as described in [“To create a JDBC resource” on page 88](#).
- 2 Enter the JNDI name of the resource in the Timer DataSource field.
- 3 Click Save.
- 4 Restart the Application Server.

Sample timer database creation files are provided for PointBase and Oracle at `install-dir/lib/install/databases/`.

Configuring Security

This chapter describes some core application server security concepts, and describes how to configure security for the Application Server. This chapter contains the following topics:

- “About Application Server Security” on page 145
- “Admin Console Tasks for Security” on page 162
- “Admin Console Tasks for Realms” on page 164
- “Admin Console Tasks for JACC Providers” on page 176
- “Admin Console Tasks for Audit Modules” on page 179
- “Admin Console Tasks for Security” on page 162
- “Admin Console Tasks for Listeners and JMX Connectors” on page 183
- “Admin Console Tasks for Connector Connection Pools” on page 187
- “Working with Certificates and SSL” on page 191
- “Further Information” on page 199

About Application Server Security

- “Overview of Security” on page 145
- “About Authentication and Authorization” on page 151
- “Understanding Users, Groups, Roles, and Realms” on page 153
- “Introduction to Certificates and SSL” on page 156
- “About Firewalls” on page 159
- “Managing Security With the Admin Console” on page 159

Overview of Security

Security is about protecting data: how to prevent unauthorized access or damage to it in storage or transit. The Application Server has a dynamic, extensible security architecture based on the J2EE standard. Built in security features include cryptography, authentication and authorization, and public key infrastructure. The Application Server is built on the Java security

model, which uses a sandbox where applications can run safely, without potential risk to systems or users. The following topics are discussed:

- [“Understanding Application and System Security”](#) on page 146
- [“Tools for Managing Security”](#) on page 146
- [“Managing Security of Passwords”](#) on page 147
- [“Assigning Security Responsibilities”](#) on page 150

Understanding Application and System Security

Broadly, there are two kinds of application security:

- In *programmatic security*, application code written by the developer handles security chores. As an administrator, you don't have any control over this mechanism. Generally, programmatic security is discouraged since it hard-codes security configurations in the application instead of managing it through the J2EE containers.
- In *declarative security*, the container (the Application Server) handles security through an application's deployment descriptors. You can control declarative security by editing deployment descriptors directly or with a tool such as `deploytool`. Because deployment descriptors can change after an application is developed, declarative security allows for more flexibility.

In addition to application security, there is also *system security*, which affects all the applications on an Application Server system.

Programmatic security is controlled by the application developer, so this document does not discuss it; declarative security is somewhat less so, and this document touches on it occasionally. This document is intended primarily for system administrators, and so focuses on system security.

Tools for Managing Security

The Application Server provides the following tools for managing security:

- `Admin Console`, a browser-based tool used to configure security for the entire server, to manage users, groups, and realms, and to perform other system-wide security tasks. For a general introduction to the Admin Console, see [“Tools for Administration”](#) on page 33. For an overview of the security tasks you can perform with the Admin Console, see [“Managing Security With the Admin Console”](#) on page 159.
- `asadmin`, a command-line tool that performs many of the same tasks as the Admin Console. You may be able to do some things with `asadmin` that you cannot do with Admin Console. You perform `asadmin` commands from either a command prompt or from a script, to automate repetitive tasks. For a general introduction to `asadmin`, see [“Tools for Administration”](#) on page 33.

- `deploytool`, a graphical packaging and deployment tool for editing application deployment descriptors to control individual applications' security. Because `deploytool` is intended for application developers, this document does not describe its use in detail. For instructions on using `deploytool`, see the tool's online help and *The J2EE 1.4 Tutorial* at <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.

The Java 2 Platform, Standard Edition (J2SE) provides two tools for managing security:

- `keytool`, a command-line utility for managing digital certificates and key pairs. Use `keytool` to manage users in the certificate realm.
- `policytool`, a graphical utility for managing system-wide Java security policies. As an administrator, you will rarely need to use `policytool`.

For more information on using `keytool`, `policytool`, and other Java security tools, see *Java 2 SDK Tools and Utilities* at <http://java.sun.com/j2se/1.4.2/docs/tooldocs/tools.html#security>.

In the Enterprise Edition, two other tools that implement Network Security Services (NSS) are available for managing security. For more information on NSS, go to <http://www.mozilla.org/projects/security/pki/nss/>. The tools for managing security include the following:

- `certutil`, a command-line utility for managing certificates and key databases.
- `pk12util`, a command-line utility used to import and export keys and certificates between the certificate/key databases and files in PKCS12 format.

For more information on using `certutil`, `pk12util`, and other NSS security tools, see *NSS Security Tools* at <http://www.mozilla.org/projects/security/pki/nss/tools>.

Managing Security of Passwords

In this release of the Application Server, the file `domain.xml`, which contains the specifications for a particular domain, initially contains the password of the Sun Java System Message Queue broker in clear text. The element in the `domain.xml` file that contains this password is the `admin-password` attribute of the `jms-host` element. Because this password is not changeable at installation time, it is not a significant security impact.

However, use the Admin Console to add users and resources and assign passwords to these users and resources. Some of these passwords are written to the `domain.xml` file in clear text, for example, passwords for accessing a database. Having these passwords in clear text in the `domain.xml` file can present a security hazard. You can encrypt any password in `domain.xml`, including the `admin-password` attribute or a database password using the following procedure.

▼ To encrypt a password in domain.xml

- 1 **From the directory where the `domain.xml` file resides (*domain-dir/config* by default), run the following `asadmin` command:**

```
asadmin create-password-alias --user admin alias-name
```

For example,

```
asadmin create-password-alias --user admin jms-password
```

A password prompt appears (admin in this case). Refer to the man pages for the `create-password-alias`, `list-password-aliases`, `delete-password-alias` commands for more information.

- 2 **Remove and replace the password in `domain.xml`. This is accomplished using the `asadmin set` command. An example of using the `set` command for this purpose is as follows:**

```
asadmin set --user admin server.jms-service.jms-host.  
default_JMS_host.admin-password='${ALIAS=jms-password}'
```

- 3 **Restart the Application Server for the relevant domain.**

Protecting files with encoded passwords

Some files contain encoded passwords that need protecting using file system permissions. These files include the following:

- *domain-dir/master-password*
This file contains the encoded master password and should be protected with file system permissions 600.
- Any password file created to pass as an argument using the `--passwordfile` argument to `asadmin` should be protected with file system permissions 600.

▼ To change the master password

The master password (MP) is an overall shared password. It is never used for authentication and is never transmitted over the network. This password is the choke point for overall security; the user can choose to enter it manually when required, or obscure it in a file. It is the most sensitive piece of data in the system. The user can force prompting for the MP by removing this file. When the master password is changed, it is re-saved in the `master-password` keystore.

- 1 **Stop the Application Server for the domain. Use the `asadmin change-master-password` command, which prompts for the old and new passwords, then re-encrypts all dependent items. For example,**

```
asadmin change-master-password>  
Please enter the master password>
```

Please enter the new master password>
Please enter the the new master password again>

2 Restart the Application Server.



Caution – At this point in time, server instances that are running must not be started and running server instances must not be restarted until the SMP on their corresponding node agent has been changed. If a server instance is restarted before changing its SMP, it will fail to come up.

- 3 **Stop each node agent and its related servers one at a time. Run the `asadmin change-master-password` command again, and then restart the node agent and its related servers.**
- 4 **Continue with the next node agent until all node agents have been addressed. In this way, a rolling change is accomplished.**

▼ To change the admin password

Encrypting the admin password was discussed in “[Managing Security of Passwords](#)” on [page 147](#). Encrypting the admin password is strongly encouraged. If you want to change the admin password before encrypting it, use the `asadmin set` command. An example of using the `set` command for this purpose is as follows:

```
asadmin set --user admin
server.jms-service.jms-host.default_JMS_host.admin-password=new_pwd
```

It is also possible to change the admin password using the Admin Console as in the following procedure.

- 1 **In the Admin Console tree component, expand the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, expand the instance’s config node. For example, the default instance, server, expand the server-config node.**
 - **To configure the default settings for all instances, expand the default-config node.**
- 3 **Expand the Security node.**
- 4 **Expand the Realms node.**
- 5 **Select the `admin-realm` node.**
- 6 **Click the Manage Users button from the Edit Realm page.**

- 7 Select the user named `admin`.
- 8 Enter the new password and confirm the password.
- 9 Click **Save** to save or click **Close** to close without saving.

Assigning Security Responsibilities

Security responsibilities are assigned to the following:

- “Application Developer” on page 150
- “Application Deployer” on page 150
- “System Administrator” on page 150

Application Developer

The application developer is responsible for the following:

- Specifying roles and role-based access restrictions for application components.
- Defining an application’s authentication method and specifying the parts of the application that are secured.

An application developer can use tools such as `deploytool` to edit application deployment descriptors. These security tasks are discussed in more detail in the *Security* chapter of *The J2EE 1.4 Tutorial*, which can be viewed at <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.

Application Deployer

The application deployer is responsible for:

- Mapping users or groups (or both) to security roles.
- Refining the privileges required to access component methods to suit the requirements of the specific deployment scenario.

An application deployer can use tools such as `deploytool` to edit application deployment descriptors. These security tasks are discussed in more detail in the *Security* chapter of *The J2EE 1.4 Tutorial*, which can be viewed at <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.

System Administrator

The system administrator is responsible for:

- Configuring security realms.
- Managing user accounts and groups.

- Managing audit logs.
- **Managing server certificates and configuring the server's use of secure sockets layer (SSL).**
- Handling other miscellaneous system-wide security features, such as **security maps for connector connection pools**, additional JACC Providers, and so on.

A system administrator uses the Admin Console to manage server security settings and `certutil` to manage certificates. This document is intended primarily for system administrators.

About Authentication and Authorization

Authentication and authorization are central concepts of application server security. The following topics are discussed related to authentication and authorization:

- “Authenticating Entities” on page 151
- “Authorizing Users” on page 152
- “Specifying JACC Providers” on page 152
- “Auditing Authentication and Authorization Decisions” on page 153
- “Configuring Message Security” on page 153

Authenticating Entities

Authentication is the way an entity (a user, an application, or a component) determines that another entity is who it claims to be. An entity uses *security credentials* to authenticate itself. The credentials may be a user name and password, a digital certificate, or something else.

Typically, authentication means a user logging in to an application with a user name and password; but it might also refer to an EJB providing security credentials when it requests a resource from the server. Usually, servers or applications require clients to authenticate; additionally, clients can require servers to authenticate themselves, too. When authentication is bidirectional, it is called mutual authentication.

When an entity tries to access a protected resource, the Application Server uses the authentication mechanism configured for that resource to determine whether to grant access. For example, a user can enter a user name and password in a Web browser, and if the application verifies those credentials, the user is authenticated. The user is associated with this authenticated security identity for the remainder of the session.

The Application Server supports four types of authentication, as outlined in “Authenticating Entities” on page 151. An application specifies the type of authentication it uses within its deployment descriptors. For more information on using `deploytool` to configure the authentication method for an application, see *The J2EE 1.4 Tutorial* at <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.

TABLE 9-1 Application Server Authentication Methods

Authentication Method	Communication Protocol	Description	User Credential Encryption
Basic	HTTP (SSL optional)	Uses the server's built-in pop-up login dialog box.	None, unless using SSL.
Form-based	HTTP (SSL optional)	Application provides its own custom login and error pages.	None, unless using SSL.
Client Certificate	HTTPS (HTTP over SSL)	Server authenticates the client using a public key certificate.	SSL

Verifying Single Sign-On

Single sign-on enables multiple applications in one virtual server instance to share user authentication state. With single sign-on, a user who logs in to one application becomes implicitly logged in to other applications that require the same authentication information.

Single sign-on is based on groups. All Web applications whose deployment descriptor defines the same *group* and use the same authentication method (basic, form, digest, certificate) share single sign-on.

Single sign-on is enabled by default for virtual servers defined for the Application Server. For information on disabling single sign-on, see [“To configure single sign-on \(SSO\)” on page 186](#).

Authorizing Users

Once a user is authenticated, the level of *authorization* determines what operations can be performed. A user's authorization is based on his *role*. For example, a human resources application may authorize managers to view personal employee information for all employees, but allow employees to view only their own personal information. For more on roles, see [“Understanding Users, Groups, Roles, and Realms” on page 153](#).

Specifying JACC Providers

JACC (Java Authorization Contract for Containers) is part of the J2EE 1.4 specification that defines an interface for pluggable authorization providers. This enables the administrator to set up third-party plug-in modules to perform authorization.

By default, the Application Server provides a simple, file-based authorization engine that complies with the JACC specification. It is also possible to specify additional third-party JACC providers.

JACC providers use the Java Authentication and Authorization Service (JAAS) APIs. JAAS enables services to authenticate and enforce access controls upon users. It implements a Java technology version of the standard Pluggable Authentication Module (PAM) framework.

Auditing Authentication and Authorization Decisions

The Application Server can provide an audit trail of all authentication and authorization decisions through *audit modules*. The Application Server provides a default audit module, as well as the ability to customize the audit modules. For information on developing custom audit modules, see the Application Server *Developer's Guide*.

Configuring Message Security

Message Security enables a server to perform end-to-end authentication of web service invocations and responses at the message layer. The Application Server implements message security using message security providers on the SOAP layer. The message security providers provide information such as the type of authentication that is required for the request and response messages. The types of authentication that are supported include the following:

- Sender authentication, including username-password authentication.
- Content authentication, including XML Digital Signatures.

Two message security providers are included with this release. The message security providers can be configured for authentication for the SOAP layer. The providers that can be configured include `ClientProvider` and `ServerProvider`.

Support for message layer security is integrated into the Application Server and its client containers in the form of (pluggable) authentication modules. By default, message layer security is disabled on the Application Server.

Message level security can be configured for the entire Application Server or for specific applications or methods. Configuring message security at the Application Server level is discussed in [Chapter 10, “Configuring Message Security.”](#) Configuring message security at the application level is discussed in the *Developer's Guide* chapter titled *Securing Applications*.

Understanding Users, Groups, Roles, and Realms

The Application Server enforces its authentication and authorization policies upon the following entities:

- **“Users” on page 154:** An individual identity *defined in the Application Server*. In general, a user is a person, a software component such as an enterprise bean, or even a service. A user who has been authenticated is sometimes called a *principal*. Users are sometimes referred to as *subjects*.
- **“Groups” on page 154:** A set of users *defined in the Application Server*, classified by common traits.
- **“Roles” on page 155:** A named authorization level *defined by an application*. A role can be compared to a key that opens a lock. Many people might have a copy of the key. The lock doesn't care who seeks access, only that the right key is used.

- “Realms” on page 155: A repository containing user and group information and their associated security credentials. A realm is also called a *security policy domain*.

Note – Users and groups are designated for the entire Application Server, whereas each application defines its own roles. When the application is being packaged and deployed, the application specifies mappings between users/groups and roles, as illustrated in the following figure.

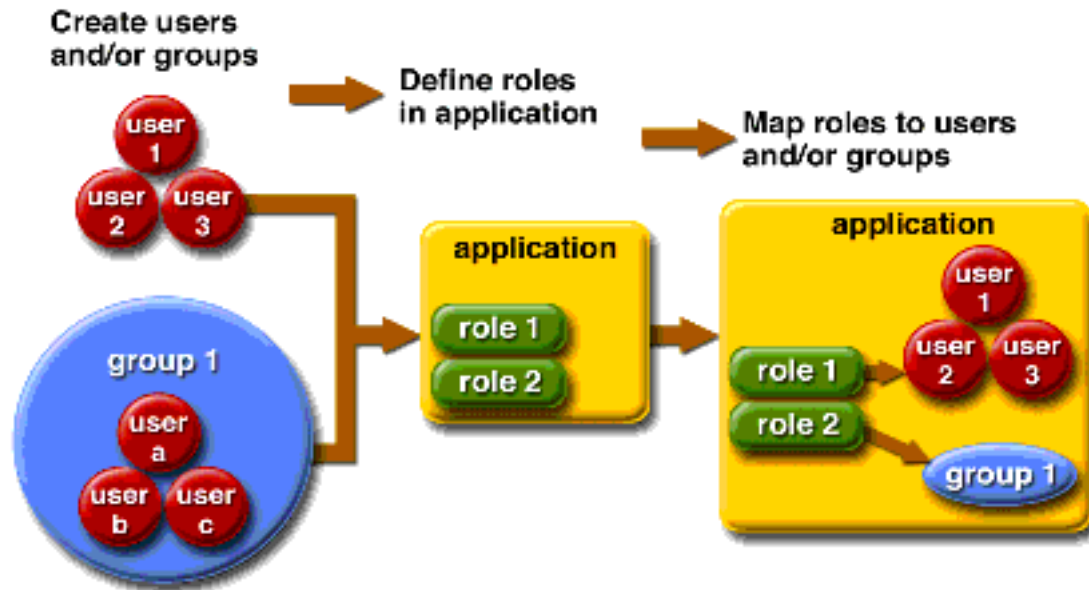


FIGURE 9-1 Role Mapping

Users

A *user* is an individual (or application program) identity that has been defined in the Application Server. A user can be associated with a group. The Application Server authentication service can govern users in multiple realms.

Groups

A *J2EE group* (or simply group) is a category of users classified by common traits, such as job title or customer profile. For example, users of an e-commerce application might belong to the customer group, but the big spenders would belong to the preferred group. Categorizing users into groups makes it easier to control the access of large numbers of users.

Roles

A *role* defines which applications and what parts of each application users can access and what they can do. In other words, roles determine users' authorization levels.

For example, in a personnel application all employees might have access to phone numbers and email addresses, but only managers would have access to salary information. The application might define at least two roles: `employee` and `manager`; only users in the `manager` role are allowed to view salary information.

A role is different from a user group in that a role defines a function in an application, while a group is a set of users who are related in some way. For example, in the personnel application there might be groups such as `full-time`, `part-time`, and `on-leave`, but users in all these groups would still be in the `employee` role.

Roles are defined in application deployment descriptors. In contrast, groups are defined for an entire server and realm. The application developer or deployer maps roles to one or more groups for each application in its deployment descriptor.

Realms

A *realm*, also called a *security policy domain* or *security domain*, is a scope over which the server defines and enforces a common security policy. In practical terms, a realm is a repository where the server stores user and group information.

The Application Server comes preconfigured with three realms: `file` (the initial default realm), `certificate`, and `admin-realm`. It is possible to also set up `ldap`, `solaris`, or `custom` realms. Applications can specify the realm to use in their deployment descriptor. If they do not specify a realm, the Application Server uses its default realm.

In the `file` realm, the server stores user credentials locally in a file named `keyfile`. You can use the Admin Console to manage users in the `file` realm. For more information, see [“Managing file Realm Users” on page 173](#).

In the `certificate` realm, the server stores user credentials in a certificate database. When using the `certificate` realm, the server uses certificates with the HTTPS protocol to authenticate Web clients. For more information about certificates, see [“Introduction to Certificates and SSL” on page 156](#).

The `admin-realm` is also a `FileRealm` and stores administrator user credentials locally in a file named `admin-keyfile`. Use the Admin Console to manage users in this realm in the same way you manage users in the `file` realm. For more information, see [“Managing file Realm Users” on page 173](#).

In the `ldap` realm the server gets user credentials from a Lightweight Directory Access Protocol (LDAP) server such as the Sun Java System Directory Server. LDAP is a protocol for enabling anyone to locate organizations, individuals, and other resources such as files and devices in a

network, whether on the public Internet or on a corporate intranet. Consult your LDAP server documentation for information on managing users and groups in the `ldap` realm.

In the `solaris` realm the server gets user credentials from the Solaris operating system. This realm is supported on the Solaris 9 OS and later. Consult your Solaris documentation for information on managing users and groups in the `solaris` realm.

A custom realm is any other repository of user credentials, such as a relational database or third-party component. For more information, see [“Creating a Custom Realm” on page 170](#).

Introduction to Certificates and SSL

The following topics are discussed in this section:

- [“About Digital Certificates” on page 156](#)
- [“About Secure Sockets Layer” on page 157](#)

About Digital Certificates

Digital certificates (or simply certificates) are electronic files that uniquely identify people and resources on the Internet. Certificates also enable secure, confidential communication between two entities.

There are different kinds of certificates, such as personal certificates, used by individuals, and server certificates, used to establish secure sessions between the server and clients through secure sockets layer (SSL) technology. For more information on SSL, see [“About Secure Sockets Layer” on page 157](#).

Certificates are based on *public key cryptography*, which uses pairs of digital *keys* (very long numbers) to *encrypt*, or encode, information so it can be read only by its intended recipient. The recipient then *decrypts* (decodes) the information to read it.

A key pair contains a public key and a private key. The owner distributes the public key and makes it available to anyone. But the owner never distributes the private key; it is always kept secret. Because the keys are mathematically related, data encrypted with one key can be decrypted only with the other key in the pair.

A certificate is like a passport: it identifies the holder and provides other important information. Certificates are issued by a trusted third party called a *Certification Authority* (CA). The CA is analogous to passport office: it validates the certificate holder’s identity and signs the certificate so that it cannot be forged or tampered with. Once a CA has signed a certificate, the holder can present it as proof of identity and to establish encrypted, confidential communications.

Most importantly, a certificate binds the owner’s public key to the owner’s identity. Like a passport binds a photograph to personal information about its holder, a certificate binds a public key to information about its owner.

In addition to the public key, a certificate typically includes information such as:

- The name of the holder and other identification, such as the URL of the Web server using the certificate, or an individual's email address.
- The name of the CA that issued the certificate.
- An expiration date.

Digital Certificates are governed by the technical specifications of the X.509 format. To verify the identity of a user in the certificate realm, the authentication service verifies an X.509 certificate, using the common name field of the X.509 certificate as the principal name.

About Certificate Chains

Web browsers are preconfigured with a set of *root* CA certificates that the browser automatically trusts. Any certificates from elsewhere must come with a *certificate chain* to verify their validity. A certificate chain is series of certificates issued by successive CA certificates, eventually ending in a root CA certificate.

When a certificate is first generated, it is a *self-signed* certificate. A self-signed certificate is one for which the issuer (signer) is the same as the subject (the entity whose public key is being authenticated by the certificate). When the owner sends a certificate signing request (CSR) to a CA, then imports the response, the self-signed certificate is replaced by a chain of certificates. At the bottom of the chain is the certificate (reply) issued by the CA authenticating the subject's public key. The next certificate in the chain is one that authenticates the CA's public key. Usually, this is a self-signed certificate (that is, a certificate from the CA authenticating its own public key) and the last certificate in the chain.

In other cases, the CA can return a chain of certificates. In this case, the bottom certificate in the chain is the same (a certificate signed by the CA, authenticating the public key of the key entry), but the second certificate in the chain is a certificate signed by a different CA, authenticating the public key of the CA to which you sent the CSR. Then, the next certificate in the chain is a certificate authenticating the second CA's key, and so on, until a self-signed *root* certificate is reached. Each certificate in the chain (after the first) thus authenticates the public key of the signer of the previous certificate in the chain.

About Secure Sockets Layer

Secure Sockets Layer (SSL) is the most popular standard for securing Internet communications and transactions. Web applications use HTTPS (HTTP over SSL), which uses digital certificates to ensure secure, confidential communications between server and clients. In an SSL connection, both the client and the server encrypt data before sending it, then decrypt it upon receipt.

When a Web browser (client) wants to connect to a secure site, an *SSL handshake* happens:

- The browser sends a message over the network requesting a secure session (typically, by requesting a URL that begins with `https` instead of `http`).

- The server responds by sending its certificate (including its public key).
- The browser verifies that the server's certificate is valid and is signed by a CA whose certificate is in the browser's database (and who is trusted). It also verifies that the CA certificate has not expired.
- If the certificate is valid, the browser generates a one time, unique *session key* and encrypts it with the server's public key. The browser then sends the encrypted session key to the server so that they both have a copy.
- The server decrypts the message using its private key and recovers the session key.

After the handshake, the client has verified the identity of the Web site, and only the client and the Web server have a copy of the session key. From this point forward, the client and the server use the session key to encrypt all their communications with each other. Thus, their communications are ensured to be secure.

The newest version of the SSL standard is called TLS (Transport Layer Security). The Application Server supports the Secure Sockets Layer (SSL) 3.0 and the Transport Layer Security (TLS) 1.0 encryption protocols.

To use SSL, the Application Server must have a certificate for each external interface, or IP address, that accepts secure connections. The HTTPS service of most Web servers will not run unless a digital certificate has been installed. Use the procedure described in [“To generate a certificate using the keytool utility” on page 193](#) to set up a digital certificate that your Web server can use for SSL.

About Ciphers

A *cipher* is a cryptographic algorithm used for encryption or decryption. SSL and TLS protocols support a variety of ciphers used to authenticate the server and client to each other, transmit certificates, and establish session keys.

Some ciphers are stronger and more secure than others. Clients and servers can support different cipher suites. Choose ciphers from the SSL3 and TLS protocols. During a secure connection, the client and the server agree to use the strongest cipher they both have enabled for communication, so it is usually sufficient to enable all ciphers.

Using Name-based Virtual Hosts

Using name-based virtual hosts for a secure application can be problematic. This is a design limitation of the SSL protocol itself. The SSL handshake, where the client browser accepts the server certificate, must occur before the HTTP request is accessed. As a result, the request information containing the virtual host name cannot be determined prior to authentication, and it is therefore not possible to assign multiple certificates to a single IP address.

If all virtual hosts on a single IP address need to authenticate against the same certificate, the addition of multiple virtual hosts probably will not interfere with normal SSL operations on the server. Be aware, however, that most browsers will compare the server's domain name against

the domain name listed in the certificate, if any (applicable primarily to official, CA-signed certificates). If the domain names do not match, these browsers display a warning. In general, only address-based virtual hosts are commonly used with SSL in a production environment.

About Firewalls

A *firewall* controls the flow of data between two or more networks, and manages the links between the networks. A firewall can consist of both hardware and software elements. This section describes some common firewall architectures and their configuration. The information here pertains primarily to the Application Server. For details about a specific firewall technology, refer to the documentation from your firewall vendor.

In general, configure the firewalls so that clients can access the necessary TCP/IP ports. For example, if the HTTP listener is operating on port 8080, configure the firewall to allow HTTP requests on port 8080 only. Likewise, if HTTPS requests are setup for port 8181, you must configure the firewalls to allow HTTPS requests on port 8181.

If direct Remote Method Invocations over Internet Inter-ORB Protocol (RMI-IIOP) access from the Internet to EJB modules are required, open the RMI-IIOP listener port as well, but this is strongly discouraged because it creates security risks.

In double firewall architecture, you must configure the outer firewall to allow for HTTP and HTTPS transactions. You must configure the inner firewall to allow the HTTP server plug-in to communicate with the Application Server behind the firewall.

Managing Security With the Admin Console

The Admin Console provides the means to manage the following aspects of security:

- “Server Security Settings” on page 159
- “Realms and file Realm Users” on page 160
- “JACC Providers” on page 160
- “Audit Modules” on page 160
- “Message Security” on page 160
- “HTTP and IIOB Listener Security” on page 161
- “Admin Service Security” on page 161
- “Security Maps” on page 161

Server Security Settings

On the Security Settings page, set properties for the entire server, including specifying the default realm, the anonymous role, and the default principal user name and password. For more information, see “To configure security settings” on page 162.

Realms and file Realm Users

The concept of realms was introduced in [“Understanding Users, Groups, Roles, and Realms” on page 153](#).

- Create a new realm
- Delete an existing realm
- Modify the configuration of an existing realm
- Add, modify, and delete users in the file realm
- Set the default realm

See [“Admin Console Tasks for Realms” on page 164](#) for details on these tasks.

JACC Providers

JACC providers were introduced in [“Specifying JACC Providers” on page 152](#). Use the Admin Console to perform the following tasks:

- Add a new JACC provider
- Delete or modify an existing JACC provider

See [“Admin Console Tasks for JACC Providers” on page 176](#) for details on these tasks.

Audit Modules

Audit modules were introduced in [“Auditing Authentication and Authorization Decisions” on page 153](#). Auditing is the method by which significant events, such as errors or security breaches, are recorded for subsequent examination. All authentication events are logged to the Application Server logs. A complete access log provides a sequential trail of Application Server access events.

Use the Admin Console to perform the following tasks:

- Add a new audit module
- Delete or modify an existing audit module

See [“Admin Console Tasks for Audit Modules” on page 179](#) for details on these tasks.

Message Security

The concept of message security was introduced in [“Configuring Message Security” on page 153](#). Use the Admin Console to perform the following tasks:

- Enable message security
- Configure a message security provider
- Delete or configure an existing message security configuration or provider

See [Chapter 10, “Configuring Message Security,”](#) for details on these tasks.

HTTP and IIOP Listener Security

Each virtual server in the HTTP service provides network connections through one or more *HTTP listeners*. For general information about the HTTP service and HTTP listeners, see [“What Is the HTTP Service?” on page 229](#).

The Application Server supports CORBA (Common Object Request Broker Architecture) objects, which use the Internet Inter-Orb Protocol (IIOP) to communicate across the network. An *IIOP listener* accepts incoming connections from remote clients of EJB components and from other CORBA-based clients. For general information on IIOP listeners, see [“IIOP Listeners” on page 250](#).

With the Admin Console, perform the following tasks:

- Create a new HTTP or IIOP listener, and specify the security it uses.
- Modify the security settings for an existing HTTP or IIOP listener.

See [“Admin Console Tasks for Listeners and JMX Connectors” on page 183](#) for details on these tasks.

Admin Service Security

The Admin Service determines whether the server instance is a regular instance, a domain administration server (DAS), or a combination. Use the Admin Service to configure a JSR-160 compliant remote JMX connector, which handles communication between the domain administration server and the node agents, which manage server instances on a host machine, for remote server instances.

With the Admin Console, perform the following tasks:

- Manage the Admin Service
- Edit the JMX connector
- Modify the security settings of the JMX connector

See [“To configure security for the Admin Service’s JMX connector” on page 185](#) for details on these tasks.

Security Maps

The concept of security maps for connector connection pools is introduced in [“About Security Maps” on page 188](#). Use the Admin Console to perform the following tasks:

- Add a security map to an existing connector connection pool
- Delete or configure an existing security map

See [“Admin Console Tasks for Connector Connection Pools” on page 187](#) for details on these tasks.

Admin Console Tasks for Security

- “To configure security settings” on page 162
- “To grant access to administration tools” on page 163
- “Configuring Mutual Authentication” on page 175
- “To configure single sign-on (SSO)” on page 186

▼ To configure security settings

The Security page in the Admin Console enables you to set a variety of system-wide security settings.

1 In the Admin Console tree component, expand the Configurations node.

2 Select the instance to configure:

- **To configure a particular instance, expand the instance’s config node. For example, the default instance, server, expand the server-config node.**
- **To configure the default settings for all instances, expand the default-config node.**

3 Select the Security node.

The Security page displays.

4 Modify the values as necessary.

The general security options are discussed in the following table.

Setting	Description
Audit Logging	Select to enable audit logging. If enabled, the server will load and run all the audit modules specified in the Audit Modules setting. If disabled, the server does not access audit modules. Disabled by default.
Default Realm	The active (default) realm the server uses for authentication. Applications use this realm unless they specify a different realm in their deployment descriptor. All configured realms appear in the list. The initial default realm is the <code>file</code> realm.
Anonymous Role	The name for the default or anonymous role. The anonymous role is assigned to all users. Applications can use this role in their deployment descriptors to grant authorization to anyone.

Setting	Description
Default Principal	Specifies the default user name. The server uses this when no principal is provided. If you enter a value in this field, enter a corresponding value in the Default Principal Password field. This attribute is not required for normal server operation.
Default Principal Password	Password of the default principal specified in the Default Principal field. This attribute is not required for normal server operation.
JACC	Class name of a configured JACC provider. See “To create a JACC provider” on page 176
Audit Modules	List of audit module provider classes, delimited by commas. A module listed here must already be configured. If Audit Logging is enabled, this setting must list audit modules. By default, the server uses an audit module named default. For information on creating new audit modules, see “To create an audit module” on page 179 .

5 Enter additional properties to pass to the Java Virtual Machine (JVM) in the Additional Properties section.

Valid properties are dependent upon the type of realm selected in the Default Realm field. Valid properties are discussed in the following sections:

- [“Editing the file and admin-realm Realms” on page 172](#)
- [“Editing the certificate Realm” on page 171](#)
- [“Creating the solaris Realm” on page 170](#)
- [“Creating an ldap Realm” on page 168](#)
- [“Creating a Custom Realm” on page 170](#)

6 Select Save to save the changes or Load Defaults to restore the default values.

▼ To grant access to administration tools

Only users in the `asadmin` group are able to access Admin Console and the `asadmin` command line utility.

To give a user access to these administration tools, add them to the `asadmin` group in the `admin-realm`.

1 In the Admin Console tree component, expand the Configurations node.

2 Select the instance to configure:

- To configure a particular instance, expand the instance’s config node. For example, the default instance, `server`, expand the `server-config` node.

- **To configure the default settings for all instances, expand the `default-config` node.**
- 3 Expand the Security node.**
- 4 Expand the Realms node.**
- 5 Select the `admin-realm` node.**
- 6 Click the Manage Users button from the Edit Realm page.**

Initially after installation, the administrator user name and password entered during installation are listed in a file named `admin-keyfile`. By default, this user belongs to the group `asadmin`, which gives rights to modify the Application Server. Assign users to this group only if you want to grant them administrator privileges for the Application Server.

If you add users to the `admin-realm` realm, but assign the user to a group other than `asadmin`, the user information will still be written to the file named `admin-keyfile`, but the user will have no access to administrative tools or to applications in the `file` realm.
- 7 Click New to add a new user to the `admin-realm` realm.**
- 8 Enter the correct information into the User ID, Password, and Group List fields.**

To authorize a user to make modifications to the Application Server, include the `asadmin` group in the Group List.
- 9 Click OK to add this user to the `admin-realm` realm or click Cancel to quit without saving.**

Admin Console Tasks for Realms

- [“To create a realm” on page 164](#)
- [“To edit a realm” on page 166](#)
- [“To delete a realm” on page 167](#)
- [“To set the default realm” on page 167](#)
- [“Additional Information for Specific Realms” on page 168](#)

▼ To create a realm

The Application Server comes preconfigured with three realms: `file`, `certificate`, and `admin-realm`. It is also possible to create `ldap`, `solaris`, and `custom` realms. Generally, you will have one realm of each type on a server, but on the Application Server there are two `file` realms: `file` and `admin-realm`. These are two realms of the same type used for two different purposes. It is also possible to have a different certificate database for each virtual server on your system.

- 1 In the Admin Console tree component, expand the Configurations node.**

2 Select the instance to configure:

- To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
- To configure the default settings for all instances, expand the default-config node.

3 Expand the Security node.**4 Select the Realms node.****5 On the Realms page, click New.**

The Create Realm page is displayed.

6 Enter a name for the realm in the Name field.**7 Specify the class name for the realm being created.**

Valid choices are shown in the following table.

Realm Name	Class Name
file	com.sun.enterprise.security.auth.realm.file.FileRealm
certificate	com.sun.enterprise.security.auth.realm.certificate.CertificateRealm
ldap	com.sun.enterprise.security.auth.realm.ldap.LDAPRealm
solaris	com.sun.enterprise.security.auth.realm.solaris.SolarisRealm
custom	Name of login realm class

8 Add the required properties and any desired optional properties for the realm.**a. Click Add Property.****b. In the Name field, enter the name of the property.**

- For a description of file realm properties, see [“Editing the file and admin-realm Realms” on page 172.](#)
- For a description of certificate realm properties, see [“Editing the certificate Realm” on page 171.](#)
- For a description of ldap realm properties, see [“Creating an ldap Realm” on page 168.](#)
- For a description of solaris realm properties, see [“Creating the solaris Realm” on page 170.](#)
- For a description of custom realm properties, see [“Creating a Custom Realm” on page 170.](#)

c. In the **Value** field, enter the value of the property.

9 Click **OK**.

More Information Equivalent `asadmin` command

```
create-auth-realm
```

▼ To edit a realm

1 In the **Admin Console** tree component, expand the **Configurations** node.

2 Select the instance to configure:

- To configure a particular instance, expand the instance's **config** node. For example, the default instance, `server`, expand the `server-config` node.
- To configure the default settings for all instances, expand the `default-config` node.

3 Expand the **Security** node.

4 Expand the **Realms** node.

5 Select the name of an existing realm.

The **Edit Realm** page displays.

6 Edit existing properties and their values as desired.

7 To add additional properties, click the **Add Properties** button.

The page displays a new row. Enter a valid property name and property value.

- For a description of `file` realm properties, see [“Editing the file and admin-realm Realms” on page 172](#).
- For a description of `certificate` realm properties, see [“Editing the certificate Realm” on page 171](#).
- For a description of `ldap` realm properties, see [“Creating an ldap Realm” on page 168](#).
- For a description of `solaris` realm properties, see [“Creating the solaris Realm” on page 170](#).
- For a description of `custom` realm properties, see [“Creating a Custom Realm” on page 170](#).

8 Click **Save** to save the changes.

▼ To delete a realm

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Expand the Security node.
- 4 Select the Realms node.
- 5 Click in the box beside the realm to be deleted.
- 6 Click Delete.

More Information Equivalent `asadmin` command

```
delete-auth-realm
```

▼ To set the default realm

The *default realm* is the realm that the Application Server uses for authentication and authorization if an application's deployment descriptor does not specify a realm.

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Select the Security node.

The Security page displays.
- 4 In the Default Realm field, pick the desired realm from the drop-down list.

- 5 Click **Save** to save the changes or **Load Defaults** to delete changes and restore the Application Server default values.
- 6 Restart the server if **Restart Required** displays in the console.

Additional Information for Specific Realms

This section covers the following topics:

- “Creating an ldap Realm” on page 168
- “Creating the solaris Realm” on page 170
- “Creating a Custom Realm” on page 170
- “Editing the certificate Realm” on page 171
- “Editing the file and admin-realm Realms” on page 172
- “Managing Users with Network Security Services (NSS)” on page 172
- “Managing file Realm Users” on page 173
- “Configuring Mutual Authentication” on page 175

Creating an ldap Realm

The ldap realm performs authentication using information from an LDAP server. User information includes user name, password, and the groups to which the user belongs. To use an LDAP realm, the users and groups must already be defined in your LDAP directory.

To create an LDAP realm, follow the steps in “[To create a realm](#)” on page 164 for adding a new realm, and add the properties shown in the following table.

TABLE 9-2 Required properties for ldap realm

Property Name	Description	Value
directory	LDAP URL of the directory server.	LDAP URL of the form <code>ldap://hostname:port</code> For example, <code>ldap://myldap.foo.com:389</code> .
base-dn	Base Distinguished Name (DN) for the location of user data, which can be at any level above the user data, since a tree scope search is performed. The smaller the search tree, the better the performance.	Domain for the search, for example: <code>dc=siliconvalley, dc=BayArea, dc=sun, dc=com</code> .
jaas-context	Type of login module to use for this realm.	Must be <code>ldapRealm</code> .

Optional properties for the ldap realm are shown in the following table.

TABLE 9-3 Optional properties for ldap realm

Property Name	Description	Default
search-filter	Search filter to use to find the user.	uid=%s (%s expands to the subject name).
group-base-dn	Base DN for the location of group data.	Same as the base-dn, but it can be tuned if necessary.
group-search-filter	Search filter to find group memberships for the user.	uniqueMember=%d (%d expands to the user element DN).
group-target	LDAP attribute name that contains group name entries.	CN
search-bind-dn	Optional DN used to authenticate to the directory for performing the search-filter lookup. Only required for directories that do not allow anonymous search.	
search-bind-password	LDAP password for the DN given in search-bind-dn.	

Example

For example, suppose an LDAP user, Joe Java, is defined in the LDAP directory as follows:

```
uid=jjava,ou=People,dc=acme,dc=com
uid=jjava
givenName=joe
objectClass=top
objectClass=person
objectClass=organizationalPerson
objectClass=inetorgperson
sn=java
cn=Joe Java
```

Using the example code, when creating or editing the ldap realm, you can enter the values as shown in the following table.

TABLE 9-4 Example ldap realm values

Property Name	Property Value
directory	LDAP URL to your server, for example: ldap://ldap.acme.com:389

TABLE 9-4 Example ldap realm values (Continued)

Property Name	Property Value
base-dn	ou=People, dc=acme, dc=com. Can be rooted higher, for example dc=acme, dc=com, but searches would traverse a larger part of the tree, reducing performance.
jaas-context	ldapRealm

Creating the solaris Realm

The `solaris` realm gets user and group information from the underlying Solaris user database, as determined by the system's configuration. The `solaris` realm invokes the underlying PAM infrastructure for authenticating. If the configured PAM modules require root privileges, the domain must run as root to use this realm. For details, see the Solaris documentation for security services.

The `solaris` realm has one required property, `jaas-context` that specifies the type of login module to use. The property value must be `solarisRealm`.

Note – The `solaris` realm is supported only for Solaris 9 or later.

Creating a Custom Realm

In addition to the four built-in realms, you can also create custom realms that store user data in some other way, such as in a relational database. Development of a custom realm is outside the scope of this document. For more information, see the *Application Server Developer's Guide* chapter titled *Securing Applications*.

As an administrator, the main thing you need to know is that a custom realm is implemented by a class (called the `LoginModule`) derived from the Java Authentication and Authorization Service (JAAS) package.

▼ To create a custom realm

- 1 Follow the procedure outline in [“To create a realm” on page 164](#), entering the name of the custom realm and the name of the `LoginModule` class.

Any unique name can be used for the custom realm, for example `myCustomRealm`.

- 2 Add the properties for a custom realm shown in the following table.

Property Name	Property Value
jaas-context	LoginModule class name, for example <code>simpleCustomRealm</code>
auth-type	Description of the realm, for example “A simple example custom realm”.

- 3 Click OK.
- 4 Edit the domain's login configuration file, *domain-dir/config/login.conf*, and add the fully-qualified class name of the JAAS LoginModule at the end of the file, as follows:

```
realmName {
    fully-qualified-LoginModule-classname required;
};
```

For example,

```
myCustomRealm {
    com.foo.bar.security.customrealm.simpleCustomLoginModule required;
};
```

- 5 Copy the LoginModule class and all dependent classes into the directory *domain-dir/lib/classes*.
- 6 Restart the Server if Restart Required displays in the console.
- 7 Make sure that the realm is properly loaded.

Check *domain-dir/logs/server.log* to make sure the server loaded the realm. The server should invoke the realm's `init()` method.

Editing the certificate Realm

The `certificate` realm supports SSL authentication. This realm sets up the user identity in the Application Server's security context, and populates it with user data obtained from cryptographically verified client certificates in the truststore and keystore files (see “[About Certificate Files](#)” on page 191). Add users to these files using `certutil`.

With the `certificate` realm, J2EE containers handle authorization processing based on each user's Distinguished Name (DN) from his or her certificate. The DN is the name of the entity whose public key the certificate identifies. This name uses the X.500 standard, so it is intended to be unique across the Internet. For more information on key stores and trust stores, refer to the `certutil` documentation at “[Using Network Security Services \(NSS\) Tools](#)” on page 196.

The following table lists the optional properties for the `certificate` realm.

TABLE 9-5 Optional properties for certificate realm

Property	Description
assign-groups	A comma-separated list of group names. All clients who present valid certificates are assigned to these groups. For example, <code>employee,manager</code> , where these are the names of user groups.
jaas-context	Type of login module to use for this realm. For the certificate realm, the value must be <code>certificateRealm</code> .

Editing the file and admin-realm Realms

The server maintains all user, group, and password information in a file named `keyfile` for the `file` realm and `admin-keyfile` for the `admin-realm`. For both, the `file` property specifies the location of the keyfile. The following table shows required properties for a `file` realm.

TABLE 9-6 Required properties for file realms

Property name	Description	Default Value
file	Full path and name of the keyfile.	<code>domain-dir/config/keyfile</code>
jaas-context	Type of login module to use for this realm.	<code>fileRealm</code> is the only valid value

The `keyfile` is initially empty, so users must be added before the `file` realm is used. For instructions, see [“Managing file Realm Users” on page 173](#).

The `admin-keyfile` initially contains the admin user name, the admin password in an encrypted format, and the group to which this user belongs, which is `asadmin` by default. For more information on adding users to the `admin-realm`, read [“To grant access to administration tools” on page 163](#).

Note – Users in the group `asadmin` in the `admin-realm` are authorized to use the Admin Console and `asadmin` tools. Add only users to this group that have server administrative privileges.

Managing Users with Network Security Services (NSS)

In the **Enterprise Edition only**, you can manage users using the Admin Console as discussed in [“Managing file Realm Users” on page 173](#) or you can manage users using NSS tools. Network Security Services (NSS) is a set of libraries designed to support cross-platform development of security-enabled client and server applications. Applications built with NSS can support SSL v2 and v3, TLS, PKCS #5, PKCS #7, PKCS #11, PKCS #12, S/MIME, X.509 v3 certificates, and other security standards. For detailed information, link to the following URLs:

- Network Security Services (NSS) at <http://www.mozilla.org/projects/security/pki/nss/>

- NSS Security Tools at <http://www.mozilla.org/projects/security/pki/nss/tools/>
- Overview of NSS at <http://www.mozilla.org/projects/security/pki/nss/overview.html>

Managing file Realm Users

Manage file realm users with the Admin Console. Users and groups in the file realm are listed in the keyfile, whose location is specified by the file property.

Note – It is also possible to use these steps to add users to any file realm, including the `admin-realm`. Simply substitute the name of the target realm in place of the file realm referenced in this section.

A user in the file realm can belong to a *J2EE group*, a category of users classified by common traits. For example, customers of an e-commerce application might belong to the `CUSTOMER` group, but the big spenders would belong to the `PREFERRED` group. Categorizing users into groups makes it easier to control the access of large numbers of users.

Initially after installation of the Application Server, the only user is the administrator entered during installation. By default, this user belongs to the group `asadmin`, in the realm `admin-realm`, which gives rights to modify the Application Server. Any users assigned to this group will have administrator privileges, that is, they will have access to the `asadmin` tool and the Admin Console.

To manage file realm users, perform these tasks:

- “To access the File Users page” on page 173
- “To add a user” on page 174
- “To edit user information” on page 174
- “To delete a user” on page 175

▼ To access the File Users page

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance’s config node. For example, the default instance, `server`, expand the `server-config` node.
 - To configure the default settings for all instances, expand the `default-config` node.
- 3 Expand the Security node.
- 4 Expand the Realms node.

5 **Select the file node.**

6 **Click the Manage Users button from the Edit Realm page.**

The File Users page displays. In this page, perform the following tasks:

- [“To add a user” on page 174](#)
- [“To edit user information” on page 174](#)
- [“To delete a user” on page 175](#)

▼ **To add a user**

1 **Click New to add a new user to the file realm.**

2 **Enter the following information on the File Users page:**

- **User ID** (*required*) - The name of the user.
- **Password** (*required*) - The user password.
- **Confirm Password** (*required*) - The user password again, for verification.
- **Group List** (*optional*) - A comma-separated list of the groups to which the user belongs. These groups do not need to be defined elsewhere.

3 **Click OK to add this user to the list of users in the file realm. Click Cancel to quit without saving.**

More Information Equivalent asadmin command

```
create-file-user
```

▼ **To edit user information**

1 **In the User ID column, click the name of the user to be modified.**

The Edit File Realm User page displays.

2 **Change the user password by entering a new password in the Password and Confirm Password fields.**

3 **Change the groups to which the user belongs by adding or deleting groups in the Group List field.**

Separate group names with commas. Groups need not be previously defined.

4 **Click Save to save this user to the list of users in the file realm.**

Click Close to quit without saving.

▼ To delete a user

- 1 Select the checkbox to the left of the name of the user(s) to be deleted.
- 2 Click Delete.
- 3 Click Close to return to the Edit Realm page.

More Information Equivalent `asadmin` command

```
delete-file-user
```

Configuring Mutual Authentication

- “Enabling Mutual SSL Authentication in an Application” on page 175
- “To enable mutual authentication for all the applications” on page 175

In mutual authentication, both server and client-side authentication are enabled. To test mutual authentication, a client with a valid certificate must exist. For information on mutual authentication, see the *Security* chapter of *The J2EE 1.4 Tutorial* at <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.

Enabling Mutual SSL Authentication in an Application

To enable mutual authentication for a specific application, use `deploytool` to set the method of authentication to `Client-Certificate`. For more information about using `deploytool`, refer to the *Security* chapter of *The J2EE 1.4 Tutorial* at <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.

▼ To enable mutual authentication for all the applications

The Application Server uses the `certificate` realm for HTTPS authentication.

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance’s config node. For example, the default instance, `server`, expand the `server-config` node.
 - To configure the default settings for all instances, expand the `default-config` node.
- 3 Expand the Security node.
- 4 Expand the Realms node.

- 5 **Select the certificate realm.**
- 6 **Click the Add Property button.**
 - a. In the Name field, enter `clientAuth`.
 - b. In the Value field, enter `true`.
- 7 **Click Save.**
- 8 **Restart the Application Server if Restart Required displays in the console.**

After restarting the server, client authentication is required for all applications that use the certificate realm.

Admin Console Tasks for JACC Providers

- [“To create a JACC provider” on page 176](#)
- [“To edit a JACC provider” on page 177](#)
- [“To delete a JACC provider” on page 178](#)
- [“To set the active JACC provider” on page 178](#)

▼ To create a JACC provider

JACC (Java Authorization Contract for Containers) is part of the J2EE 1.4 specification that defines an interface for pluggable authorization providers. This enables the administrator to set up third-party *plug in* modules to perform authorization. By default, the Application Server provides a simple, JACC-compliant file-based authorization engine.

- 1 **In the Admin Console tree component, expand the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, expand the instance’s config node. For example, the default instance, `server`, expand the `server-config` node.**
 - **To configure the default settings for all instances, expand the `default-config` node.**
- 3 **Expand the Security node.**
- 4 **Select the JACC Providers node.**
- 5 **On the JACC Providers page, click New.**

- 6 On the **Create JACC Provider** page, enter the following:
 - **Name** – The name to use to identify this provider.
 - **Policy Configuration** – The name of the class that implements the policy configuration factory. The default provider uses `com.sun.enterprise.security.provider.PolicyConfigurationFactoryImpl`.
 - **Policy Provider** – The name of the class that implements the policy factory. The default provider uses `com.sun.enterprise.security.provider.PolicyWrapper`.
- 7 Add properties to the provider by clicking the **Add Property** button. Valid properties include:
 - `repository` – the directory that contains the policy file. For the default provider, this value is `${com.sun.aas.instanceRoot}/generated/policy`.
- 8 Click **OK** to save this configuration, or click **Cancel** to quit without saving.

▼ To edit a JACC provider

- 1 In the Admin Console tree component, expand the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's `config` node. For example, the default instance, `server`, expand the `server-config` node.
 - To configure the default settings for all instances, expand the `default-config` node.
- 3 Expand the **Security** node.
- 4 Expand the **JACC Providers** node.
- 5 Select the node of the JACC provider to be edited.
- 6 On the **Edit JACC Provider** page, modify the provider information as desired:
 - **Policy Configuration** – The name of the class that implements the policy configuration factory.
 - **Policy Provider** – The name of the class that implements the policy factory.
- 7 To add properties, click the **Add** button. Enter the name and value for the property. Valid entries include:
 - `repository` – the directory that contains the policy file. For the default provider, this value is `${com.sun.aas.instanceRoot}/generated/policy`.

- 8 To delete an existing property, click in the checkbox to the left of the property, then click Delete Properties.
- 9 Click Save to save or click the browser's back button to cancel without saving.

▼ To delete a JACC provider

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Expand the Security node.
- 4 Select the JACC Providers node.
- 5 Click in the checkbox to the left of the JACC provider to be deleted.
- 6 Click Delete.

▼ To set the active JACC provider

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Select the Security node.

The Security page displays.
- 4 In the JACC field, enter the name of the JACC provider to be used by the server.

If you don't know which JACC providers are available, expand the JACC Provider component in the tree to view all configured JACC providers.

- 5 Select **Save** to save the changes or **Load Defaults** to return to the default values.
- 6 Restart the Application Server if **Restart Required** displays in the console.

Admin Console Tasks for Audit Modules

- “To create an audit module” on page 179
- “To edit an audit module” on page 180
- “To delete an audit module” on page 181
- “To set the active audit module” on page 182
- “To enable or disable audit logging” on page 181

▼ To create an audit module

The Application Server provides a simple default audit module; for more information, see “To use the default audit module” on page 182.

- 1 In the Admin Console tree component, expand the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance’s **config** node. For example, the default instance, **server**, expand the **server-config** node.
 - To configure the default settings for all instances, expand the **default-config** node.
- 3 Expand the **Security** node.
- 4 Select the **Audit Modules** node.
- 5 On the **Audit Modules** page, click **New**.
- 6 On the **Create Audit Module** page, enter the following information:
 - **Name** – The name used to identify this audit module.
 - **Classname** – The fully-qualified name of the class that implements this module. The class name for the default audit module is `com.sun.enterprise.security.Audit`.
- 7 To add **JVM** properties to this module, click **Add Property**. Specify a name and value for each property. Valid properties include:
 - **auditOn** - Specifies whether or not to enable this implementation class. Valid values are `true` and `false`.

- 8 Click OK to save entries, or click Cancel to quit without saving.

▼ To edit an audit module

Audit modules are not turned on by default. For more information on how to activate audit modules, read [“To enable or disable audit logging” on page 181](#).

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance’s config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Expand the Security node.
- 4 Expand the Audit Modules node.
- 5 Click the node of the audit module to be edited.
- 6 On the Edit Audit Module page, modify the class name, if needed.
- 7 Enter any additional properties for the module by selecting the Add button and entering the name and value of the property. Valid properties include:
 - auditOn - Specifies whether or not to use this audit module. Valid values are true and false.
- 8 Modify any existing properties by selecting the name or value to be modified, and entering the changes directly into the text field.
- 9 Delete a property by selecting the checkbox to the left of the property and clicking Delete Properties.
- 10 Click Save to save or click the Back button on the browser to cancel without saving.

▼ To delete an audit module

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Expand the Security node.
- 4 Select the Audit Modules node.
- 5 Click in the checkbox to the left of the audit module to be deleted.
- 6 Click Delete.

▼ To enable or disable audit logging

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.
- 3 Select the Security node.

The Security page displays.
- 4 To enable logging, select the Audit Logging check box. To disable it, deselect it.

Selecting this option causes the loading of the audit modules and ensures they are called by the Application Server's audit library at audit points.
- 5 If you are enabling audit logging, specify a default audit module as described in ["To set the active audit module" on page 182](#).
- 6 Select Save to save the changes.

- 7 Restart the Application Server if Restart Required displays in the console.

▼ To set the active audit module

Before You Begin To specify the audit module that the server uses, first enable audit logging as described in [“To enable or disable audit logging” on page 181](#).

- 1 **In the Audit Modules field, enter the name of the audit module to be used by the server.**
The preconfigured audit module is called `default`. Make sure that this audit module has `auditOn` set to true as described in [“To use the default audit module” on page 182](#).
- 2 **Select Save to save the changes, Load Defaults to cancel.**
- 3 **Restart the Application Server if Restart Required displays in the console.**

▼ To use the default audit module

The `default` audit module logs authentication and authorization requests to the server log file. For information on changing the location of the log file, see [“To configure general logging settings” on page 262](#).

Authentication log entries include the following information:

- Names of users who attempted to authenticate.
- The realm that processed the access request.
- The requested Web module URI or EJB component.
- Success or failure of the request.

Regardless of whether audit logging is enabled, the Application Server logs all denied authentication events.

Authorization log entries include the following information:

- Names of authenticated users, if any.
- The requested Web URI or EJB component.
- Success or failure of the requests.

- 1 **In the Admin Console tree component, expand the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, expand the instance’s config node. For example, the default instance, `server`, expand the `server-config` node.**
 - **To configure the default settings for all instances, expand the `default-config` node.**

- 3 Expand the Security node.
- 4 Expand the Audit Modules node.
- 5 Click the default node.
- 6 Set the value of the `auditOn` property to `true`.
- 7 Select Save to save the changes.
- 8 Restart the Application Server if Restart Required displays in the console.

Admin Console Tasks for Listeners and JMX Connectors

- “To configure security for HTTP listeners” on page 183
- “To configure security for IIOP listeners” on page 184
- “To configure security for the Admin Service’s JMX connector” on page 185
- “To set listener security properties” on page 185

▼ To configure security for HTTP listeners

Each virtual server in the HTTP service provides network connections through one or more *HTTP listeners*. With the Admin Console, create new HTTP listeners and edit the security settings of existing HTTP listeners.

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance’s config node. For example, the default instance, `server`, expand the `server-config` node.
 - To configure the default settings for all instances, expand the `default-config` node.
- 3 Expand the HTTP Service node.
- 4 Select the HTTP Listeners node.
- 5 Select an HTTP listener to edit an existing listener or click New and follow the procedure in “To create an HTTP listener” on page 244 to create a new listener.
- 6 Follow the procedure in “To set listener security properties” on page 185 to set security properties.

- 7 Click **Save** to save the changes, or click the browser's **Back** button to cancel without saving.

More Information Equivalent `asadmin` command

```
create-http-listener
```

▼ To configure security for IIOP listeners

The Application Server supports CORBA (Common Object Request Broker Architecture) objects, which use the Internet Inter-Orb Protocol (IIOP) to communicate across the network. An *IIOP listener* accepts incoming connections from remote clients of EJB components and from other CORBA-based clients. With the Admin Console, create new IIOP listeners and edit the settings of existing IIOP listeners.

- 1 In the **Admin Console** tree component, expand the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's **config** node. For example, the default instance, `server`, expand the `server-config` node.
 - To configure the default settings for all instances, expand the `default-config` node.
- 3 Expand the **ORB** node.
- 4 Select the **IIOP Listeners** node.
- 5 Select an IIOP listener to edit that listener or click **New** and follow the procedure in [“To create an IIOP listener” on page 251](#) to create a new listener.
- 6 Follow the procedure in [“To set listener security properties” on page 185](#) to set security properties.
- 7 Click **Save** to save the changes, or click **Load Defaults** to restore the properties to their default values.

If a new listener was created, it will now be listed in the **Current Listeners** table on the **IIOP Listeners** page.

More Information Equivalent `asadmin` command

```
create-iiop-listener
```


▼ To configure security for the Admin Service's JMX connector

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, expand the instance's config node. For example, the default instance, `server`, expand the `server-config` node.
 - To configure the default settings for all instances, expand the `default-config` node.
- 3 Expand the Admin Service node.
- 4 Select the admin service to be modified.
- 5 Follow the procedure in [“To set listener security properties” on page 185](#) to set security properties.
- 6 Click Save to save the changes, or click Load Defaults to restore the properties to their default values.

▼ To set listener security properties

This procedure applies to HTTP listener, IIOP listener, and JMX Connector security properties.

- 1 In the Edit HTTP Listener, Edit IIOP Listener, or Edit JMX Connector page, go to the section labeled SSL.
- 2 Check the Enabled box in the Security field to enable security for this listener. When this option is selected, you must select SSL3 or TLS to specify which type of security is enabled, and you must enter a certificate nickname.
- 3 Check the Enabled box in the Client Authentication field if clients are to authenticate themselves to the Application Server when using this listener.
- 4 Enter the keystore alias in the Certificate Nickname field if the Enabled box is checked. The keystore alias is a single value that identifies an existing server key pair and certificate. The certificate nickname for the default keystore is `slas`.

To find the Certificate Nickname, use the `certutil` utility, as described in [“Using Network Security Services \(NSS\) Tools” on page 196](#).

- 5 Select SSL3 and/or TLS if the Enabled box is checked. By default, both SSL3 and TLS are enabled.
- 6 Enable individual cipher suites, if needed. By default, all supported cipher suites are enabled. Ciphers are discussed in [“About Ciphers” on page 158](#).
- 7 Select Save to save the changes or Load Defaults to cancel.

Admin Console Security Tasks for Virtual Servers

- [“To configure single sign-on \(SSO\)” on page 186](#)

▼ To configure single sign-on (SSO)

Single sign-on enables multiple applications to share user sign-on information, rather than requiring each application to have separate user sign-on. Applications using single sign-on authenticate the user one time, and the authentication information is propagated to all other involved applications.

Single sign-on applies to Web applications configured for the same realm and virtual server.

Note – Single sign-on uses an HTTP cookie to transmit a token that associates each request with the saved user identity, so it can be used only when the browser client supports cookies.

Single sign-on operates according to the following rules:

- When a user accesses a protected resource in a Web application, the server requires the user to authenticate himself or herself, using the method defined for that Web application.
 - Once authenticated, the Application Server uses the roles associated with the user for authorization decisions across all Web applications on the virtual server, without challenging the user to authenticate to each application individually.
 - When the user logs out of one Web application (explicitly, or because of session expiration), the user’s sessions in all Web applications become invalid. Thereafter, the user is required to log in to access a protected resource in any application.
- 1 In the Admin Console tree component, expand the Configurations node.
 - 2 Select the instance to configure:
 - To configure a particular instance, expand the instance’s config node. For example, the default instance, server, expand the server-config node.
 - To configure the default settings for all instances, expand the default-config node.

- 3 **Expand the HTTP Service node.**
- 4 **Expand the Virtual Servers node, and select the virtual server to be configured for single sign-on support.**
- 5 **Click Add Property.**
A blank property entry is added to the bottom of the list.
- 6 **Enter `sso-enable` in the Name field.**
- 7 **Enter `false` in the Value field to disable, enter `true` to enable SSO.**
SSO is enabled by default.
- 8 **Add or change any other single sign-on properties by clicking Add Property and configuring any applicable SSO properties.**
Valid SSO properties for virtual servers are discussed in the following table.

Property Name	Description	Values
<code>sso-max-inactive-seconds</code>	Number of seconds after which a user's single sign-on record becomes eligible for purging, if no client activity is received. Access to any of the applications on the virtual server keeps the single sign-on record active.	Default is 300 seconds (5 minutes). A higher value provides longer persistence for users, but consumes more memory on the server.
<code>sso-reap-interval-seconds</code>	Interval (in seconds) between purges of expired single sign-on records.	Default is 60.

- 9 **Click Save.**
- 10 **Restart the Application Server if Restart Required displays in the console.**

Admin Console Tasks for Connector Connection Pools

- [“About Connector Connection Pools” on page 188](#)
- [“About Security Maps” on page 188](#)

About Connector Connection Pools

A *connector module* (also called a resource adapter) enables J2EE applications to interact with enterprise information systems (EIS). A *connector resource* provides an application with a connection to an EIS. A *connector connection pool* is a group of reusable connections for a particular EIS.

Security maps enables the creation of a mapping between J2EE users and groups and EIS users and groups. Use the Admin Console to create, update, list, and delete security maps for connector connection pools.

Note – In this context, users are referred to as principals. The enterprise information system (EIS) is any system that holds the information. It can be a mainframe, a messaging system, a database system, or an application.

About Security Maps

Use security maps to map the caller identity of the application (principal or user group) to a suitable EIS principal in container-managed transaction-based scenarios. When an application principal initiates a request to an EIS, the application server first checks for an exact principal using the security map defined for the connector connection pool to determine the mapped back end EIS principal. If there is no exact match, then the application server uses the wild card character specification, if any, to determine the mapped back end EIS principal. Security maps are used when an application user needs to execute EIS operations that require to be executed as a specific identity in the EIS.

Use the following procedures in the Admin Console to manage security maps:

- [“To create a security map” on page 188](#)
- [“To edit a security map” on page 189](#)
- [“To delete a security map” on page 190](#)

▼ To create a security map

A security map for a connector connection pool maps application users and groups (principals) to EIS principals. Use a security map when an application user needs to execute EIS operations that require a specific identity in the EIS.

- 1 **Expand the Resources node**
- 2 **Expand the Connectors node.**
- 3 **Select the Connector Connection Pools node.**

- 4 Select a Connector Connection Pool by selecting its name from the list of current pools or create a new connector connection pool by selecting New from the list of current pools and following the instructions in [“To create a JDBC connection pool” on page 84.](#)
- 5 Select the Security Maps page.
- 6 Click New to create a new Security Map.
- 7 On the Create Security Map page, enter the following properties.
 - **Name** – Enter a name to be used to reference this particular security map.
 - **User Groups** – The caller identity of the application to be mapped to a suitable EIS principal. Enter a comma-separated list of application-specific user groups, or enter the wild card asterisk (*) to indicate all users or all user groups. Specify either the Principals or User Groups options, but not both.
 - **Principals** – The caller identity of the application to be mapped to a suitable EIS principal. Enter a comma-separated list of application-specific principals, or enter the wild card asterisk (*) to indicate all principals. Specify either the Principals or User Groups options, but not both.
- 8 In the Backend Principal section, enter the following properties.
 - **Username** – Enter the EIS user name. The enterprise information system (EIS) is any system that holds the information. It can be a mainframe, a messaging system, a database system, or an application.
 - **Password** – Enter the password for the EIS user.
- 9 Click OK to create the security map or Cancel to quit without saving.

More Information Equivalent asadmin command
`create-connector-security-map`

▼ To edit a security map

- 1 Expand the Resources node.
- 2 Expand the Connectors node.
- 3 Select the Connector Connection Pools node.
- 4 Select a Connector Connection Pool by selecting its name from the list of current pools.
- 5 Select the Security Maps page.

- 6 On the **Security Maps** page, select a security map from the list of current security maps.
- 7 On the **Edit Security Map** page, modify the following properties where needed.
 - **User Groups** – The caller identity of the application to be mapped to a suitable EIS principal. Enter a comma-separated list of application-specific user groups, or enter the wild card asterisk (*) to indicate all users or all user groups. Specify either the Principals or User Groups options, but not both.
 - **Principals** – The caller identity of the application to be mapped to a suitable EIS principal. Enter a comma-separated list of application-specific principals, or enter the wild card asterisk (*) to indicate all principals. Specify either the Principals or User Groups options, but not both.
- 8 In the **Backend Principal** section, enter the following properties.
 - **Username** – Enter the EIS user name. The enterprise information system (EIS) is any system that holds the information. It can be a mainframe, a messaging system, a database system, or an application.
 - **Password** – Enter the password for the EIS user.
- 9 Click **Save** to save the changes to the security map.

More Information Helpful asadmin commands

`list-connector-security-maps` and `update-connector-security-maps`

▼ To delete a security map

- 1 Expand the **Resources** node.
- 2 Expand the **Connectors** node.
- 3 Select the **Connector Connection Pools** node.
- 4 Select a **Connector Connection Pool** by selecting its name from the list of current pools.
- 5 Select the **Security Maps** page.
- 6 On the **Security Maps** page, click the checkbox to the left of the name of the security map to be deleted.
- 7 Click **Delete**.

More Information Equivalent `asadmin` command
`delete-connector-security-map`

Working with Certificates and SSL

- “About Certificate Files” on page 191
- “To change the location of certificate files” on page 191
- “Using Java Secure Socket Extension (JSSE) Tools” on page 192
- “Using Network Security Services (NSS) Tools” on page 196

About Certificate Files

Installation of the Application Server generates a digital certificate in NSS (Network Security Services) format suitable for internal testing. By default, the Application Server stores its certificate information in a certificate database in the `domain-dir/config` directory:

- **Keystore file**, `key3.db`, contains the Application Server’s certificate, including its private key. The keystore file is protected with a password. Change the password using the `asadmin change-master-password` command. For more information about `certutil`, read “Using the `certutil` Utility” on page 196.

Each keystore entry has a unique alias. After installation, the Application Server keystore has a single entry with alias `s1as`.

- **Truststore file**, `cert8.db`, contains the Application Server’s trusted certificates, including public keys for other entities. For a trusted certificate, the server has confirmed that the public key in the certificate belongs to the certificate’s owner. Trusted certificates generally include those of certification authorities (CAs).

In the Platform Edition, on the server side, the Application Server uses the JSSE format, which uses `keytool` to manage certificates and key stores. In the Enterprise Edition, on the server side, the Application Server uses NSS, which uses `certutil` to manage the NSS database which stores private keys and certificates. In both editions, the client side (appclient or stand-alone), uses the JSSE format.

By default, the Application Server is configured with a keystore and truststore that will work with the example applications and for development purposes. For production purposes, you may wish to change the certificate alias, add other certificates to the truststore, or change the name and/or location of the keystore and truststore files.

▼ To change the location of certificate files

The keystore and truststore files provided for development are stored in the `domain-dir/config` directory.

- 1 In the Admin Console tree, expand Configurations.

- 2 Expand the server-config (Admin Config) node.
- 3 Select the JVM Settings node.
- 4 Click the JVM Options tab.
- 5 On the JVM Options page, add or modify the following values in the Value field to reflect the new location of the certificate files:
`-Dcom.sun.appserv.nss.db=${com.sun.aas.instanceRoot}/NSS-database-directory`
where *NSS-database-directory* is the location of the NSS database.
- 6 Click Save.
- 7 Restart the Application Server if Restart Required displays in the console.

Using Java Secure Socket Extension (JSSE) Tools

Use `keytool` to set up and work with JSSE (Java Secure Socket Extension) digital certificates. In the Platform Edition, the Application Server uses the JSSE format on the server side to manage certificates and key stores. In both the Platform Edition and Enterprise Edition, the client side (appclient or stand-alone) uses the JSSE format.

The J2SE SDK ships with `keytool`, which enables the administrator to administer public/private key pairs and associated certificates. It also enables users to cache the public keys (in the form of certificates) of their communicating peers.

To run `keytool`, the shell environment must be configured so that the `J2SE/bin` directory is in the path, or the full path to the tool must be present on the command line. For more information on `keytool`, see the `keytool` documentation at <http://java.sun.com/j2se/1.5.0/docs/tooldocs/solaris/keytool.html>.

Using the keytool Utility

The following examples demonstrate usage related to certificate handling using JSSE tools:

- Create a self-signed certificate in a keystore of type JKS using an RSA key algorithm. RSA is public-key encryption technology developed by RSA Data Security, Inc. The acronym stands for Rivest, Shamir, and Adelman, the inventors of the technology.

```
keytool -genkey -noprompt -trustcacerts -keyalg RSA -alias ${cert.alias}  
-dname ${dn.name} -keypass ${key.pass} -keystore ${keystore.file}  
-storepass ${keystore.pass}
```

Another example of creating a certificate is shown in “To generate a certificate using the `keytool` utility” on page 193.

- Create a self-signed certificate in a keystore of type JKS using the default key algorithm.

```
keytool -genkey -noprompt -trustcacerts -alias ${cert.alias} -dnname
${dn.name} -keypass ${key.pass} -keystore ${keystore.file} -storepass
${keystore.pass}
```

An example of signing a certificate is shown in [“To sign a digital certificate using the keytool utility” on page 195](#)

- Display available certificates from a keystore of type JKS.

```
keytool -list -v -keystore ${keystore.file} -storepass ${keystore.pass}
```

- Display certificate information from a keystore of type JKS.

```
keytool -list -v -alias ${cert.alias} -keystore ${keystore.file}
-storepass ${keystore.pass}
```

- Import an RFC/text-formatted certificate into a JKS store. Certificates are often stored using the printable encoding format defined by the Internet RFC (Request for Comments) 1421 standard instead of their binary encoding. This certificate format, also known as *Base 64 encoding*, facilitates exporting certificates to other applications by email or through some other mechanism.

```
keytool -import -noprompt -trustcacerts -alias ${cert.alias} -file
${cert.file} -keystore ${keystore.file} -storepass ${keystore.pass}
```

- Export a certificate from a keystore of type JKS in PKCS7 format. The reply format defined by the Public Key Cryptography Standards #7, Cryptographic Message Syntax Standard, includes the supporting certificate chain in addition to the issued certificate.

```
keytool -export -noprompt -alias ${cert.alias} -file ${cert.file}
-keystore ${keystore.file} -storepass ${keystore.pass}
```

- Export a certificate from a keystore of type JKS in RFC/text format.

```
keytool -export -noprompt -rfc -alias ${cert.alias} -file
${cert.file} -keystore ${keystore.file} -storepass ${keystore.pass}
```

- Delete a certificate from a keystore of type JKS.

```
keytool -delete -noprompt -alias ${cert.alias} -keystore ${keystore.file}
-storepass ${keystore.pass}
```

Another example of deleting a certificate from a keystore is shown in [“Deleting a Certificate Using the keytool Utility” on page 195](#)

▼ To generate a certificate using the keytool utility

Use `keytool` to generate, import, and export certificates. By default, `keytool` creates a keystore file in the directory where it is run.

1 Change to the directory where the certificate is to be run.

Always generate the certificate in the directory containing the keystore and truststore files, by default *domain-dir/config*. For information on changing the location of these files, see [“To change the location of certificate files” on page 191](#).

2 Enter the following keytool command to generate the certificate in the keystore file, keystore.jks:

```
keytool -genkey -alias keyAlias-keyalg RSA
-keypass changeit
-storepass changeit
-keystore keystore.jks
```

Use any unique name as your *keyAlias*. If you have changed the keystore or private key password from their default, then substitute the new password for *changeit* in the above command.

A prompt appears that asks for your name, organization, and other information that keytool uses to generate the certificate.

3 Enter the following keytool command to export the generated certificate to the file server.cer (or client.cer if you prefer):

```
keytool -export -alias keyAlias-storepass changeit
-file server.cer
-keystore keystore.jks
```

4 If a certificate signed by a certificate authority is required, see [“To sign a digital certificate using the keytool utility” on page 195](#).**5 To create the truststore file cacerts.jks and add the certificate to the truststore, enter the following keytool command:**

```
keytool -import -v -trustcacerts
-alias keyAlias
-file server.cer
-keystore cacerts.jks
-keypass changeit
```

If you have changed the keystore or private key password from their default, then substitute the new password for *changeit* in the above command.

The tool displays information about the certificate and prompts whether you want to trust the certificate.

6 Type yes, then press Enter.

Then keytool displays something like this:

```
Certificate was added to keystore
[Saving cacerts.jks]
```

7 Restart the Application Server.

▼ To sign a digital certificate using the `keytool` utility

After creating a digital certificate, the owner must sign it to prevent forgery. E-commerce sites, or those for which authentication of identity is important can purchase a certificate from a well-known Certificate Authority (CA). If authentication is not a concern, for example if private secure communications is all that is required, save the time and expense involved in obtaining a CA certificate and use a self-signed certificate.

1 Follow the instructions on the CA's Web site for generating certificate key pairs.

2 Download the generated certificate key pair.

Save the certificate in the directory containing the keystore and truststore files, by default `domain-dir/config` directory. See “[To change the location of certificate files](#)” on page 191.

3 In your shell, change to the directory containing the certificate.

4 Use `keytool` to import the certificate into the local keystore and, if necessary, the local truststore.

```
keytool -import -v -trustcacerts
-alias keyAlias
-file server.cer
-keystore cacerts.jks
-keypass changeit
-storepass changeit
```

If the keystore or private key password is not the default password, then substitute the new password for `changeit` in the above command.

5 Restart the Application Server.

Deleting a Certificate Using the `keytool` Utility

To delete an existing certificate, use the `keytool -delete` command, for example:

```
keytool -delete
-alias keyAlias
-keystore keystore-name
-storepass password
```

For a complete list of possible options for the `-delete` command, refer to the `keytool` documentation at <http://java.sun.com/j2se/1.5.0/docs/tooldocs/solaris/keytool.html>.

Using Network Security Services (NSS) Tools

In the Enterprise Edition, use Network Security Services (NSS) digital certificates on the server-side to manage the database that stores private keys and certificates. For the client side (appliant or stand-alone), use the JSSE format as discussed in [“Using Java Secure Socket Extension \(JSSE\) Tools” on page 192](#).

The tools for managing security with Network Security Services (NSS) include the following:

- `certutil`, a command-line utility for managing certificates and key databases. Some examples using the `certutil` utility are shown in [“Using the `certutil` Utility” on page 196](#).
- `pk12util`, a command-line utility used to import and export keys and certificates between the certificate/key databases and files in PKCS12 format. Some examples using the `pk12util` utility are shown in [“Importing and Exporting Certificates Using the `pk12util` Utility” on page 198](#).
- `modutil`, a command-line utility for managing PKCS #11 module information within `secmod.db` files or within hardware tokens. Some examples using the `modutil` utility are shown in [“Adding and Deleting PKCS11 Modules using `modutil`” on page 199](#).

The tools are located in the `install-dir/lib/` directory. The following environment variables are used to point to the location of the NSS security tools:

- `LD_LIBRARY_PATH = ${install-dir}/lib`
- `${os.nss.path}`

In the examples, the certificate common name (CN) is the name of the client or server. The CN is also used during SSL handshake for comparing the certificate name and the host name from which it originates. If the certificate name and the host name do not match, warnings or exceptions are generated during SSL handshake. In some examples, the certificate common name `CN=localhost` is used for convenience so that all users can use that certificate instead of creating a new one with their real host name.

The examples in the following sections demonstrate usage related to certificate handling using NSS tools:

- [“Using the `certutil` Utility” on page 196](#)
- [“Importing and Exporting Certificates Using the `pk12util` Utility” on page 198](#)
- [“Adding and Deleting PKCS11 Modules using `modutil`” on page 199](#)

Using the `certutil` Utility

The certificate database tool, `certutil`, is an NSS command-line utility that can create and modify the Netscape Communicator `cert8.db` and `key3.db` database files. It can also list, generate, modify, or delete certificates within the `cert8.db` file and create or change the password, generate new public and private key pairs, display the contents of the key database, or delete key pairs within the `key3.db` file.

The key and certificate management process generally begins with creating keys in the key database, then generating and managing certificates in the certificate database. The following document discusses certificate and key database management with NSS, including the syntax for the certutil utility: <http://www.mozilla.org/projects/security/pki/nss/tools/certutil.html>.

Each of the items in the list below gives an example using NSS and JSSE security tools to create and/or manage certificates.

- Generate a self-signed server and client certificate. In this example, the CN must be of the form `hostname.domain.[com|org|net|...]`.

In this example, `domain-dir/config`. The `serverseed.txt` and `clientseed.txt` files can contain any random text. This random text will be used for generating the key pair.

```
certutil -S -n $SERVER_CERT_NAME -x -t "u,u,u"
-s "CN=$HOSTNAME.$HOSTDOMAIN, OU=Java Software, O=Sun Microsystems Inc.,
  L=Santa Clara, ST=CA, C=US"
-m 25001 -o $CERT_DB_DIR/Server.crt
-d $CERT_DB_DIR -f passfile <$CERT_UTIL_DIR/serverseed.txt
```

Generate the client certificate. This certificate is also a self-signed certificate.

```
certutil -S -n $CLIENT_CERT_NAME -x -t "u,u,u"
-s "CN=MyClient, OU=Java Software, O=Sun Microsystems Inc.,
  L=Santa Clara, ST=CA, C=US"
-m 25002 -o $CERT_DB_DIR/Client.crt
-d $CERT_DB_DIR -f passfile <$CERT_UTIL_DIR/clientseed.txt
```

- Verify the certificates generated in the previous bullet.

```
certutil -V -u V -n $SERVER_CERT_NAME -d $CERT_DB_DIR
certutil -V -u C -n $CLIENT_CERT_NAME -d $CERT_DB_DIR
```

- Display available certificates.

```
certutil -L -d $CERT_DB_DIR
```

- Import an RFC text-formatted certificate into an NSS certificate database.

```
<appserver_install>/lib/certutil -A -n <Intermediate_CA_cert_name_to_display_at_DAS_certDB> -i <Intermediate_CA_c
```

- Export a certificate from an NSS certificate database in RFC format.

```
certutil -L -a -n ${cert.nickname} -f ${pass.file}
-d ${admin.domain.dir}/${admin.domain}/config > cert.rfc
```

- Delete a certificate from an NSS certificate database.

```
certutil -D -n ${cert.nickname} -f ${pass.file}
-d ${admin.domain.dir}/${admin.domain}/config
```

- Move a certificate from an NSS database to JKS format

```
certutil -L -a -n ${cert.nickname}
-d ${admin.domain.dir}/${admin.domain}/config > cert.rfc
```

```
keytool -import -noprompt -trustcacerts -keystore ${keystore.file}
-storepass ${keystore.pass} -alias ${cert.alias} -file cert.rfc
```

Importing and Exporting Certificates Using the pk12util Utility

The command-line utility used to import and export keys and certificates between the certificate/key databases and files in PKCS12 format is `pk12util`. PKCS12 is Public-Key Cryptography Standards (PKCS) #12, Personal Information Exchange Syntax Standard. More description of the `pk12util` utility can be read at <http://www.mozilla.org/projects/security/pki/nss/tools/pk12util.html>.

- Import a PKCS12-formatted certificate into an NSS certificate database.

```
pk12util -i ${cert.pkcs12.file} -k ${certdb.pass.file}
-w ${cert.pass.file} -d ${admin.domain.dir}/${admin.domain}/config
```

- Import a PKCS12-formatted certificate into an NSS certificate database token module.

```
pk12util -i ${cert.pkcs12.file} -h ${token.name} -k ${certdb.pass.file}
-w ${cert.pass.file} -d ${admin.domain.dir}/${admin.domain}/config
```

- Export a certificate from an NSS certificate database in PKCS12 format.

```
pk12util -o -n ${cert.nickname} -k ${pass.file} -w${cert.pass.file}
-d ${admin.domain.dir}/${admin.domain}/config
```

- Export a certificate from an NSS certificate database token module in PKCS12 format (useful for hardware accelerator configuration).

```
pk12util -o -n ${cert.nickname} -h ${token.name} -k ${pass.file}
-w ${cert.pass.file} -d ${admin.domain.dir}/${admin.domain}/config
```

- Convert a PKCS12 certificate into JKS format (requires a Java source):

```
<target name="convert-pkcs12-to-jks" depends="init-common">
  <delete file="${jks.file}" failonerror="false"/>
  <java classname="com.sun.enterprise.security.KeyTool">
    <arg line="-pkcs12"/>
    <arg line="-pkcsFile ${pkcs12.file}"/>
    <arg line="-pkcsKeyStorePass ${pkcs12.pass}"/>
    <arg line="-pkcsKeyPass ${pkcs12.pass}"/>
    <arg line="-jksFile ${jks.file}"/>
    <arg line="-jksKeyStorePass ${jks.pass}"/>
    <classpath>
      <pathelement path="${slas.classpath}"/>
      <pathelement path="${env.JAVA_HOME}/jre/lib/jsse.jar"/>
    </classpath>
  </java>
</target>
```

Adding and Deleting PKCS11 Modules using `modutil`

The *Security Module Database Tool*, `modutil`, is a command-line utility for managing PKCS #11 (Cryptographic Token Interface Standard) module information within `secmod.db` files or within hardware tokens. You can use the tool to add and delete PKCS #11 modules, change passwords, set defaults, list module contents, enable or disable slots, enable or disable FIPS-140-1 compliance, and assign default providers for cryptographic operations. This tool can also create `key3.db`, `cert7.db`, and `secmod.db` security database files. For more information on this tool, see <http://www.mozilla.org/projects/security/pki/nss/tools/modutil.html>.

- Add a new PKCS11 module or token.

```
modutil -add ${token.module.name} -nocertdb -force -mechanisms RSA:DSA:RC4:DES  
-libfile ${SCA.lib.path} -dbdir ${admin.domain.dir}/${admin.domain}/config
```

- Delete a PKCS11 module from an NSS store.

```
modutil -delete ${token.module.name} -nocertdb -force -mechanisms RSA:DSA:RC4:DES  
-libfile ${SCA.lib.path} -dbdir ${admin.domain.dir}/${admin.domain}/config
```

- List available token modules in an NSS store.

```
modutil -list -dbdir ${admin.domain.dir}/${admin.domain}/config
```

Further Information

- The Java 2 Standard Edition discussion of security can be viewed from <http://java.sun.com/j2se/1.5.0/docs/guide/security/index.html>.
- The *J2EE 1.4 Tutorial* chapter titled *Security* can be viewed from <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.
- The *Administration Guide* chapter titled **Chapter 10, “Configuring Message Security”**
- The *Developer’s Guide* chapter titled “Securing Applications.”

Configuring Message Security

This chapter describes the configuration of message layer security for web services in the Application Server. This chapter contains the following topics:

- [“About Application Server Security” on page 145](#)
- [“Admin Console Tasks for Security” on page 162](#)

Some of the material in this chapter assumes a basic understanding of security and web services concepts. To learn more about these concepts, explore the resources listed in [“Further Information” on page 199](#) before beginning this chapter.

About Message Security

- [“Overview of Security” on page 145](#)
- [“Understanding Message Security in the Application Server” on page 202](#)
- [“Securing a Web Service” on page 206](#)
- [“Securing the Sample Application” on page 207](#)
- [“Configuring the Application Server for Message Security” on page 207](#)

Overview of Message Security

In *message security*, security information is inserted into messages so that it travels through the networking layers and arrives with the message at the message destination(s). Message security differs from transport layer security (which is discussed in the *Security* chapter of the *J2EE 1.4 Tutorial*) in that message security can be used to decouple message protection from message transport so that messages remain protected after transmission.

Web Services Security: SOAP Message Security (WS-Security) is an international standard for interoperable Web Services Security that was developed in OASIS by a collaboration of all the major providers of web services technology (including Sun Microsystems). WS-Security is a message security mechanism that uses XML Encryption and XML Digital Signature to secure

web services messages sent over SOAP. The WS-Security specification defines the use of various security tokens including X.509 certificates, SAML assertions, and username/password tokens to authenticate and encrypt SOAP web services messages.

The WS-Security specification can be viewed at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf>.

Understanding Message Security in the Application Server

The Application Server offers integrated support for the WS-Security standard in its web services client and server-side containers. This functionality is integrated such that web services security is enforced by the containers of the Application Server on behalf of applications, and such that it can be applied to protect any web service application without requiring changes to the implementation of the application. The Application Server achieves this effect by providing facilities to bind SOAP layer message security providers and message protection policies to containers and to applications deployed in containers.

Assigning Message Security Responsibilities

In the Application Server, the “[System Administrator](#)” on page 150 and “[Application Deployer](#)” on page 150 roles are expected to take primary responsibility for configuring message security. In some situations, the “[Application Developer](#)” on page 150 may also contribute, although in the typical case either of the other roles may secure an existing application without changing its implementation without involving the developer. The responsibilities of the various roles are defined in the following sections:

- “[System Administrator](#)” on page 150
- “[Application Deployer](#)” on page 150
- “[Application Developer](#)” on page 150

System Administrator

The system administrator is responsible for:

- Configuring message security providers on the Application Server.
- Managing user databases.
- Managing keystore and truststore files.
- Configuring a Java Cryptography Extension (JCE) provider if using encryption and running a version of the Java SDK prior to version 1.5.0.
- Installing the samples server. This is only done if the `xms` sample application will be used to demonstrate the use of message layer web services security.

A system administrator uses the Admin Console to manage server security settings and uses a command line tool to manage certificate databases. In Platform Edition, certificates and private keys are stored in key stores and are managed with `keytool`. Standard Edition and Enterprise Edition store certificates and private keys in an NSS database, where they are managed using `certutil`. This document is intended primarily for system administrators. For an overview of message security tasks, see [“Configuring the Application Server for Message Security” on page 207](#).

Application Deployer

The application deployer is responsible for:

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

These security tasks are discussed in the *Securing Applications* chapter of the *Developers' Guide*. For a link to this chapter, see [“Further Information” on page 199](#).

Application Developer

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

The application developer or assembler is responsible for the following:

- Determining if an application-specific message protection policy is required by the application. If so, ensuring that the required policy is specified at application assembly which may be accomplished by communicating with the Application Deployer.

About Security Tokens and Security Mechanisms

The WS-Security specification provides an extensible mechanism for using security tokens to authenticate and encrypt SOAP web services messages. The SOAP layer message security providers installed with the Application Server may be used to employ username/password and X.509 certificate security tokens to authenticate and encrypt SOAP web services messages. Additional providers that employ other security tokens including SAML assertions will be installed with subsequent releases of the Application Server.

About Username Tokens

The Application Server uses *Username tokens* in SOAP messages to establish the authentication identity of the message *sender*. The recipient of a message containing a Username token (within embedded password) validates that the message sender is authorized to act as the user (identified in the token) by confirming that the sender knows the secret (the password) of the user.

When using a Username token, a valid user database must be configured on the Application Server. For more information on this topic, read [“To edit a realm” on page 166](#).

About Digital Signatures

The Application Server uses XML Digital signatures to bind an authentication identity to message *content*. Clients use digital signatures to establish their caller identity, analogous to the way basic authentication or SSL client certificate authentication have been used to do the same thing when transport layer security is being used. Digital signatures are verified by the message receiver to authenticate the source of the message content (which may be different from the sender of the message.)

When using digital signatures, valid keystore and truststore files must be configured on the Application Server. For more information on this topic, read [“About Certificate Files” on page 191](#).

About Encryption

The purpose of encryption is to modify the data such that it can only be understood by its intended audience. This is accomplished by substituting an encrypted element for the original content. When predicated on public key cryptography, encryption can be used to establish the identity of the parties that can read a message.

When using Encryption, you must have an installed JCE provider that supports encryption. For more information on this topic, read [“To configure a JCE Provider” on page 209](#).

About Message Protection Policies

Message protection policies are defined for request message processing and response message processing and are expressed in terms of requirements for source and/or recipient authentication. A source authentication policy represents a requirement that the identity of the entity that sent a message or that defined the content of a message be established in the message such that it can be authenticated by the message receiver. A recipient authentication policy represents a requirement that the message be sent such that the identity of the entities that can receive the message can be established by the message sender. The providers apply specific message security mechanisms to cause the message protection policies to be realized in the context of SOAP web services messages.

Request and response message protection policies are defined when a provider is configured into a container. Application-specific message protection policies (at the granularity of the web service port or operation) may also be configured within the Sun-specific deployment descriptors of the application or application client. In any case, where message protection policies are defined, the request and response message protection policies of the client must match (be equivalent to) the request and response message protection policies of the server. For more information on defining application-specific message protection policies, refer to the *Securing Applications* chapter of the *Developers' Guide*.

Glossary of Message Security Terminology

The terminology used in this document is described below. The concepts are also discussed in [“Configuring the Application Server for Message Security” on page 207](#).

- Authentication Layer

The *authentication layer* is the message layer on which authentication processing must be performed. The Application Server enforces web services message security at the SOAP layer.

- Authentication Provider

In this release of the Application Server, the Application Server invokes *authentication providers* to process SOAP message layer security.

- A *client-side provider* establishes (by signature or username/password) the source identity of request messages and/or protects (by encryption) request messages such that they can only be viewed by their intended recipients. A client-side provider also establishes its container as an authorized recipient of a received response (by successfully decrypting it) and validates passwords or signatures in the response to authenticate the source identity associated with the response. Client-side providers configured in the Application Server can be used to protect the request messages sent and the response messages received by server-side components (servlets and EJB components) acting as clients of other services.

- A *server-side provider* establishes its container as an authorized recipient of a received request (by successfully decrypting it) and validates passwords or signatures in the request to authenticate the source identity associated with the request. A server-side provider also establishes (by signature or username/password) the source identity of response messages and/or protects (by encryption) response messages such that they can only be viewed by their intended recipients. *Server-side providers* are only invoked by server-side containers.

- Default Server Provider

The *default server provider* is used to identify the server provider to be invoked for any application for which a specific server provider has not been bound. The *default server provider* is sometimes referred to as the *default provider*.

- Default Client Provider

The *default client provider* is used to identify the client provider to be invoked for any application for which a specific client provider has not been bound.

- Request Policy

The *request policy* defines the authentication policy requirements associated with request processing performed by the authentication provider. Policies are expressed in message sender order such that a requirement that encryption occur after content would mean that the message receiver would expect to decrypt the message before validating the signature.

- Response Policy

The *response policy* defines the authentication policy requirements associated with response processing performed by the authentication provider. Policies are expressed in message sender order such that a requirement that encryption occur after content would mean that the message receiver would expect to decrypt the message before validating the signature.

Securing a Web Service

Web services deployed on the Application Server are secured by binding SOAP layer message security providers and message protection policies to the containers in which the applications are deployed or to web service endpoints served by the applications. SOAP layer message security functionality is configured in the client-side containers of the Application Server by binding SOAP layer message security providers and message protection policies to the client containers or to the portable service references declared by client applications.

When the Application Server is installed, SOAP layer message security providers are configured in the client and server-side containers of the Application Server, where they are available for binding for use by the containers, or by individual applications or clients deployed in the containers. During installation, the providers are configured with a simple message protection policy that, if bound to a container, or to an application or client in a container, would cause the source of the content in all request and response messages to be authenticated by XML digital signature.

The administrative interfaces of the Application Server can be employed to bind the existing providers for use by the server-side containers of the Application Server, to modify the message protection policies enforced by the providers, or to create new provider configurations with alternative message protection policies. These operations are defined in [“Admin Console Tasks for Security” on page 162](#). Analogous administrative operations can be performed on the SOAP message layer security configuration of the application client container as defined in [“To enable message security for application clients” on page 219](#).

By default, message layer security is disabled on the Application Server. To configure message layer security for the Application Server follow the steps outlined in [“Configuring the Application Server for Message Security” on page 207](#). If you want to cause web services security to be used to protect all web services applications deployed on the Application Server, follow the steps in [“To enable providers for message security” on page 211](#).

Once you have completed the above steps (which may include restarting the Application Server), web services security will be applied to all web services applications deployed on the Application Server.

Configuring Application-Specific Web Services Security

Application-specific web services security functionality is configured (at application assembly) by defining message-security-binding elements in the Sun-specific deployment descriptors of the application. These message-security-binding elements are used to associate a specific provider or message protection policy with a web services endpoint or service reference, and may be qualified so that they apply to a specific port or method of the corresponding endpoint or referenced service.

For more information on defining application specific message protection policies, refer to the *Securing Applications* chapter of the *Developers' Guide*. There is a link to this chapter in [“Further Information” on page 199](#).

Securing the Sample Application

The Application Server ships with a sample application named `xms`. The `xms` application features a simple web service that is implemented by both a J2EE EJB endpoint and a Java Servlet endpoint. Both endpoints share the same service endpoint interface. The service endpoint interface defines a single operation, `sayHello`, which takes a string argument, and returns a `String` composed by pre-pending `Hello` to the invocation argument.

The `xms` sample application is provided to demonstrate the use of the Application Server's WS-Security functionality to secure an existing web services application. The instructions which accompany the sample describe how to enable the WS-Security functionality of the Application Server such that it is used to secure the `xms` application. The sample also demonstrates the binding of WS-Security functionality directly to the application (as described in [“Configuring Application-Specific Web Services Security” on page 207](#) application.

The `xms` sample application is installed in the directory:
`install-dir/samples/webservices/security/ejb/apps/xms/`.

For information on compiling, packaging, and running the `xms` sample application, refer to the *Securing Applications* chapter of the *Developers' Guide*.

Configuring the Application Server for Message Security

- [“Actions of Request and Response Policy Configurations” on page 208](#)
- [“To configure other security facilities” on page 208](#)
- [“To configure a JCE Provider” on page 209](#)

Actions of Request and Response Policy Configurations

The following table shows message protection policy configurations and the resulting message security operations performed by the WS-Security SOAP message security providers for that configuration.

TABLE 10-1 Message protection policy to WS-Security SOAP message security operation mapping

Message Protection Policy	Resulting WS-Security SOAP message protection operations
auth-source="sender"	The message contains a <code>wsse:Security</code> header that contains a <code>wsse:UsernameToken</code> (with password).
auth-source="content"	The content of the SOAP message Body is signed. The message contains a <code>wsse:Security</code> header that contains the message Body signature represented as a <code>ds:Signature</code> .
auth-source="sender" auth-recipient="before-content" OR auth-recipient="after-content"	The content of the SOAP message Body is encrypted and replaced with the resulting <code>xenc:EncryptedData</code> . The message contains a <code>wsse:Security</code> header that contains a <code>wsse:UsernameToken</code> (with password) and an <code>xenc:EncryptedKey</code> . The <code>xenc:EncryptedKey</code> contains the key used to encrypt the SOAP message body. The key is encrypted in the public key of the recipient.
auth-source="content" auth-recipient="before-content"	The content of the SOAP message Body is encrypted and replaced with the resulting <code>xenc:EncryptedData</code> . The <code>xenc:EncryptedData</code> is signed. The message contains a <code>wsse:Security</code> header that contains an <code>xenc:EncryptedKey</code> and a <code>ds:Signature</code> . The <code>xenc:EncryptedKey</code> contains the key used to encrypt the SOAP message body. The key is encrypted in the public key of the recipient.
auth-source="content" auth-recipient="after-content"	The content of the SOAP message Body is signed, then encrypted, and then replaced with the resulting <code>xenc:EncryptedData</code> . The message contains a <code>wsse:Security</code> header that contains an <code>xenc:EncryptedKey</code> and a <code>ds:Signature</code> . The <code>xenc:EncryptedKey</code> contains the key used to encrypt the SOAP message body. The key is encrypted in the public key of the recipient.
auth-recipient="before-content" OR auth-recipient="after-content"	The content of the SOAP message Body is encrypted and replaced with the resulting <code>xenc:EncryptedData</code> . The message contains a <code>wsse:Security</code> header that contains an <code>xenc:EncryptedKey</code> . The <code>xenc:EncryptedKey</code> contains the key used to encrypt the SOAP message body. The key is encrypted in the public key of the recipient.
No policy specified.	No security operations are performed by the modules.

▼ To configure other security facilities

The Application Server implements message security using message security providers integrated in its SOAP processing layer. The message security providers depend on other security facilities of Application Server.

- 1 **If using a version of the Java SDK prior to version 1.5.0, and using encryption technology, configure a JCE provider.**

Configuring a JCE provider is discussed in “[To configure a JCE Provider](#)” on page 209.

- 2 **If using a username token, configure a user database, if necessary. When using a username/password token, an appropriate realm must be configured and an appropriate user database must be configured for the realm.**

Configuring a user database is discussed in “[To edit a realm](#)” on page 166.

- 3 **Manage certificates and private keys, if necessary.**

Managing certificates and private keys is discussed in “[About Certificate Files](#)” on page 191.

Next Steps Once the facilities of the Application Server are configured for use by message security providers, then the providers installed with the Application Server may be enabled as described in “[To enable providers for message security](#)” on page 211.

▼ To configure a JCE Provider

The Java Cryptography Extension (JCE) provider included with J2SE 1.4.x does not support RSA encryption. Because the XML Encryption defined by WS-Security is typically based on RSA encryption, in order to use WS-Security to encrypt SOAP messages you must download and install a JCE provider that supports RSA encryption.

Note – RSA is public-key encryption technology developed by RSA Data Security, Inc. The acronym stands for Rivest, Shamir, and Adelman, the inventors of the technology.

If you are running the Application Server on version 1.5 of the Java SDK, the JCE provider is already configured properly. If you are running the Application Server on version 1.4.x of the Java SDK, you can add a JCE provider statically as part of your JDK environment, as follows.

- 1 **Download and install a JCE provider JAR (Java ARchive) file.**

The following URL provides a list of JCE providers that support RSA encryption:
http://java.sun.com/products/jce/jce14_providers.html.

- 2 **Copy the JCE provider JAR file to `java-home/jre/lib/ext/`.**

- 3 **Stop the Application Server.**

If the Application Server is not stopped and then restarted later in this process, the JCE provider will not be recognized by the Application Server.

4 Edit the `java-home/jre/lib/security/java.security` properties file in any text editor. Add the JCE provider you've just downloaded to this file.

The `java.security` file contains detailed instructions for adding this provider. Basically, you need to add a line of the following format in a location with similar properties:

```
security.provider.n=provider-class-name
```

In this example, *n* is the order of preference to be used by the Application Server when evaluating security providers. Set *n* to 2 for the JCE provider you've just added.

For example, if you've downloaded The Legion of the Bouncy Castle JCE provider, you would add this line.

```
security.provider.2=org.bouncycastle.jce.provider.  
    BouncyCastleProvider
```

Make sure that the Sun security provider remains at the highest preference, with a value of 1.

```
security.provider.1=sun.security.provider.Sun
```

Adjust the levels of the other security providers downward so that there is only one security provider at each level.

The following is an example of a `java.security` file that provides the necessary JCE provider and keeps the existing providers in the correct locations.

```
security.provider.1=sun.security.provider.Sun  
security.provider.2=org.bouncycastle.jce.provider.  
    BouncyCastleProvider  
security.provider.3=com.sun.net.ssl.internal.ssl.Provider  
security.provider.4=com.sun.rsa.jca.Provider  
security.provider.5=com.sun.crypto.provider.SunJCE  
security.provider.6=sun.security.jgss.SunProvider
```

5 Save and close the file.

6 Restart the Application Server.

Admin Console Tasks for Message Security

Most of the steps for setting up the Application Server for using message security can be accomplished using the Admin Console, the `asadmin` command-line tool, or by manually editing system files. In general, editing system files is discouraged due to the possibility of making unintended changes that prevent the Application Server from running properly, therefore, where possible, steps for configuring the Application Server using the Admin Console are shown first, with the `asadmin` tool command shown after. Steps for manually editing system files are shown only when there is no Admin Console or `asadmin` equivalent.

Support for message layer security is integrated into the Application Server and its client containers in the form of (pluggable) authentication modules. By default, message layer security is disabled on the Application Server. The following sections provide the details for enabling, creating, editing, and deleting message security configurations and providers.

- [“To enable providers for message security” on page 211](#)
- [“To configure a message security provider” on page 213](#)
- [“Creating a Message Security Provider” on page 215](#)
- [“To delete a message security configuration” on page 217](#)
- [“To delete a message security provider” on page 218](#)
- [“To enable message security for application clients” on page 219](#)

In most cases, it will be necessary to restart the Application Server after performing the administrative operations listed above. This is especially the case if you want the effects of the administrative change to be applied to applications that were already deployed on the Application Server at the time the operation was performed.

▼ To enable providers for message security

To enable message security for web services endpoints deployed in the Application Server, you must specify a provider to be used by default on the server side. If you enable a default provider for message security, you also need to enable providers to be used by clients of the web services deployed in the Application Server. Information for enabling the providers used by clients is discussed in [“To enable message security for application clients” on page 219](#).

To enable message security for web service invocations originating from deployed endpoints, you must specify a default client provider. If you enabled a default client provider for the Application Server, you must ensure that any services invoked from endpoints deployed in the Application Server are compatibly configured for message layer security.

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Expand the Security node.
- 4 Expand the Message Security node.

- 5 Select the SOAP node.
- 6 Select the Message Security tab.
- 7 On the Edit Message Security Configuration page, specify a provider to be used on the server side and a provider to be used on the client side for all applications for which a specific provider has not been bound.

This is accomplished by modifying the following optional properties:

- **Default Provider** – The identity of the server provider to be invoked for any application for which a specific server provider has not been bound.

By default, no provider configuration is selected for the Application Server. To identify a server-side provider, select `ServerProvider`. Selecting the null option means that no message security provider will be invoked (by default) on the server side.

You would generally select `ServerProvider` for this field.

- **Default Client Provider** – The identity of the client provider to be invoked for any application for which a specific client provider has not been bound.

By default, no provider configuration is selected for the Application Server. To identify a client-side provider, select `ClientProvider`. Selecting the null option means that no message security provider will be invoked (by default) on the client side.

You would generally select null for this field. You would select `ClientProvider` if you wanted to enable a default provider and message protection policy to apply to the web services invocations originating from web services endpoints deployed on the Application Server.

- 8 Click Save.
- 9 If you enabled a client or server provider and you want to modify the message protection policies of the enabled providers, refer to [“To configure a message security provider” on page 213](#) for information on modifying the message security providers enabled in this step.

More Information Equivalent `asadmin` commands

- To specify the default server provider:

```
asadmin set --user admin-user --port admin-port
server-config.security-service.message-security-config.SOAP.
default_provider=ServerProvider
```

- To specify the default client provider:

```
asadmin set --user admin-user --port admin-port
server-config.security-service.message-security-config.SOAP.
default_client_provider=ClientProvider
```

▼ To configure a message security provider

Typically, a provider would be reconfigured to modify its message protection policies, although the provider type, implementation class, and provider-specific configuration properties may also be modified.

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Expand the Security node.
- 4 Expand the Message Security node.
- 5 Select the SOAP node.
- 6 Select the Providers tab.
- 7 Select the message security provider to edit.

`ClientProvider` and `ServerProvider` ship with the Application Server.
- 8 In the Provider Config section of the Edit Provider Config page, the following properties are available for modification:
 - **Provider Type** – Select `client`, `server`, or `client-server` to establish whether the provider is to be used as a client authentication provider, a server authentication provider, or both (a client-server provider).
 - **Class Name** - Enter the Java implementation class of the provider. Client authentication providers must implement the `com.sun.enterprise.security.jauth.ClientAuthModule` interface. Server-side providers must implement the `com.sun.enterprise.security.jauth.ServerAuthModule` interface. A provider may implement both interfaces, but it must implement the interface corresponding to its provider type.
- 9 In the Request Policy section of the Create a Provider Configuration page, enter the following optional values, if needed.

These properties are optional, but if not specified, no authentication is applied to request messages.

The *request policy* defines the authentication policy requirements associated with request processing performed by the authentication provider. Policies are expressed in message sender order such that a requirement that encryption occur after content would mean that the message receiver would expect to decrypt the message before validating the signature.

- **Authentication Source**– Select sender, content, or null (the blank option) to define a requirement for message-layer sender authentication (for example, username password), content authentication (for example, digital signature), or no authentication be applied to request messages. When null is specified, source authentication of the request is not required.
- **Authentication Recipient**– Select beforeContent or afterContent to define a requirement for message-layer authentication of the receiver of the request message to its sender (by XML encryption). When the value is not specified it defaults to afterContent.

For a description of the actions performed by the SOAP message security providers as a result of the following message protection policies see [“Actions of Request and Response Policy Configurations” on page 208](#).

10 In the Response Policy section of the Create a Provider Configuration page, enter the following optional properties, if needed.

These properties are optional, but if not specified, no authentication is applied to response messages.

The *response policy* defines the authentication policy requirements associated with response processing performed by the authentication provider. Policies are expressed in message sender order such that a requirement that encryption occur after content would mean that the message receiver would expect to decrypt the message before validating the signature.

- **Authentication Source** – Select sender, content, or null (the blank option) to define a requirement for message-layer sender authentication (for example, username password) or content authentication (for example, digital signature) to be applied to response messages. When null is specified, source authentication of the response is not required.
- **Authentication Recipient** – Select beforeContent or afterContent to define a requirement for message-layer authentication of the receiver of the response message to its sender (by XML encryption). When the value is not specified it defaults to afterContent.

For a description of the actions performed by the SOAP message security providers as a result of the following message protection policies see [“Actions of Request and Response Policy Configurations” on page 208](#).

11 Add additional properties by clicking the Add Property button.

The provider that is shipped with the Application Server supports the property listed below. If other providers are used, refer to their documentation for more information on properties and valid values.

- `server.config` – The directory and file name of an XML file that contains the server configuration information. For example, `domain-dir/config/wss-server-config.xml`.

12 Click Save.

More Information Equivalent asadmin commands

To set the response policy, replace the word `request` in the following commands with `response`.

- Add a request policy to the client and set the authentication source:

```
asadmin set --user admin-user --port admin-port
server-config.security-service.message-security-config.SOAP.
provider-config.ClientProvider.request-policy.auth_source=
sender | content
```

- Add a request policy to the server and set the authentication source:

```
asadmin set --user admin-user --port admin-port
server-config.security-service.message-security-config.SOAP.
provider-config.ServerProvider.request-policy.auth_source=
sender | content
```

- Add a request policy to the client and set the authentication recipient:

```
asadmin set --user admin-user --port admin-port
server-config.security-service.message-security-config.SOAP.
provider-config.ClientProvider.request-policy.auth_recipient=
before-content | after-content
```

- Add a request policy to the server and set the authentication recipient:

```
asadmin set --user admin-user --port admin-port
server-config.security-service.message-security-config.SOAP.
provider-config.ServerProvider.request-policy.auth_recipient=
before-content | after-content
```

▼ Creating a Message Security Provider

To configure an existing provider, follow the steps in [“To configure a message security provider” on page 213](#).

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, the default instance, `server`, select the `server-config` node.

- **To configure the default settings for all instances, select the `default-config` node.**
- 3 **Expand the Security node.**
 - 4 **Expand the Message Security node.**
 - 5 **Select the SOAP node.**
 - 6 **Select the Providers tab.**
 - 7 **On the Provider Configuration page, click New.**
 - 8 **In the Provider Config section of the Create a Provider Configuration page, enter the following:**
 - **Default Provider** – Check the box beside this field to make the new message security provider the provider to be invoked for any application for which a specific provider has not been bound. Whether the provider becomes the default client provider, the default server provider, or both will be based on the value selected for Provider Type.
 - **Provider Type** – Select `client`, `server`, or `client-server` to establish whether the provider is to be used as a client authentication provider, a server authentication provider, or both (a client-server provider).
 - **Provider ID** - Enter an identifier for this provider configuration. This name will appear in the Current Provider Configurations list.
 - **Class Name** - Enter the Java implementation class of the provider. Client authentication providers must implement the `com.sun.enterprise.security.jauth.ClientAuthModule` interface. Server-side providers must implement the `com.sun.enterprise.security.jauth.ServerAuthModule` interface. A provider may implement both interfaces, but it must implement the interface corresponding to its provider type.
 - 9 **In the Request Policy section of the Create a Provider Configuration page, enter the following optional values, if needed.**

These properties are optional, but if not specified, no authentication is applied to request messages.

- **Authentication Source** – Select `sender`, `content`, or `null` (the blank option) to define a requirement for message-layer sender authentication (for example, username password), content authentication (for example, digital signature), or no authentication be applied to request messages. When `null` is specified, source authentication of the request is not required.
- **Authentication Recipient** – Select `beforeContent` or `afterContent` to define a requirement for message-layer authentication of the receiver of the request message to its sender (by XML encryption). When the value is not specified it defaults to `afterContent`.

For a description of the actions performed by the SOAP message security providers as a result of the following message protection policies see [“Actions of Request and Response Policy Configurations”](#) on page 208.

10 In the Response Policy section of the Create a Provider Configuration page, enter the following optional properties, if needed.

These properties are optional, but if not specified, no authentication is applied to response messages.

- **Authentication Source** – Select sender, content, or null (the blank option) to define a requirement for message-layer sender authentication (for example, username password) or content authentication (for example, digital signature) to be applied to response messages. When null is specified, source authentication of the response is not required.
- **Authentication Recipient** – Select beforeContent or afterContent to define a requirement for message-layer authentication of the receiver of the response message to its sender (by XML encryption). When the value is not specified it defaults to afterContent.

For a description of the actions performed by the SOAP message security providers as a result of the following message protection policies see [“Actions of Request and Response Policy Configurations”](#) on page 208.

11 Add additional properties by clicking the Add Property button.

The provider that is shipped with the Application Server supports the property listed below. If other providers are used, refer to their documentation for more information on properties and valid values.

- `server.config` – The directory and file name of an XML file that contains the server configuration information. For example, `domain-dir/config/wss-server-config.xml`.

12 Click OK to save this configuration, or click Cancel to quit without saving.

More Information Equivalent asadmin command

```
create-message-security-provider
```

▼ To delete a message security configuration

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, the default instance, server, select the `server-config` node.

- To configure the default settings for all instances, select the `default-config` node.
- 3 Expand the Security node.
 - 4 Select the Message Security node.
 - 5 Click in the checkbox to the left of the Message Security Configuration to be deleted.
 - 6 Click Delete.

▼ To delete a message security provider

- 1 In the Admin Console tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Expand the Security node.
- 4 Expand the Message Security node.
- 5 Select the SOAP node.
- 6 Select the Providers page.
- 7 Click in the checkbox to the left of the Provider Configuration to be deleted.
- 8 Click Delete.

More Information Equivalent `asadmin` command
`delete-message-security-provider`

▼ To enable message security for application clients

The message protection policies of client providers must be configured such that they are equivalent to the message protection policies of the server-side providers they will be interacting with. This is already the case for the providers configured (but not enabled) when the Application Server is installed.

To enable message security for client applications, modify the Application Server specific configuration for the application client container.

- 1 **Stop any client applications that depend on the client container descriptor.**
- 2 **In a text editor, open the Sun application client container descriptor, located in *domain-dir/config/sun-acc.xml*.**
- 3 **Add the `default-client-provider` element to the file to enable the default client provider in the application client.**

The other code is provided to show where the code to enable message security for client applications should be located. The other code may differ slightly in your installation. Do not change it.

```
<client-container>
  <target-server name="your-host" address="your-host"
    port="your-port"/>
  <log-service file="" level="WARNING"/>
  <message-security-config auth-layer="SOAP"
    default-client-provider="ClientProvider">
    <provider-config
      class-name="com.sun.enterprise.security.jauth.ClientAuthModule"
      provider-id="ClientProvider" provider-type="client">
      <request-policy auth-source="sender"/>
      <response-policy/>
      <property name="security.config"
        value="C:/Sun/AppServer/lib/appclient/wss-client-config.xml"/>
    </provider-config>
  </message-security-config>
</client-container>
```

The message security provider configured in the client container will also require access to private keys and trusted certificates. This is accomplished by defining appropriate values for the following system properties in the application client startup script.

```
-Djavax.net.ssl.keyStore
```

```
-Djavax.net.ssl.trustStore
```

Setting the Request and Response Policy for the Application Client Configuration

The *request and response policies* define the authentication policy requirements associated with request and response processing performed by the authentication provider. Policies are expressed in message sender order such that a requirement that encryption occur after content would mean that the message receiver would expect to decrypt the message before validating the signature.

To achieve message security, the request and response policies must be enabled on both the server and client. When configuring the policies on the client and server, make sure that the client policy matches the server policy for request/response protection at application-level message binding.

To set the request policy for the application client configuration, modify the Application Server specific configuration for the application client container as described in [“To enable message security for application clients” on page 219](#). In the application client configuration file, add the request-policy and response-policy elements as shown to set the request policy.

The other code is provided for reference. The other code may differ slightly in your installation. Do not change it.

```
<client-container>
  <target-server name="your-host" address="your-host"
    port="your-port"/>
  <log-service file="" level="WARNING"/>
  <message-security-config auth-layer="SOAP"
    default-client-provider="ClientProvider">
    <provider-config
      class-name="com.sun.enterprise.security.jauth.ClientAuthModule"
      provider-id="ClientProvider" provider-type="client">
      <request-policy auth-source="sender | content"
        auth-recipient="after-content | before-content"/>
      <response-policy auth-source="sender | content"
        auth-recipient="after-content | before-content"/>
      <property name="security.config"
        value="install-dir/lib/appclient/wss-client-config.xml"/>
    </provider-config>
  </message-security-config>
</client-container>
```

Valid values for auth-source include sender and content. Valid values for auth-recipient include before-content and after-content. A table describing the results of various combinations of these values can be found in [“Actions of Request and Response Policy Configurations” on page 208](#).

To not specify a request or response policy, leave the element blank, for example:

```
<response-policy/>
```

Further Information

- The Java 2 Standard Edition discussion of security can be viewed from <http://java.sun.com/j2se/1.4.2/docs/guide/security/index.html>.
- The *J2EE 1.4 Tutorial* chapter titled *Security* can be viewed from <http://java.sun.com/j2ee/1.4/docs/tutorial/doc/index.html>.
- The *Administration Guide* chapter titled *Chapter 9, “Configuring Security.”*
- The *Developer’s Guide* chapter titled *Securing Applications*.
- The Oasis Web Services Security: SOAP Message Security (WS-Security) specification, can be viewed from <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf>.
- The OASIS Web Services Security Username Token Profile 1.0, can be found at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-username-token-profile-1.0.pdf>.
- The OASIS Web Services Security X.509 Certificate Token Profile 1.0, can be found at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0.pdf>.
- The *XML-Signature Syntax and Processing* document can be viewed at <http://www.w3.org/TR/xmlsig-core/>.
- The *XML Encryption Syntax and Processing* document can be viewed at <http://www.w3.org/TR/xmlenc-core/>.

Transactions

By enclosing one or more steps in an indivisible unit of work, a transaction ensures data integrity and consistency. This chapter contains the following sections:

- “About Transactions” on page 223
- “Admin Console Tasks for Transactions” on page 224

About Transactions

- “What is a Transaction?” on page 223
- “Transactions in J2EE Technology” on page 224

What is a Transaction?

A transaction is a series of discreet actions in an application that must all complete successfully or else all the changes in each action are backed out. For example, to transfer funds from a checking account to a savings account is a transaction with the following steps:

1. Check to see if the checking account has enough money to cover the transfer.
2. If there’s enough money in the checking account debit the amount from the checking account.
3. Credit the money to the savings account.
4. Record the transfer to the checking account log.
5. Record the transfer to the savings account log.

If any of these steps fails, all changes from the preceding steps must be backed out, and the checking account and savings account must be in the same state as they were before the transaction started. This event is called a *rollback*. If all the steps complete successfully, the transaction is in a *committed* state. Transactions end in either a commit or a rollback.

Transactions in J2EE Technology

Transaction processing in J2EE technology involves the following five participants:

- Transaction Manager
- Application Server
- Resource Manager(s)
- Resource Adapter(s)
- User Application.

Each of these entities contribute to reliable transaction processing by implementing the different APIs and functionalities, discussed below:

- The Transaction Manager provides the services and management functions required to support transaction demarcation, transactional resource management, synchronization, and transaction context propagation.
- The Application Server provides the infrastructure required to support the application runtime environment that includes transaction state management.
- The Resource Manager (through a resource adapter) provides the application access to resources. The resource manager participates in distributed transactions by implementing a transaction resource interface used by the transaction manager to communicate transaction association, transaction completion and recovery work. An example of such a resource manager is a relational database server.
- A Resource Adapter is a system level software library that is used by the application server or client to connect to a Resource Manager. A Resource Adapter is typically specific to a Resource Manager. It is available as a library and is used within the address space of the client using it. An example of such a resource adapter is a JDBC driver.
- A Transactional User Application developed to operate in an application server environment looks up transactional data sources and, optionally, the transaction manager, using JNDI. The application may use declarative transaction attribute settings for enterprise beans or explicit programmatic transaction demarcation.

Admin Console Tasks for Transactions

The Application Server handles transactions based on the settings in the Admin Console.

Configuring Transactions

This section explains how to configure transaction settings:

- [“To configure how the Application Server recovers from transactions” on page 225](#)
- [“To set a transaction timeout value” on page 226](#)

- “To set the location of the transaction logs” on page 226
- “To set the keypoint interval” on page 227

For additional information about transactions, see these sections:

- “What is a Transaction?” on page 223
- “Transactions in J2EE Technology” on page 224

▼ To configure how the Application Server recovers from transactions

Transactions might be incomplete either because the server crashed or a resource manager crashed. It is essential to complete these stranded transactions and recover from the failures. Application Server is designed to recover from these failures and complete the transactions upon server startup.

While performing the recovery, if some of the resources are unreachable the server restart may be delayed as it tries to recover the transactions.

When the transaction spans across servers, the server that started the transaction can contact the other servers to get the outcome of the transactions. If the other servers are unreachable, the transaction uses the Heuristic Decision field to determine the outcome.

- 1 In the tree component select the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, the default instance, server, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the **Transaction Service** node.
- 4 To enable the recovery of incomplete transactions, check the **Recover** in the **On Restart** field.
- 5 Set the amount of time, in seconds, the Application Server tries to connect to the unreachable server in the **Retry Timeout** field. The default value is 10 minutes (600 seconds).
- 6 Set the policy for unreachable servers in a transaction in the **Heuristic Decision** field.

Unless there is a good reason to set this field to **Commit**, leave **Heuristic Decision** set to **Rollback**. Committing indeterminate transactions can compromise the data integrity of your application.
- 7 Click **Save**.
- 8 Restart the **Application Server**.

▼ To set a transaction timeout value

By default, the server does not timeout a transaction. That is, the server waits indefinitely for a transaction to complete. If you set a timeout value for transactions, if a transaction isn't completed within the configured time, the Application Server rolls back the transaction.

- 1 In the tree component, select the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the **Transaction Service** node.
- 4 Enter the number of seconds before the transaction times out, in the **Transaction Timeout** field. The default value of Transaction Timeout is 0 seconds. This disables transaction timeouts.
- 5 Click **Save**.
- 6 Restart the Application Server.

▼ To set the location of the transaction logs

The transaction log records the information about each transaction in order to maintain the data integrity of the resources involved and to recover from failures. Transaction logs are kept in the `tx` subdirectory of the directory specified by the Transaction Log Location field. These logs are not human readable.

- 1 In the tree component, select the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the **Transaction Service** node.
- 4 Enter the location of the transaction logs in the **Transaction Log Location** field. A `tx` subdirectory is created and transaction logs are kept under that directory.

The default value is `${com.sun.aas.instanceRoot}/logs`. The `${com.sun.aas.instanceRoot}` variable is the name of the instance, and is set when you start an Application Server instance. To see the value of `${com.sun.aas.instanceRoot}`, click Actual Values.

- 5 **Click Save.**
- 6 **Restart the Application Server.**

▼ **To set the keypoint interval**

Keypoint operations compress the transaction log file. The keypoint interval is the number of transactions between keypoint operations on the log. Keypoint operations can reduce the size of the transaction log files. A larger number of keypoint intervals (for example, 2048) results in larger transaction log files, but fewer keypoint operations, and potentially better performance. A smaller keypoint interval (for example, 256) results in smaller log files but slightly reduced performance due to the greater frequency of keypoint operations.

- 1 **In the tree component select the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.**
 - **To configure the default settings for all instances, select the `default-config` node.**
- 3 **Select the Transaction Service node.**
- 4 **Enter the number of transactions between keypoint operations in the Keypoint Interval field.**

The default value is 2048.
- 5 **Click Save.**
- 6 **Restart the Application Server.**

Configuring the HTTP Service

This chapter describes how to configure virtual servers and HTTP listeners for the HTTP service component of the Application Server.

- “About the HTTP Service” on page 229
- “Admin Console Tasks for the HTTP Service” on page 233
- “Admin Console Tasks for Virtual Servers” on page 240
- “Admin Console Tasks for HTTP Listeners” on page 244

About the HTTP Service

- “What Is the HTTP Service?” on page 229
- “Virtual Servers” on page 229
- “HTTP Listeners” on page 231

What Is the HTTP Service?

The HTTP service is the component of the Application Server that provides facilities for deploying web applications and for making deployed web applications accessible by HTTP clients. (See “[To deploy a web application](#)” on page 63.) These facilities are provided by means of two kinds of related objects, virtual servers and HTTP listeners.

Virtual Servers

A virtual server, sometimes called a virtual host, is an object that allows the same physical server to host multiple Internet domain names. All virtual servers hosted on the same physical server share the Internet Protocol (IP) address of that physical server. A virtual server associates a domain name for a server (such as `www.aaa.com`) with the particular server on which the Application Server is running.

Note – Do not confuse an Internet domain with the administrative domain of the Application Server.

For instance, assume you want to host these domains on your physical server:

```
www.aaa.com
www.bbb.com
www.ccc.com
```

Assume also that `www.aaa.com`, `www.bbb.com`, and `www.ccc.com` have web modules `web1`, `web2`, and `web3`, respectively, associated with them.

This means that all of these URLs are handled by your physical server:

```
http://www.aaa.com:8080/web1
http://www.bbb.com:8080/web2
http://www.ccc.com:8080/web3
```

The first URL is mapped to virtual host `www.aaa.com`, the second URL is mapped to virtual host `www.bbb.com`, and the third is mapped to virtual host `www.ccc.com`.

On the other hand, the following URL results in a 404 return code, because `web3` isn't registered with `www.bbb.com`:

```
http://www.bbb.com:8080/web3
```

For this mapping to work, make sure that `www.aaa.com`, `www.bbb.com`, and `www.ccc.com` all resolve to your physical server's IP address. They need to be registered with the DNS server for your network. In addition, on a UNIX system, add these domains to your `/etc/hosts` file (if the setting for `hosts` in your `/etc/nsswitch.conf` file includes `files`).

When the Application Server is started, it starts the following virtual servers automatically:

- A virtual server named `server`, which hosts all user-defined web modules
- A virtual server named `__asadmin`, which hosts all administration-related web modules (specifically, the Admin Console). This server is restricted; you cannot deploy web modules to this virtual server.

For development, testing, and deployment of web services in a non-production environment, `server` is often the only virtual server required. In a production environment, additional virtual servers provide hosting facilities for users and customers so that each appears to have its own web server, even though there is only one physical server.

HTTP Listeners

Each virtual server provides connections between the server and clients through one or more HTTP listeners. Each HTTP listener is a listen socket that has an IP address, a port number, a server name, and a default virtual server.

HTTP listeners must have a unique combination of port number and IP address. For example, an HTTP listener can listen on all configured IP addresses on a given port for a machine by specifying the IP address 0.0.0.0. Alternatively, the HTTP listener can specify a unique IP address for each listener, but use the same port.

Since an HTTP listener is a combination of IP address and port number, you can have multiple HTTP listeners with the same IP address and different port numbers (for example, 1.1.1.1:8081 and 1.1.1.1:8082), or with different IP addresses and the same port number (for example, 1.1.1.1:8081 and 1.2.3.4:8081, if your machine was configured to respond to both these addresses).

However, if an HTTP listener uses the 0.0.0.0 IP address, which listens on all IP addresses on a port, you cannot create HTTP listeners for additional IP addresses that listen on the same port for a specific IP address. For example, if an HTTP listener uses 0.0.0.0:8080 (all IP addresses on port 8080), another HTTP listener cannot use 1.2.3.4:8080.

Because the system running the Application Server typically has access to only one IP address, HTTP listeners typically use the 0.0.0.0 IP address and different port numbers, with each port number serving a different purpose. If the system does have access to more than one IP address, each address can serve a different purpose.

By default, when the Application Server starts, it has the following HTTP listeners:

- Two HTTP listeners named `http-listener-1` and `http-listener-2`, associated with the virtual server named `server`. The listener named `http-listener-1` does not have security enabled; `http-listener-2` has security enabled.
- An HTTP listener named `admin-listener`, associated with the virtual server named `__asadmin`. This listener has security enabled.

All these listeners use the IP address 0.0.0.0 and the port numbers specified as the HTTP server port numbers during installation of the Application Server. If the Application Server uses the default port number values, `http-listener-1` uses port 8080, `http-listener-2` uses port 8181, and `admin-listener` uses port 4849.

Each HTTP listener has a default virtual server. The default virtual server is the server to which the HTTP listener routes all request URLs whose host component does not match any of the virtual servers that are associated with the HTTP listener (a virtual server is associated with an HTTP listener by listing the HTTP listener in its `http-listeners` attribute).

In addition, specify the number of acceptor threads in the HTTP listener. Acceptor threads are threads that wait for connections. The threads accept connections and put them in a queue, called the connection queue, where they are then picked up by worker threads. Configure

enough acceptor threads so that there is always one available when a new request comes in, but few enough so that they do not provide too much of a burden on the system. The connection queue includes both the new connections just accepted by acceptor threads and persistent connections managed by the Keep-Alive connection management subsystem.

A set of request processing threads retrieves incoming HTTP requests from the connection queue and processes the requests. These threads parse the HTTP headers, select the appropriate virtual server, and run through the request processing engine to service the request. When there are no more requests to process, but the connection can be kept persistent (either by using HTTP/1.1 or sending a `Connection: keep-alive` header), the request processing thread assumes the connection to be idle and passes the connection to the Keep-Alive connection management subsystem.

The Keep-Alive subsystem periodically polls such idle connections and queues those connections with activity into the connection queue for future processing. From there, a request processing thread again retrieves the connection and processes its request. The Keep-Alive subsystem is multi-threaded, as it manages potentially tens of thousands of connections. Efficient polling techniques are used, by dividing the number of connections into smaller subsets, to determine which connections are ready with requests and which of those connections have idled for sufficient time to deem them closed (beyond a maximum permissible Keep-Alive timeout).

The HTTP listener's server name is the host name that appears in the URLs the server sends to the client as part of a redirect. This attribute affects URLs the server automatically generates; it does not affect the URLs for directories and files stored in the server. This name is normally the alias name if the server uses an alias. If a client sends a `Host:` header, that host name supersedes the HTTP listener's server name value in redirects.

Specify a redirect port to use a different port number from that specified in the original request. A *redirect* occurs in one of these situations:

- If a client tries to access a resource that no longer exists at the specified URL (that is, the resource has moved to another location), the server redirects the client to the new location (instead of returning a 404), by returning a designated response code and including the new location in the response's `Location` header.
- If a client tries to access a resource that is protected (for example, SSL) on the regular HTTP port, the server redirects the request to the SSL-enabled port. In this case, the server returns a new URL in the `Location` response header, in which the original insecure port has been replaced with the SSL-enabled port. The client then connects to this new URL.

Specify also whether security is enabled for an HTTP listener and what kind of security is used (for example, which SSL protocol and which ciphers).

To access a web application deployed on the Application Server, use the URL `http://localhost:8080/` (or `https://localhost:8181/` if it is a secure application), along

with the context root specified for the web application. To access the Admin Console, use the URL `https://localhost:4849/` or `https://localhost:4849/asadmin/` (its default context root).

Because a virtual server must specify an existing HTTP listener, and because it cannot specify an HTTP listener that is already being used by another virtual server, create at least one HTTP listener before creating a new virtual server.

Admin Console Tasks for the HTTP Service

- “To configure the HTTP Service” on page 233
- “To configure the HTTP Service access log” on page 235

▼ To configure the HTTP Service

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, for the default instance, server, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Select the HTTP Service node.
- 4 On the HTTP Service page, you can set properties that apply to all of the service’s HTTP listeners.

The following table lists these properties.

Property Name	Description	Default Value
<code>traceEnabled</code>	If set to true, enables the TRACE operation. Set this property to false to make the Application Server less susceptible to cross-site scripting attacks.	false

Property Name	Description	Default Value
monitoringCacheEnabled	<p>If set to true, the Application Server caches local values of the statistics of the HTTP Service to answer statistics queries. This value improves performance.</p> <p>If set to false, the Application Server queries the HTTP Service for each statistics value.</p>	true
monitoringCacheRefreshInMillis	Specifies the time interval, in milliseconds, after which the monitoring cache is refreshed.	5000
sslCacheEntries	Specifies the number of SSL sessions that can be cached. There is no upper limit.	10000
sslSessionTimeout	Specifies the number of seconds before an SSL2 session times out.	100
ssl3SessionTimeout	Specifies the number of seconds before an SSL3 session times out.	86400
sslClientAuthDataLimit	Specifies the maximum amount of application data, in bytes, that is buffered during the client certificate handshake phase.	1048576
sslClientAuthTimeout	Specifies the number of seconds before the client certificate handshake phase times out.	60
keepAliveQueryMeanTime	Specifies the desired keep-alive latency (in milliseconds).	100
keepAliveQueryMaxSleepTime	Specifies the upper limit to the time slept (in milliseconds) after polling keep-alive connections for further requests.	100
stackSize	Specifies the maximum stack size of the native thread.	OS/Machine dependent
statsProfilingEnabled	If set to false, disables the recording of monitoring statistics by the HTTP Service, which improves performance. If this property is set to false, enabling monitoring for the HTTP service has no effect.	true
chunkedRequestBufferSize	Specifies the default buffer size, in bytes, for unchunking request data.	8192

Property Name	Description	Default Value
chunkedRequestTimeoutSeconds	Specifies the default timeout (in seconds) for unchunking request data.	60
dnsCacheEnabled	If set to true, allows the user to monitor statistics related to DNS caching. This property takes effect only if the DNS Lookup box on the HTTP Protocol tab is selected. Otherwise, the property setting is ignored.	false

5 Click the Access Log tab to configure access log rotation.

Click the other tabs to configure request processing, the Keep-Alive subsystem, the connection pool, the HTTP protocol, and the HTTP file cache.

6 Click Save.

▼ To configure the HTTP Service access log

Use this page to enable and configure rotation for the access logs for the virtual servers. These logs are in the *domain-dir/logs/access* directory and are named as follows:
virtual-server-name_access_log%YYYY;%MM;%DD-%hh;h:mm;m:ss;s.

Click Default to load the default values.

1 Check the File Rotation box to turn on file rotation.

By default, file rotation is enabled.

2 From the Rotation Policy drop-down list, choose a policy.

The only policy available is *time*.

3 In the Rotation Interval field, type a numeric value to specify the number of minutes between rotations of the access log.

This field is valid only if the Rotation Policy is *time*. The default is 1440 minutes.

4 In the Rotation Suffix field, type a string value to specify the suffix to be added to the log file name after rotation.

The default is *%YYYY;%MM;%DD;-%hh;h:mm;m:ss;s.*

5 In the Format field, enter a string value to specify the format of the access log.

Use the formats shown in the following table. The default format is *%client.name% %auth-user-name% %datetime% %request% %status% %response.length%*.

Data	Token
Client Host Name	%client.name%
Client DNS	%client.dns%
System Date	%datetime%
Full HTTP Request line	%request%
Status	%status%
Response Content Length	%response.length%
Referer Header	%header.referer%
User-agent	%header.user-agent%
HTTP Method	%http-method%
HTTP URI	%http-uri%
HTTP Query String	%query-str%
HTTP Protocol Version	%http-version%
Accept Header	%header.accept%
Date Header	%header.date%
If-Modified-Since Header	%header.if-mod-since%
Authorization Header	%header.auth%
Any valid HTTP header value defined in RFC 2616 (any is also a valid header value; it is specified as a variable here)	%header.any%
Name of Authorized User	%auth-user-name%
Value of a Cookie	%cookie.value%
Virtual Server ID	%vs.id%

- 6 Click **Save** to save the changes, or **Load Defaults** to return to the default settings.

▼ To configure HTTP Service request processing threads

- 1 Click **Load Defaults** to load the default values.
- 2 In the **Thread Count** field, type a numeric value to specify the maximum number of request processing threads.

The default is 128.

- 3 **In the Initial Thread Count field, type the number of request processing threads that will be available when the server starts.**
The default is 48.
- 4 **In the Thread Increment field, type the number of request processing threads to be added when the number of requests exceeds the initial thread count.**
The default is 10.
- 5 **In the Request Timeout field, type the number of seconds after which requests will time out.**
The default is 30 seconds.
- 6 **In the Buffer Length field, type the size (in bytes) of the buffer used by the request processing threads to read the request data.**
The default is 4096 bytes.
- 7 **Click Save to save the changes, or Load Defaults to return to the default settings.**

▼ **To configure the HTTP Service keep-alive subsystem**

- 1 **Click Load Defaults to load the default values.**
- 2 **In the Thread Count field, type the number of keep-alive threads to be used.**
The default is 1.
- 3 **In the Max Connections field, type the maximum number of persistent connections to be maintained.**
The default is 256.
- 4 **In the Time Out field, type the maximum number of seconds for which a keep-alive connection should be kept open.**
The default is 30 seconds.
- 5 **Click Save to save the changes, or Load Defaults to return to the default settings.**

▼ **Configuring the HTTP Service Connection Pool**

- 1 **Click Load Defaults to load the default values.**

- 2 In the Max Pending Count field, type the maximum number of pending connections to be allowed for an HTTP listener.**

The default that the Application Server software specifies is 4096.

Note – The actual value for the default maximum pending count depends on the operating system. For example, for the Solaris operating system (OS), the default maximum pending count is determined by the `tcp_conn_req_max_q` kernel parameter of the TCP driver. The value of this parameter is 128. To determine the default maximum pending count on the Solaris OS, type the command `ndd /dev/tcp tcp_conn_req_max_q`.

- 3 In the Queue Size field, type the maximum number of outstanding connections the server can have**

The default is 4096.

- 4 In the Receive Buffer Size field, type the size of the receive buffer for an HTTP listener.**

The default is 4096.

- 5 In the Send Buffer Size field, type the size of the send buffer for an HTTP listener.**

The default is 8192.

- 6 Click Save to save the changes, or Load Defaults to return to the default settings.**

▼ To configure the HTTP protocol for the HTTP Service

- 1 Click Load Defaults to load the default values.**

- 2 In the Version field, type the version of the HTTP protocol to be used (HTTP/1.0 or HTTP/1.1).**

The default is HTTP/1.1.

- 3 Select the DNS Lookup box to enable lookup of the DNS entry for the client.**

The default is false.

- 4 Remove the check from the SSL box to disable security in the server globally.**

Leave this value set to true to be able to use SSL for any listener that has security enabled. The default is true.

- 5 In the Forced Response Type field, type the response type to be used if there is no MIME mapping available matching the extension.**

The default value is `text/html; charset=iso-8859-1`.

- 6 In the Default Response Type field, type the default response type.**

The default value is `text/html; charset=iso-8859-1`. The value is a semicolon-delimited string consisting of content-type, encoding, language, and charset.

- 7 Click Save to save the changes, or Load Defaults to return to the default settings.**

▼ To configure the HTTP file cache for the HTTP Service

The file cache stores static content so that the server handles requests for such content quickly.

- 1 Click Load Defaults to load the default values.**

- 2 Check the Globally box to enable file caching.**

The default is true.

- 3 Check the File Transmission box to enable the use of the `TransmitFileSystem` method on Windows.**

The default is false.

- 4 In the Max Age field, type the maximum age, in seconds, of a valid cache entry.**

The default is 30 seconds.

- 5 In the Max Files Count field, type the maximum number of files in the file cache.**

The default is 1024.

- 6 In the Hash Init Size field, type the initial number of hash buckets.**

The default is zero.

- 7 In the Medium File Size Limit field, type the maximum size, in bytes, of a file that can be cached as a memory mapped file.**

The default is 537,600 bytes.

- 8 In the Medium File Size field, type the total size, in bytes, of all files that are cached as memory mapped files.**

The default is 10,485,760 bytes.

- 9 In the Small File Size Limit field, type the maximum size, in bytes, of a file that can be read into memory.**

The default is 2048 bytes.

- 10 **In the Small File Size field, type the total size, in bytes, of all files that are read into memory.**
The default is 1,048,576 bytes.
- 11 **Choose ON or OFF from the File Caching Enabled drop-down list to enable or disable caching of file content if the size of the file is less than the medium file size limit.**
The default is ON.
- 12 **Click Save to save the changes, or Load Defaults to return to the default settings.**

Admin Console Tasks for Virtual Servers

- [“To create a virtual server” on page 240](#)
- [“To edit a virtual server” on page 243](#)
- [“To delete a virtual server” on page 243](#)

▼ To create a virtual server

- 1 **In the tree component, expand the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance’s config node. For example, for the default instance, server, select the server-config node.**
 - **To configure the default settings for future instances that use a copy of default-config, select the default-config node.**
- 3 **Expand the HTTP Service node.**
- 4 **Select the Virtual Servers node.**
- 5 **On the Virtual Servers page, click New.**
The Create Virtual Server page appears.
- 6 **In the ID field, type a unique name for the virtual server.**
This value is used to identify the virtual server internally. It is not exposed to HTTP clients. The host names that are exposed to HTTP clients must be specified in the Hosts field.
- 7 **In the Hosts field, type the host name or names for the machine on which the server is running.**
Use either actual or virtual host names that are registered with the DNS server for your network (and, on a UNIX system, in your /etc/hosts file).

8 In the area opposite State, select either On, Off, or Disabled.

The default is On.

9 Leave the HTTP Listeners field empty.

It is filled in automatically when you create an HTTP listener and associate it with this server.

Use of this field requires that you specify an existing HTTP listener. You must not, however, specify a listener that is used by another virtual server; if you do, an error appears in the server log. Since a listener must be associated with an existing virtual server when it is created, all existing listeners are used by another virtual server.

10 From the Default Web Module drop-down list, choose the deployed web module (if any) that is to respond to all requests that cannot be mapped to other web modules deployed to the virtual server.

If a Default Web Module is not specified, the web module that has an empty context root is used. If there is no web module with an empty context root, a system default web module is created and used.

11 In the Log File field, type the path name of the file where logging messages from this virtual server will appear.

Leave this field empty to send logging messages to the default server log, *domain-dir/logs/server.log*.

12 In the Additional Properties area, click Add Property to add a property for the virtual server.

Whether you specify properties or not, the new server has the default properties `docroot` and `accesslog`, set to default values.

The following table lists the available virtual server properties.

Property Name	Description
<code>docroot</code>	Absolute path to root document directory for server. Default is <i>domain-dir/docroot</i> .
<code>accesslog</code>	Absolute path to server access logs. Default is <i>domain-dir/logs/access</i> .

Property Name	Description
<code>sso-enabled</code>	<p>If false, single sign-on is disabled for this virtual server, and users must authenticate separately to every application on the virtual server.</p> <p>Single sign-on across applications on the Application Server is supported by servlets and JSP pages. This feature allows multiple applications that require the same user sign-on information to share this information, rather than have the user sign on separately for each application.</p> <p>Default is true.</p>
<code>sso-max-inactive-seconds</code>	<p>Specifies the number of seconds 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.</p> <p>Default is 300 seconds (5 minutes). Higher values provide longer single sign-on persistence for users at the expense of more memory use on the server.</p>
<code>sso-reap-interval-seconds</code>	<p>Specifies the number of seconds between purges of expired single sign-on records.</p> <p>Default is 60.</p>
<code>allowLinking</code>	<p>If true, resources that are symbolic links will be served for all web applications deployed on this virtual server. Individual web applications may override this setting by using the <code>sun-web-app</code> property <code>allowLinking</code> in the <code>sun-web.xml</code> file:</p> <pre><sun-web-app> <property name="allowLinking" value="{true false}"/> </sun-web-app></pre> <p>Default is true.</p>

13 Click OK to save the virtual server.

More Information Equivalent `asadmin` command

`create-virtual-server`

▼ To edit a virtual server

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Expand the HTTP Service node.
- 4 Select the Virtual Servers node.
- 5 Select the virtual server to be edited.
- 6 On the Edit Virtual Server page, you can perform these tasks:
 - Change the host name in the Hosts field.
 - Change the value of the State setting.
 - Add or remove an HTTP listener.
 - Change the Default Web Module selection.
 - Change the Log File value.
 - Add, remove, or modify properties.
- 7 Click Save to save the changes.

▼ To delete a virtual server

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.

- **To configure the default settings for future instances that use a copy of default-config, select the default-config node.**
- 3 **Expand the HTTP Service node.**
 - 4 **Select the Virtual Servers node.**
 - 5 **On the Virtual Servers page, check the box next to the name of the virtual server to be deleted.**
 - 6 **Click Delete.**

It is possible to delete the __asadmin virtual server, but this is not recommended. If you plan to do so, first copy the virtual-server elements of the Application Server's domain.xml file to a safe place so that the settings can be restored if needed.

More Information Equivalent asadmin command

```
delete-virtual-server
```

Admin Console Tasks for HTTP Listeners

- [“To create an HTTP listener” on page 244](#)
- [“To edit an HTTP listener” on page 246](#)
- [“To delete an HTTP listener” on page 247](#)

▼ To create an HTTP listener

- 1 **In the tree component, expand the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance's config node. For example, for the default instance, server, select the server-config node.**
 - **To configure the default settings for future instances that use a copy of default-config, select the default-config node.**
- 3 **Expand the HTTP Service node.**
- 4 **Select the HTTP Listeners node.**
- 5 **On the HTTP Listeners page, click New.**

The Create HTTP Listener page appears.

- 6 In the **Name** field, type a name for the listener.
- 7 In the **Listener** field, remove the check from the **Enabled** box if you do not want to enable the listener when the server restarts.
- 8 In the **Network Address** field, type 0.0.0.0 if you want the listener to listen on all IP addresses for the server, using a unique port value. Otherwise, type a valid IP address for the server.
- 9 In the **Listener Port** field, type a unique port value if the **Network Address** field is 0.0.0.0, or the desired port value if you are using another IP address.
- 10 Choose a virtual server from the **Default Virtual Server** drop-down list.
- 11 In the **Server Name** field, type the host name to be used in the URLs the server sends to the client. This name is the alias name if your server uses an alias.
If your server does not use an alias, leave this field empty.
- 12 In the **Advanced** area, perform any of the following tasks:
 - To redirect requests to another port, type a value in the **Redirect Port** field. The Application Server automatically redirects the request if these two conditions exist:
 - This listener is supporting non-SSL requests.
 - A request is received for which a matching security constraint requires SSL transport.
By default, the Application Server uses the port number specified in the original request.
 - Change the number of **Acceptor Threads**.
 - Remove the check from the **Powered By** box to disable the inclusion of the **X-Powered-By: Servlet/2.4** header in servlet-generated HTTP response headers.
The Java Servlet 2.4 Specification defines this header, which containers may add to servlet-generated responses. Similarly, the JavaServer Pages™ (JSP™) 2.0 Specification defines an **X-Powered-By: JSP/2.0** header to be added (on an optional basis) to responses that use JSP technology. The inclusion of the **X-Powered-By: JSP/2.0** header is enabled by default for web applications. The goal of these headers is to aid web site administrators in gathering statistical data about the use of Servlet and JSP technology.
For information on enabling and disabling the **X-Powered-By** header for JSP pages, see the chapter entitled “Deployment Descriptor Files” in the *Application Server Developer’s Guide*. See “[Further Information](#)” on page 55 for a link to this document.
Production environments might decide to omit the generation of **X-Powered-By** headers to hide their underlying technology.
- 13 To create a listener that is not secure, click **OK**.

- 14 In the SSL section of this page, you can configure the listener to use SSL, TLS, or both SSL and TLS security.
To set up a secure listener, do the following:
- 15 Check the Enabled box in the Security field.
- 16 To force clients to authenticate themselves to the server when using this listener, check the Enabled box in the Client Authentication field.
- 17 Enter the name of an existing server key pair and certificate in the Certificate NickName field. See the Security chapter for more information.
- 18 In the SSL3/TLS section:
 - a. Check the security protocol(s) to be enabled on the listener. Check either SSL3 or TLS, or both.
 - b. Check the cipher suite used by the protocol(s). To enable all cipher suites, check All Supported Cipher Suites. You can also enable individual cipher suites.
- 19 The listener is now listed in the HTTP Listeners field for the virtual server that is specified as the Default Virtual Server.

More Information Equivalent asadmin commands

`create-http-listener` and `create-ssl`

▼ To edit an HTTP listener

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Expand the HTTP Service node.
- 4 Select the HTTP Listeners node.

- 5 Select the HTTP listener to be edited.
- 6 On the Edit HTTP Listener page, modify any of the settings.
- 7 Click Save to save the changes.

▼ To delete an HTTP listener

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Expand the HTTP Service node.
- 4 Select the HTTP Listeners node.
- 5 On the HTTP Listeners page, check the box next to the name of the HTTP listener to be deleted.
- 6 Click Delete.

It is possible to delete the `http-listener-1`, `http-listener-2`, and `admin-listener` HTTP listeners, but this is not recommended. If you plan to do so, first copy the `http-listener` elements of the Application Server's `domain.xml` file to a safe place so that the settings can be restored if needed.

More Information Equivalent `asadmin` command

```
delete-http-listener
```


Configuring the Object Request Broker

This chapter describes how to configure the Object Request Broker (ORB) and IIOP listeners. It has the following sections:

- “About the Object Request Broker” on page 249
- “Admin Console Tasks for the ORB” on page 250
- “Admin Console Tasks for IIOP Listeners” on page 251

About the Object Request Broker

- “CORBA” on page 249
- “What is the ORB?” on page 250
- “IIOP Listeners” on page 250

CORBA

The Application Server supports a standard set of protocols and formats that ensure interoperability. Among these protocols are those defined by CORBA.

The CORBA (Common Object Request Broker Architecture) model is based on clients requesting services from distributed objects or servers through a well-defined interface by issuing requests to the objects in the form of remote method requests. A remote method request carries information about the operation that needs to be performed, including the object name (called an object reference) of the service provider and parameters, if any, for the invoked method. CORBA automatically handles network programming tasks such as object registration, object location, object activation, request de-multiplexing, error-handling, marshalling, and operation dispatching.

What is the ORB?

The Object Request Broker (ORB) is the central component of CORBA. The ORB provides the required infrastructure to identify and locate objects, handle connection management, deliver data, and request communication.

A CORBA object never talks directly with another. Instead, the object makes requests through a remote stub to the ORB running on the local machine. The local ORB then passes the request to an ORB on the other machine using the Internet Inter-Orb Protocol (IIOP for short). The remote ORB then locates the appropriate object, processes the request, and returns the results.

IIOP can be used as a Remote Method Invocation (RMI) protocol by applications or objects using RMI-IIOP. Remote clients of enterprise beans (EJB modules) communicate with the Application Server via RMI-IIOP.

IIOP Listeners

An IIOP listener is a listen socket that accepts incoming connections from the remote clients of enterprise beans and from other CORBA-based clients. Multiple IIOP listeners can be configured for the Application Server. For each listener, specify a port number, a network address, and optionally, security attributes. For more information, see [“To create an IIOP listener” on page 251](#).

Admin Console Tasks for the ORB

- [“To configure the ORB” on page 250](#)

▼ To configure the ORB

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, for the default instance, server, select the server-config node.
 - To configure the default settings for future instances that use a copy of default-config, select the default-config node.
- 3 Select the ORB node.

- 4 **Choose the thread pool the ORB uses from the Thread Pool ID drop-down list.**
The ORB uses thread pools to respond to requests from remote clients of enterprise beans and other clients that communicate via RMI-IIOP. For more information, see [“About Thread Pools” on page 255](#) and [“To create a thread pool” on page 256](#).
- 5 **In the Max Message Fragment Size field, set the maximum fragment size for IIOP messages.**
Messages larger than this size are fragmented.
- 6 **In the Total Connections field, set the maximum number of incoming connections for all IIOP listeners.**
- 7 **Select the Required checkbox if IIOP client authentication is required.**
- 8 **Click Save to save the changes, or Load Defaults to load the default values.**
- 9 **Restart the server.**

Admin Console Tasks for IIOP Listeners

- [“To create an IIOP listener” on page 251](#)
- [“To edit an IIOP listener” on page 253](#)
- [“To delete an IIOP listener” on page 253](#)

▼ To create an IIOP listener

- 1 **In the tree component, expand the Configurations node.**
- 2 **Select the instance to configure:**
 - **To configure a particular instance, select the instance’s config node. For example, for the default instance, server, select the server-config node.**
 - **To configure the default settings for future instances that use a copy of default-config, select the default-config node.**
- 3 **Expand the ORB node.**
- 4 **Select IIOP Listeners.**
- 5 **Click New.**
- 6 **Enter a name to identify the listener in the Name field.**

7 Enter the network address of the listener in the Network Address field.

This can be an IP address or a DNS resolvable host name.

8 In the Listener Port field, enter the port number upon which the listener is to listen.

9 Check the Enabled box in the Listener field to enable the listener.

10 In the Additional Properties area, provide values for properties required by applications.

11 To create a listener:

- **To create a listener that is not secure, click OK.**
- **To set up a secure listener, do the following:**
 - a. **Check the Enabled box in the Security field.**
 - b. **To force clients to authenticate themselves to the server when using this listener, check the Enabled box in the Client Authentication field.**
 - c. **Enter the name of an existing server key pair and certificate in the Certificate NickName field.**
 - d. **In the SSL3/TLS section:**
 - i. **Check the security protocol(s) to enable on the listener. Check either SSL3 or TLS, or enable both protocols.**
 - ii. **Check the cipher suite used by the protocol(s).**

To enable all cipher suites, check All Supported Cipher Suites. You can also enable individual cipher suites.
 - e. **Click OK.**

The listener is now listed in the Current Listeners table on the IIOP Listeners page.

More Information Equivalent asadmin commands

`create-iiop-listener` and `create-ssl`

▼ To edit an IIOP listener

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Expand the ORB node.
- 4 Select the IIOP Listeners node.
- 5 Select the listener to be modified in the Current Listeners table.
- 6 Modify the listener's settings.
See [“To create an IIOP listener” on page 251](#) for descriptions of the fields that are modifiable.
- 7 If you changed the port number of the listener, restart the server.

▼ To delete an IIOP listener

- 1 In the tree component, expand the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 Expand the ORB node.
- 4 Select the IIOP Listeners node.
- 5 Check the listener(s) to be deleted in the Current Listeners table.
- 6 Click Delete.

More Information Equivalent asadmin command

`delete-iiop-listener`

Thread Pools

This chapter describes how to create, edit, and delete thread pools. It has the following sections:

- [“About Thread Pools” on page 255](#)
- [“Admin Console Tasks for Thread Pools” on page 256](#)

About Thread Pools

This section describes thread pools and how the work in the Application Server.

The Java Virtual Machine (JVM) can support many threads of execution at once. To help performance, the Application Server maintains one or more thread pools. It is possible to assign specific thread pools to connector modules and to the ORB.

One thread pool can serve multiple connector modules and enterprise beans. Request threads handle user requests for application components. When the server receives a request, it assigns the request to a free thread from the thread pool. The thread executes the client's requests and returns results. For example, if the request needs to use a system resource that is currently busy, the thread waits until that resource is free before allowing the request to use that resource.

Specify the minimum and maximum number of threads that are reserved for requests from applications. The thread pool is dynamically adjusted between these two values. The minimum thread pool size that is specified signals the server to allocate at least that many threads in reserve for application requests. That number is increased up to the maximum thread pool size that is specified.

Increasing the number of threads available to a process allows the process to respond to more application requests simultaneously.

Avoid thread starvation, where one resource adapter or application occupies all threads in the Application Server, by dividing the Application Server threads into different thread-pools.

Admin Console Tasks for Thread Pools

- [“To create a thread pool” on page 256](#)
- [“To edit a thread pool” on page 257](#)
- [“To delete a thread pool” on page 258](#)

▼ To create a thread pool

- 1 In the tree component select the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance’s config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the **Thread Pools** node.
- 4 Under **Current Pools** click **New**.
- 5 Enter the name of the thread pool in the **Thread Pool ID** field.
- 6 Enter the minimum number of threads in the thread pool servicing requests in this queue in the **Minimum Thread Pool Size** field.

These threads are created up front when this thread pool is instantiated.
- 7 Enter the maximum number of threads in the thread pool servicing requests in this queue in the **Maximum Thread Pool Size** field.

This is the upper limit on the number of threads that exist in the thread pool.
- 8 Enter the number, in seconds, after which idle threads are removed from pool in the **Idle Timeout** field.
- 9 Enter the total number of work queues that are serviced by this thread pool in the **Number of Work Queues** field.
- 10 Click **OK**.
- 11 Restart the **Application Server**.

More Information Equivalent asadmin command

```
create-threadpool
```

▼ To edit a thread pool

- 1 In the tree component select the Configurations node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the Thread Pools node.
- 4 Under Current Pools select the name of the thread pool to be changed.
- 5 Enter the minimum number of threads in the thread pool servicing requests in this queue in the Minimum Thread Pool Size field.

These threads are created up front when this thread pool is instantiated.
- 6 Enter the maximum number of threads in the thread pool servicing requests in this queue in the Maximum Thread Pool Size field.

This is the upper limit on the number of threads that exist in the thread pool.
- 7 Enter the number, in seconds, after which idle threads are removed from pool in the Idle Timeout field.
- 8 Enter the total number of work queues that are serviced by this thread pool in the Number of Work Queues field.
- 9 Click Save.
- 10 Restart the Application Server.

▼ To delete a thread pool

- 1 In the tree component, select the **Configurations** node.
- 2 Select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, the default instance, `server`, select the `server-config` node.
 - To configure the default settings for all instances, select the `default-config` node.
- 3 Select the **Thread Pools** node.
- 4 Check the thread pool name to be deleted in the **Current Pools** table.
- 5 Click **Delete**.
- 6 Restart the **Application Server**.

More Information Equivalent `asadmin` command
`delete-threadpool`

Configuring Logging

This chapter briefly describes how to use the Admin Console to configure logging and view the server log. It contains the following sections:

- “About Logging” on page 259
- “Admin Console Tasks for Logging” on page 262

About Logging

- “Log Records” on page 259
- “The Logger Namespace Hierarchy” on page 260

Log Records

The Application Server uses the Java 2 platform Logging API specified in JSR 047. Application Server logging messages are recorded in the server log, normally found at *domain-dir/logs/server.log*.

The *domain-dir/logs* directory contains two other kinds of logs in addition to the server log. In the *access* subdirectory are the HTTP Service access logs, and in the *tx* subdirectory are the Transaction Service logs. For information about these logs, see “[To configure the HTTP Service access log](#)” on page 235 and “[Configuring Transactions](#)” on page 224.

The components of the Application Server generate logging output. Application components can also generate logging output.

Application components may use the Apache Commons Logging Library to log messages. The platform standard JSR 047 API, however, is recommended for better log configuration.

Log records follow a uniform format:

```
[#|yyyy-mm-ddThh:mm:ss.SSS-Z|Log Level|ProductName-Version|LoggerName|Key Value Pairs|Message|#]
```

For example:

```
[#|2004-10-21T13:25:53.852-0400|INFO|sun-appserver-e8.1|javax.enterprise.
system.core|_ThreadID=13;|CORE5004: Resource Deployed:
[cr:jms/DurableConnectionFactory].|#]
```

In this example,

- [# and #] mark the beginning and end of the record.
- The vertical bar (|) separates the record fields.
- 2004-10-21T13:25:53.852-0400 specifies the date and time.
- The *Log Level* is INFO. This level may have any of the following values: SEVERE, WARNING, INFO, CONFIG, FINE, FINER, and FINEST.
- The *ProductName-Version* is sun-appserver-ee8.1.
- The *LoggerName* is a hierarchical logger namespace that identifies the source of the log module, in this case javax.enterprise.system.core.
- The *Key Value Pairs* are key names and values, typically a thread ID such as _ThreadID=14;.
- The *Message* is the text of the log message. For all Application Server SEVERE and WARNING messages and many INFO messages, it begins with a message ID that consists of a module code and a numerical value (in this case, CORE5004).

The log record format might be changed or enhanced in future releases.

The Logger Namespace Hierarchy

The Application Server provides a logger for each of its modules. The following table lists the names of the modules and the namespace for each logger in alphabetical order, as they appear on the Log Levels page of the Admin Console (see [“To configure log levels” on page 263](#)). The last three modules in the table do not appear on the Log Levels page.

TABLE 15-1 Application Server Logger Namespaces

Module Name	Namespace
Admin	javax.enterprise.system.tools.admin
ClassLoader	javax.enterprise.system.core.classloading
CMP	javax.enterprise.system.container.cmp
Configuration	javax.enterprise.system.core.config
Connector	javax.enterprise.resource.resourceadapter

TABLE 15-1 Application Server Logger Namespaces (Continued)

Module Name	Namespace
CORBA	<code>javax.enterprise.resource.corba</code>
Deployment	<code>javax.enterprise.system.tools.deployment</code>
EJB Container	<code>javax.enterprise.system.container.ejb</code>
JavaMail	<code>javax.enterprise.resource.javamail</code>
JAXR	<code>javax.enterprise.resource.webservices.registry</code>
JAX-RPC	<code>javax.enterprise.resource.webservices.rpc</code>
JDO	<code>javax.enterprise.resource.jdo</code>
JMS	<code>javax.enterprise.resource.jms</code>
JTA	<code>javax.enterprise.resource.jta</code>
JTS	<code>javax.enterprise.system.core.transaction</code>
MDB Container	<code>javax.enterprise.system.container.ejb.mdb</code>
Naming	<code>javax.enterprise.system.core.naming</code>
Node Agent (Enterprise Edition only)	<code>javax.ee.enterprise.system.nodeagent</code>
Root	<code>javax.enterprise</code>
SAAJ	<code>javax.enterprise.resource.webservices.saaj</code>
Security	<code>javax.enterprise.system.core.security</code>
Server	<code>javax.enterprise.system</code>
Synchronization (Enterprise Edition only)	<code>javax.ee.enterprise.system.tools.synchronization</code>
Util	<code>javax.enterprise.system.util</code>
Verifier	<code>javax.enterprise.system.tools.verifier</code>
Web Container	<code>javax.enterprise.system.container.web</code>
Core	<code>javax.enterprise.system.core</code>
System Output (<code>System.out.println</code>)	<code>javax.enterprise.system.stream.out</code>
System Error (<code>System.err.println</code>)	<code>javax.enterprise.system.stream.err</code>

Admin Console Tasks for Logging

- “To configure general logging settings” on page 262
- “To configure log levels” on page 263
- “To view the server log” on page 264

▼ To configure general logging settings

You can use the `com.sun.enterprise.server.logging.max_history_files` system property to limit the number of rotated log files for both access logging and the server log. The property specifies the maximum number of access log files to keep, starting with the most recent ones.

- If the property is not set, then no history files cleanup is done.
- If the property is set and it's either an invalid number or null, then a default of 10 history files is kept.
- If the property is set to 0, then no history files are kept.

In the Admin Console, access the `config/server` that this system property needs to be set for. On the command line, use the `create-jvm-option` command.

- 1 In the tree component, expand the **Node Agents or Configurations** node.
- 2 Select a node agent, or select the instance to configure:
 - To configure a particular instance, select the instance's config node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 For a node agent, select the **Logger Settings** tab. For a configuration, select the **Logger Settings** node.
- 4 On the **Logging Settings** page, use the following fields to customize logging:
 - **Log File** — To specify an alternative name or location for the server log file, type the new path name in the text field. By default, the location is `domain-dir/logs/server.log`.
 - **Alarms** — To route SEVERE and WARNING messages through the JMX framework, select the Enabled checkbox.
 - **Write to System Log** — On Solaris and Linux systems only, to send logging output to the `syslog` facility in addition to the server log, select the Enabled checkbox.

- **Log Handler** — To send logs to a destination other than `server.log` or `syslog`, you can plug in a custom log handler. The custom handler must extend the class `java.util.logging.Handler` (a JSR 047 compliant API). Type the absolute class name of the handler in the Log Handler field. Also put the handler class in the Application Server classpath so that the handler is installed during server startup. The log records from the custom handler will have the format described in “Log Records” on page 259.
- **Log Filter** — To filter log records that are sent to destinations such as `server.log`, `syslog`, or a destination specified by a custom log handler, you can plug in a custom log filter. The custom filter must implement the interface `java.util.logging.Filter`. Type the absolute class name of the filter in the Log Filter field. Also put the filter class in the Application Server classpath so that the filter is installed during server startup.
- **File Rotation Limit**: When the server log reaches the specified size in bytes, the logs are rotated and renamed to `server.log_date`, where *date* is the date and time when the file was rotated. The default value is 2 megabytes. The minimum value for the limit is 500 kilobytes; if you specify a lower value, the file rotates when it reaches 500 KB. To turn off log file rotation, set the value to 0.
- **File Rotation Time Limit** — Rotate the server log after the specified number of minutes is reached. The default value is zero, which means that the file is rotated when it reaches the size specified in the File Rotation Limit field. *If you specify one or more minutes, the time limit takes precedence over the size limit.*

- 5 Click **Save** to save changes. Click **View Log Files** to view the server log.

▼ To configure log levels

- 1 In the tree component, expand the **Node Agents** or **Configurations** node.
- 2 Select a node agent, or select the instance to configure:
 - To configure a particular instance, select the instance’s **config** node. For example, for the default instance, `server`, select the `server-config` node.
 - To configure the default settings for future instances that use a copy of `default-config`, select the `default-config` node.
- 3 For a node agent, select the **Log Levels** tab. For a configuration, select the **Logger Settings** node, then select the **Log Levels** tab.
- 4 On the **Module Log Levels** page, choose a new value from the drop-down list opposite the module or modules whose log level is to be changed.

The default level is `INFO`, meaning that messages at that level or higher (`WARNING`, `SEVERE`) appear in the log. Choose any of the following values (listed from highest to lowest):

- SEVERE
- WARNING
- INFO
- CONFIG
- FINE
- FINER
- FINEST
- OFF

5 Use the Additional Properties area to configure log levels for any application loggers.

The property name is the logger namespace, and the value is one of the eight possible levels. For example, the property name could be `samples.logging.simple.servlet`, and the value could be `FINE`.

Also use this area to change the log level for a submodule, such as the `transport` submodule of the CORBA module:

```
javax.enterprise.resource.corba.ORBId.transport
```

6 Click Save to save the changes, or click Default to restore the default values.

Calls to `System.out.println` are logged at the `INFO` level using the logger name `javax.enterprise.system.stream.out`. Calls to `System.err.println` are logged at the `WARNING` level using the logger name `javax.enterprise.system.stream.err`. To turn off the logs from these sources, specify the logger name with the value `OFF` in the Additional Properties area.

Changes to the Log Level settings take effect immediately. They are also saved in the `domain.xml` file for use when the server restarts.

▼ To view the server log

- 1 In the tree component, expand the node for the server instance whose log you want to view.**
- 2 On the General Information page, click View Log Files.**
- 3 Use the Search Criteria area to customize and filter the log viewer. Use the basic fields as follows:**
 - **Instance Name** — Choose an instance name from the drop-down list to view the log for that server instance. The default is the current server instance.
 - **Log File** — Choose a log file name from the drop-down list to view the contents of that log. The default is `server.log`.

- **Timestamp** — To view the most recent messages, select Most Recent (the default). To view messages only from a certain period of time, select Specific Range and type a date and time value in the From and To fields that appear. For the Time value, the syntax must take the following form (SSS stands for milliseconds):

hh:mm:ss.SSS

For example:

17:10:00.000

If the From value is later than the To value, an error message appears.

- **Log Level** — To filter messages by log level, choose a log level from the drop-down list. By default, the display includes all messages that appear in the server log at the chosen log level and more severe levels. Select the checkbox labeled “Do not include more severe messages” to display messages at only the chosen level.

To ensure that the messages you want to view appear in the server log, first set the appropriate log levels on the Log Levels page. See [“To configure log levels” on page 263](#)

If you choose to filter log messages based on log level, only messages matching the specified filter criteria are shown. However, this filtering does not affect which messages are logged to the server log.

The most recent 40 entries in the server log appear, with the settings specified on the Logging Settings and Log Levels pages.

Click the triangle next to the Timestamp header to sort the messages so that the most recent one appears last.

To view a formatted version of any message, click the link marked

(details)

A window labeled Log Entry Detail appears, with a formatted version of the message.

At the end of the list of entries, click the buttons to view earlier or later entries in the log file.

Click Advanced Search in the Search Criteria area to make additional refinements to the log viewer. Use the Advanced Options fields as follows:

- **Logger** — To filter by module, choose one or more namespaces from the drop-down list. Use shift-click or control-click to choose multiple namespaces.

Selecting a namespace at a higher level selects all the namespaces below it. For example, selecting `javax.enterprise.system` also selects the loggers for all the modules under that namespace: `javax.enterprise.system.core`, `javax.enterprise.system.tools.admin`, and so on.
- **Custom Logger** — To view messages from loggers specific to a particular application, type the logger names in the text field, one per line. If the application has several modules, you can view any or all of them. For example, suppose the application has loggers with the following names:

```
com.mycompany.myapp.module1  
com.mycompany.myapp.module2  
com.mycompany.myapp.module3
```

To view messages from all modules in the application, type `com.mycompany.myapp`. To view messages from `module2` only, type `com.mycompany.myapp.module2`.

When you specify one or more custom loggers, messages from Application Server modules appear only if you specify them explicitly in the Logger area.

- **Name-Value Pairs** — To view output from a specific thread, type the key name and value for that thread in the text field. The key name is `_ThreadID`. For example:

```
_ThreadID=13
```

Suppose that `com.mycompany.myapp.module2` runs in several threads. To refine the log viewer to show only the output from a single thread, specify that module's logger in the Custom Logger field, and then specify the thread ID in this field.

- **Display** — To view more than 40 messages at a time (the default), choose another of the available values from the drop-down list (100, 250, or 1000).

To view stack traces, deselect the “Limit excessively long messages” checkbox. By default, stack traces do not appear in the viewer; to view them, click the `(details)` link for a message.

Click Basic Search to hide the Advanced Options area.

Monitoring Components and Services

This chapter contains information about monitoring components using the Application Server Admin Console. This chapter contains the following sections:

- “About Monitoring” on page 267
- “Admin Console Tasks for Enabling and Disabling Monitoring” on page 294
- “Admin Console Tasks for Viewing Monitoring Data” on page 296
- “Using JConsole” on page 314

About Monitoring

- “Monitoring in the Application Server” on page 267
- “Overview of Monitoring” on page 267
- “About the Tree Structure of Monitorable Objects” on page 268
- “About Statistics for Monitored Components and Services” on page 271

Monitoring in the Application Server

Use monitoring to observe the runtime state of various components and services deployed in a server instance of the Application Server. With the information on the state of runtime components and processes, it is possible to identify performance bottlenecks for tuning purposes, aid capacity planning, predict failures, do root cause analysis in case of failures, and ensure that everything is functioning as expected.

Turning monitoring on reduces performance by increasing overhead.

Overview of Monitoring

To monitor the Application Server, perform these steps:

1. Enable the monitoring of specific services and components using either the Admin Console or the `asadmin` tool.

For more information on this step, refer to [“Admin Console Tasks for Enabling and Disabling Monitoring”](#) on page 294.

2. View monitoring data for the specified services or components using either the Admin Console or the `asadmin` tool.

For more information on this step, refer to [“Admin Console Tasks for Viewing Monitoring Data”](#) on page 296.

About the Tree Structure of Monitorable Objects

The Application Server uses a tree structure to track monitorable objects. Because the tree of monitoring objects is dynamic, it changes as components are added, updated, or removed in the instance. The root object in the tree is the server instance name, for example, `server`. (In the Platform Edition, just one server instance is permitted.)

The following command displays the top level of the tree:

```
asadmin> list --user adminuser --monitor server
server.applications
server.http-service
server.connector-service
server.jms-service
server.jvm
server.orb
server.resources
server.thread-pools
```

The following sections describe these sub-trees:

- [“The Applications Tree”](#) on page 268
- [“The HTTP Service Tree”](#) on page 269
- [“The Resources Tree”](#) on page 270
- [“The Connector Service Tree”](#) on page 270
- [“The JMS Service Tree”](#) on page 270
- [“The ORB Tree”](#) on page 271
- [“The Thread Pool Tree”](#) on page 271

The Applications Tree

The following schematic shows the top and child nodes for the various components of enterprise applications. The nodes at which monitoring statistics are available are marked with an asterisk (*). For more information, refer to [“EJB Container Statistics”](#) on page 272.

EXAMPLE 16-1 Applications Node Tree Structure

```

applications
|--- application1
|   |--- ejb-module-1
|   |   |--- ejb1 *
|   |       |--- cache (for entity/sfsb) *
|   |       |--- pool (for slsb/mdb/entity) *
|   |       |--- methods
|   |           |---method1 *
|   |           |---method2 *
|   |           |--- stateful-session-store (for sfsb)*
|   |           |--- timers (for slsb/entity/mdb) *
|   |--- web-module-1
|   |   |--- virtual-server-1 *
|   |       |---servlet1 *
|   |       |---servlet2 *
|--- standalone-web-module-1
|   |   |----- virtual-server-2 *
|   |       |---servlet3 *
|   |       |---servlet4 *
|   |   |----- virtual-server-3 *
|   |       |---servlet3 *(same servlet on different vs)
|   |       |---servlet5 *
|--- standalone-ejb-module-1
|   |   |--- ejb2 *
|   |       |--- cache (for entity/sfsb) *
|   |       |--- pool (for slsb/mdb/entity) *
|   |       |--- methods
|   |           |--- method1 *
|   |           |--- method2 *
|--- application2

```

The HTTP Service Tree

The nodes of the HTTP service are shown in the following schematic. The nodes at which monitoring information is available are marked with an asterisk (*). See [“HTTP Service Statistics” on page 277](#).

EXAMPLE 16-2 HTTP Service Schematic (Platform Edition version)

```

http-service
|--- virtual-server-1
|   |--- http-listener-1 *
|   |--- http-listener-2 *
|--- virtual-server-2
|   |--- http-listener-1 *
|   |--- http-listener-2 *

```

EXAMPLE 16-3 HTTP Service Schematic (Enterprise Edition version)

```

http-service *
  |-- connection-queue *
  |-- dns *
  |-- file-cache *
  |-- keep-alive *
  |-- pwc-thread-pool *
  |-- virtual-server-1*
  |         |-- request *
  |-- virtual-server-2*
  |         |-- request *

```

The Resources Tree

The resources node holds monitorable attributes for pools such as the JDBC connection pool and connector connection pool. The following schematic shows the top and child nodes for the various resource components. The nodes at which monitoring statistics are available are marked with an asterisk (*). See [“JDBC Connection Pools Statistics” on page 278](#).

EXAMPLE 16-4 Resources Schematic

```

resources
  |--connection-pool1(either connector-connection-pool or jdbc)*
  |--connection-pool2(either connector-connection-pool or jdbc)*

```

The Connector Service Tree

The connector services node holds monitorable attributes for pools such as the connector connection pool. The following schematic shows the top and child nodes for the various connector service components. The nodes at which monitoring statistics are available are marked with an asterisk (*). See [“JMS/Connector Service Statistics” on page 280](#).

EXAMPLE 16-5 Connector Service Schematic

```

connector-service
  |-- resource-adapter-1
  |         |-- connection-pools
  |         |         |-- pool-1 (All pool stats for this pool)
  |         |-- work-management (All work mgmt stats for this RA)

```

The JMS Service Tree

The JMS services node holds monitorable attributes for pools such as the connector connection pool. The following schematic shows the top and child nodes for the various JMS service components. The nodes at which monitoring statistics are available are marked with an asterisk (*).

EXAMPLE 16-6 JMS Service Schematic

```
jms-service
  |-- connection-factories [AKA conn. pools in the RA world]
  |      |-- connection-factory-1 (All CF stats for this CF)
  |-- work-management (All work mgmt stats for the MQ-RA)
```

The ORB Tree

The ORB node holds monitorable attributes for connection managers. The following schematic shows the top and child nodes for the ORB components. The nodes at which monitoring statistics are available are marked with an asterisk (*). See [“Statistics for Connection Managers in an ORB” on page 281](#).

EXAMPLE 16-7 ORB Schematic

```
orb
  |-- connection-managers
  |      |-- connection-manager-1 *
  |      |-- connection-manager-1 *
```

The Thread Pool Tree

The thread pool node holds monitorable attributes for connection managers. The following schematic shows the top and child nodes for the ORB components. The nodes at which monitoring statistics are available are marked with an asterisk (*). See [“Thread Pools Statistics” on page 282](#).

EXAMPLE 16-8 Thread Pool Schematic

```
thread-pools
  |      |-- thread-pool-1 *
  |      |-- thread-pool-2 *
```

About Statistics for Monitored Components and Services

This section describes the monitoring statistics that are available:

- [“EJB Container Statistics” on page 272](#)
- [“Web Container Statistics” on page 276](#)
- [“HTTP Service Statistics” on page 277](#)
- [“JDBC Connection Pools Statistics” on page 278](#)
- [“JMS/Connector Service Statistics” on page 280](#)
- [“Statistics for Connection Managers in an ORB” on page 281](#)

- “Thread Pools Statistics” on page 282
- “Transaction Service Statistics” on page 282
- “Java Virtual Machine (JVM) Statistics” on page 283
- “JVM Statistics in J2SE 5.0” on page 283
- “Production Web Container (PWC) Statistics” on page 288

EJB Container Statistics

EJB statistics are described in the following table.

TABLE 16-1 EJB Statistics

Attribute Name	Data Type	Description
createcount	CountStatistic	Number of times an EJB's create method is called.
removecount	CountStatistic	Number of times an EJB's remove method is called.
pooledcount	RangeStatistic	Number of entity beans in pooled state.
readycount	RangeStatistic	Number of entity beans in ready state.
messagecount	CountStatistic	Number of messages received for a message-driven bean.
methodreadycount	RangeStatistic	Number of stateful or stateless session beans that are in the MethodReady state.
passivecount	RangeStatistic	Number of stateful session beans that are in Passive state.

The statistics available for EJB method invocations are listed in the following table.

TABLE 16-2 EJB Method Statistics

Attribute Name	Data Type	Description
methodstatistic	TimeStatistic	Number of times an operation is called; the total time that is spent during the invocation, and so on.

TABLE 16-2 EJB Method Statistics (Continued)

Attribute Name	Data Type	Description
totalnumerrors	CountStatistic	Number of times the method execution resulted in an exception. This is collected for stateless and stateful session beans and entity beans if monitoring is enabled for the EJB container.
totalnumsuccess	CountStatistic	Number of times the method successfully executed. This is collected for stateless and stateful session beans and entity beans if monitoring enabled is true for EJB container.
executiontime	CountStatistic	Time (ms) spent executing the method for the last successful/unsuccessful attempt to execute the operation. This is collected for stateless and stateful session beans and entity beans if monitoring is enabled on the EJB container.

The statistics for EJB Session Stores are listed in the following table.

TABLE 16-3 EJB Session Store Statistics

Attribute Name	Data Type	Description
currentSize	RangeStatistic	Number of passivated or checkpointed sessions currently in the store.
activationCount	CountStatistic	Number of sessions activated from the store.
activationSuccessCount	CountStatistic	Number of sessions successfully activated from the store
activationErrorCount	CountStatistic	Time (ms) spent executing the method for the last successful/unsuccessful attempt to execute the operation. This is collected for stateless and stateful session beans and entity beans, if monitoring is enabled on the EJB container.

TABLE 16-3 EJB Session Store Statistics (Continued)

Attribute Name	Data Type	Description
passivationCount	CountStatistic	Number of sessions passivated (inactivated) using this store.
passivationSuccessCount	CountStatistic	Number of sessions successfully passivated using this store.
passivationErrorCount	CountStatistic	Number of sessions that could not be passivated using this store.
expiredSessionCount	CountStatistic	Number of expired sessions that were removed by this store.
passivatedBeanSize	CountStatistic	Total number of bytes passivated by this store, including total, minimum, and maximum.
passivationTime	CountStatistic	Time spent on passivating beans to the store, including the total, minimum, and maximum.
checkpointCount (EE only)	CountStatistic	Number of sessions checkpointed using this store.
checkpointSuccessCount (EE only)	CountStatistic	Number of sessions checkpointed successfully.
checkpointErrorCount (EE only)	CountStatistic	Number of sessions that couldn't be checkpointed.
checkpointedBeanSize (EE only)	ValueStatistic	Total number of beans checkpointed by the store.
checkpointTime (EE only)	TimeStatistic	Time spent on checkpointing beans to the store.

The statistics available for EJB pools are listed in the following table.

TABLE 16-4 EJB Pool Statistics

Attribute Name	Data Type	Description
numbeansinpool	BoundedRangeStatistic	Number of EJB's in the associated pool, providing an idea about how the pool is changing.
numthreadswaiting	BoundedRangeStatistic	Number of threads waiting for free beans, giving an indication of possible congestion of requests.

TABLE 16-4 EJB Pool Statistics (Continued)

Attribute Name	Data Type	Description
totalbeanscreated	CountStatistic	Number of beans created in associated pool since the gathering of data started.
totalbeansdestroyed	CountStatistic	Number of beans destroyed from associated pool since the gathering of data started.
jmsmaxmessagesload	CountStatistic	The maximum number of messages to load into a JMS session at one time for a message-driven bean to serve. Default is 1. Applies only to pools for message driven beans.

The statistics available for EJB caches are listed in the following table.

TABLE 16-5 EJB Cache Statistics

Attribute Name	Data Type	Description
cachemisses	BoundedRangeStatistic	The number of times a user request does not find a bean in the cache.
cachehits	BoundedRangeStatistic	The number of times a user request found an entry in the cache.
numbeansincache	BoundedRangeStatistic	The number of beans in the cache. This is the current size of the cache.
numpassivations	CountStatistic	Number of passivated beans. Applies only to stateful session beans.
numpassivationerrors	CountStatistic	Number of errors during passivation. Applies only to stateful session beans.
numexpiredsessionsremoved	CountStatistic	Number of expired sessions removed by the cleanup thread. Applies only to stateful session beans.
numpassivationsuccess	CountStatistic	Number of times passivation completed successfully. Applies only to stateful session beans.

The statistics available for Timers are listed in the following table.

TABLE 16-6 Timer Statistics

Statistic	Data Type	Description
numtimerscreated	CountStatistic	Number of timers created in the system.
numtimersdelivered	CountStatistic	Number of timers delivered by the system.
numtimersremoved	CountStatistic	Number of timers removed from the system.

Web Container Statistics

The web container fits into the tree of objects as shown in [“The Applications Tree” on page 268](#). Web container statistics are displayed for each individual web application. Statistics available for the web container for Servlets are shown in [“Web Container Statistics” on page 276](#) and statistics available for web modules are shown in [“Web Container Statistics” on page 276](#).

TABLE 16-7 Web Container (Servlet) Statistics

Statistic	Units	Data Type	Comments
errorcount	Number	CountStatistic	Cumulative number of cases where the response code is greater than or equal to 400.
maxtime	Milliseconds	CountStatistic	The maximum amount of time the web container waits for requests.
processingtime	Milliseconds	CountStatistic	Cumulative value of the amount of time required to process each request. The processing time is the average of request processing times divided by the request count.
requestcount	Number	CountStatistic	The total number of requests processed so far.

Statistics available for web modules are shown in [“Web Container Statistics” on page 276](#).

TABLE 16-8 Web Container (Web Module) Statistics

Statistic	Data Type	Comments
jspcount	CountStatistic	Number of JSP pages that have been loaded in the web module.
jspreloadcount	CountStatistic	Number of JSP pages that have been reloaded in the web module.

TABLE 16-8 Web Container (Web Module) Statistics (Continued)

Statistic	Data Type	Comments
sessionstotal	CountStatistic	Total number of sessions that have been created for the web module.
activesessionscurrent	CountStatistic	Number of currently active sessions for the web module.
activesessionshigh	CountStatistic	Maximum number of concurrently active sessions for the web module.
rejectedsessionstotal	CountStatistic	Total number of rejected sessions for the web module. This is the number of sessions that were not created because the maximum allowed number of sessions were active.
expiredsessionstotal	CountStatistic	Total number of expired sessions for the web module.
sessionsize (EE only)	AverageRangeStatistic	Size of the session for the web module. Value is either high, low, or average, or is in bytes for serialized sessions.
containerlatency (EE only)	AverageRangeStatistic	Latency for the web container's part of the overall latency request. Value is either high, low, or average.
sessionpersisttime (EE only)	AverageRangeStatistic	Time (in ms, low, high, or average) taken to persist HTTP session state to back-end store for the web module.
cachedsessionscurrent (EE only)	CountStatistic	Current number of sessions cached in memory for the web module.
passivatedsessionscurrent (EE only)	CountStatistic	Current number of sessions passivated for the web module.

HTTP Service Statistics

The statistics available for the HTTP service are shown in “[HTTP Service Statistics](#)” on [page 277](#). These statistics are applicable to the Platform Edition only. For statistics for the HTTP Service on the Enterprise Edition, see “[Production Web Container \(PWC\) Statistics](#)” on [page 288](#).

TABLE 16-9 HTTP Service Statistics (applicable to Platform Edition only)

Statistic	Units	Data Type	Comments
bytesreceived	Bytes	CountStatistic	The cumulative value of the bytes received by each of the request processors.
bytessent	Bytes	CountStatistic	The cumulative value of the bytes sent by each of the request processors.
currentthreadcount	Number	CountStatistic	The number of processing threads currently in the listener thread pool.
currentthreadsbusy	Number	CountStatistic	The number of request processing threads currently in use in the listener thread pool serving requests.
errorcount	Number	CountStatistic	The cumulative value of the error count, which represents the number of cases where the response code is greater than or equal to 400.
maxsparethreads	Number	CountStatistic	The maximum number of unused response processing threads that can exist.
minsparethreads	Number	CountStatistic	The minimum number of unused response processing threads that can exist.
maxthreads	Number	CountStatistic	The maximum number of request processing threads created by the listener.
maxtime	Milliseconds	CountStatistic	The maximum amount of time for processing threads.
processing-time	Milliseconds	CountStatistic	The cumulative value of the times taken to process each request. The processing time is the average of request processing times divided by the request count.
request-count	Number	CountStatistic	The total number of requests processed so far.

JDBC Connection Pools Statistics

Monitor JDBC resources to measure performance and capture resource usage at runtime. As the creation of JDBC connections are expensive and frequently cause performance bottlenecks in applications, it is crucial to monitor how a JDBC connection pool is releasing and creating new connections and how many threads are waiting to retrieve a connection from a particular pool.

The statistics available for the JDBC connection pool are shown in the following table.

TABLE 16-10 JDBC Connection Pool Statistics

Statistic	Units	Data Type	Description
numconnfailedvalidation	Number	CountStatistic	The total number of connections in the connection pool that failed validation from the start time until the last sample time.
numconnused	Number	RangeStatistic	Provides connection usage statistics. The total number of connections that are currently being used, as well as information about the maximum number of connections that were used (the high water mark).
numconnfree	Number	RangeStatistic	The total number of free connections in the pool as of the last sampling.
numconntimedout	Number	BoundedRangeStatistic	The total number of connections in the pool that timed out between the start time and the last sample time.
averageconnwaittime	Number	CountStatistic	Indicates the average wait time of connections for successful connection request attempts to the connector connection pool.
waitqueuelength	Number	CountStatistic	Number of connection requests in the queue waiting to be serviced.
connectionrequestwaittime		RangeStatistic	The longest and shortest wait times of connection requests. The current value indicates the wait time of the last request that was serviced by the pool.
numconncreated	Milliseconds	CountStatistic	The number of physical connections that were created since the last reset.
numconndestroyed	Number	CountStatistic	Number of physical connections that were destroyed since the last reset.

TABLE 16-10 JDBC Connection Pool Statistics (Continued)

Statistic	Units	Data Type	Description
numconnacquired	Number	CountStatistic	Number of logical connections acquired from the pool.
numconnreleased	Number	CountStatistic	Number of logical connections released to the pool.

JMS/Connector Service Statistics

The statistics available for the connector connection pools are shown in “[JMS/Connector Service Statistics](#)” on page 280. Statistics for Connector Work Management are shown in “[JMS/Connector Service Statistics](#)” on page 280.

TABLE 16-11 Connector Connection Pool Statistics

Statistic	Units	Data Type	Description
numconnfailedvalidation	Number	CountStatistic	The total number of connections in the connection pool that failed validation from the start time until the last sample time.
numconnused	Number	RangeStatistic	Provides connection usage statistics. The total number of connections that are currently being used, as well as information about the maximum number of connections that were used (the high water mark).
numconnfree	Number	RangeStatistic	The total number of free connections in the pool as of the last sampling.
numconntimedout	Number	CountStatistic	The total number of connections in the pool that timed out between the start time and the last sample time.
averageconnwaittime	Number	CountStatistic	Average wait time of connections before they are serviced by the connection pool.
waitqueuelenght	Number	CountStatistic	Number of connection requests in the queue waiting to be serviced.
connectionrequestwaittime		RangeStatistic	The longest and shortest wait times of connection requests. The current value indicates the wait time of the last request that was serviced by the pool.
numconncreated	Milliseconds	CountStatistic	The number of physical connections that were created since the last reset.

TABLE 16-11 Connector Connection Pool Statistics (Continued)

Statistic	Units	Data Type	Description
numconndestroyed	Number	CountStatistic	Number of physical connections that were destroyed since the last reset.
numconnacquired	Number	CountStatistic	Number of logical connections acquired from the pool.
numconnreleased	Number	CountStatistic	Number of logical connections released to the pool.

Statistics available for Connector Work Management are listed in [“JMS/Connector Service Statistics” on page 280](#).

TABLE 16-12 Connector Work Management Statistics

Statistic	Data Type	Description
activeworkcount	RangeStatistic	Number of work objects executed by the connector.
waitqueuelength	RangeStatistic	Number of work objects waiting in the queue before executing.
workrequestwaittime	RangeStatistic	Longest and shortest wait of a work object before it gets executed.
submittedworkcount	CountStatistic	Number of work objects submitted by a connector module.
rejectedworkcount	CountStatistic	Number of work objects rejected by the Application Server.
completedworkcount	CountStatistic	Number of work objects that were completed.

Statistics for Connection Managers in an ORB

The statistics available for the connection manager in an ORB are listed in [“Statistics for Connection Managers in an ORB” on page 281](#).

TABLE 16-13 Connection Manager (in an ORB) Statistics

Statistic	Units	Data Type	Description
connectionsidle	Number	CountStatistic	Provides total number of connections that are idle to the ORB.

TABLE 16-13 Connection Manager (in an ORB) Statistics (Continued)

Statistic	Units	Data Type	Description
connectionsinuse	Number	CountStatistic	Provides total number of connections in use to the ORB.
totalconnections	Number	BoundedRangeStatistic	Total number of connections to the ORB.

Thread Pools Statistics

The statistics available for the thread pool are shown in the following table.

TABLE 16-14 Thread Pool Statistics

Statistic	Units	Data Type	Description
averagetimeinqueue	Milliseconds	RangeStatistic	The average amount of time in milliseconds a request waited in the queue before getting processed.
averageworkcompletion-time	Milliseconds	RangeStatistic	The average amount of time taken to complete an assignment, in milliseconds.
currentnumberofthreads	Number	BoundedRangeStatistic	Current number of request processing threads.
numberofavailablethreads	Number	CountStatistic	The number of threads that are available.
numberofbusythreads	Number	CountStatistic	The number of threads that are busy.
totalworkitemsadded	Number	CountStatistic	The total number of work items added so far to the work queue.

Transaction Service Statistics

The transaction service allows the client to freeze the transaction subsystem in order to roll back transactions and determine the transactions that are in process at the time of the freeze. The statistics available for the transaction service are shown in the following table.

TABLE 16-15 Transaction Service Statistics

Statistic	Data Type	Description
activecount	CountStatistic	Number of transactions currently active.

TABLE 16-15 Transaction Service Statistics (Continued)

Statistic	Data Type	Description
activeids	StringStatistic	The ID's of the transactions that are currently active. Every such transaction can be rolled back after freezing the transaction service.
committedcount	CountStatistic	Number of transactions that have been committed.
rolledbackcount	CountStatistic	Number of transactions that have been rolled back.
state	StringStatistic	Indicates whether or not the transaction has been frozen.

Java Virtual Machine (JVM) Statistics

The JVM has monitorable attributes that are always enabled. The statistics available for the JVM are shown in the following table.

TABLE 16-16 JVM Statistics

Statistic	Data Type	Description
heapsize	BoundedRangeStatistic	The resident memory footprint with the higher and lower bounds of the JVM's memory heap size.
uptime	CountStatistic	The amount of time the JVM has been running.

JVM Statistics in J2SE 5.0

If the Application Server is configured to run on J2SE version 5.0 or higher, additional monitoring information can be obtained from the JVM. Set the monitoring level to LOW to enable the display of this additional information. Set the monitoring level to HIGH to also view information pertaining to each live thread in the system. More information on the additional monitoring features available in J2SE 5.0 is available in a document titled *Monitoring and Management for the Java Platform*, which is available from <http://java.sun.com/j2se/1.5.0/docs/guide/management/>.

The J2SE 5.0 monitoring tools are discussed at <http://java.sun.com/j2se/1.5.0/docs/tooldocs/#manage>.

The statistics available for class loading in the JVM in J2SE 5.0 are shown in “JVM Statistics in J2SE 5.0” on page 283.

TABLE 16-17 JVM Statistics for J2SE 5.0 - Class Loading

Statistic	Data Type	Description
loadedclasscount	CountStatistic	Number of classes that are currently loaded in the JVM.
totalloadedclasscount	CountStatistic	Total number of classes that have been loaded since the JVM began execution.
unloadedclasscount	CountStatistic	Number of classes that have been unloaded from the JVM since the JVM began execution.

The statistics available for compilation in the JVM in J2SE 5.0 are shown in [“JVM Statistics in J2SE 5.0” on page 283](#).

TABLE 16-18 JVM Statistics for J2SE 5.0 - Compilation

Statistic	Data Type	Description
totalcompilationtime	CountStatistic	Accumulated time (in milliseconds) spent in compilation.

The statistics available for garbage collection in the JVM in J2SE 5.0 are shown in [“JVM Statistics in J2SE 5.0” on page 283](#).

TABLE 16-19 JVM Statistics for J2SE 5.0 - Garbage Collection

Statistic	Data Type	Description
collectioncount	CountStatistic	Total number of collections that have occurred.
collectiontime	CountStatistic	Accumulated collection time (in milliseconds).

The statistics available for memory in the JVM in J2SE 5.0 are shown in [“JVM Statistics in J2SE 5.0” on page 283](#).

TABLE 16-20 JVM Statistics for J2SE 5.0 - Memory

Statistic	Data Type	Description
objectpendingfinalizationcount	CountStatistic	Approximate number of objects that are pending finalization.

TABLE 16–20 JVM Statistics for J2SE 5.0 - Memory (Continued)

Statistic	Data Type	Description
initheapsize	CountStatistic	Size of the heap initially requested by the JVM.
usedheapsize	CountStatistic	Size of the heap currently in use.
maxheapsize	CountStatistic	Maximum amount of memory (in bytes) that can be used for memory management.
committedheapsize	CountStatistic	Amount of memory (in bytes) that is committed for the JVM to use.
initnonheapsize	CountStatistic	Size of the non-heap area initially requested by the JVM.
usednonheapsize	CountStatistic	Size of the non-heap area currently in use.
maxnonheapsize	CountStatistic	Maximum amount of memory (in bytes) that can be used for memory management.
committednonheapsize	CountStatistic	Amount of memory (in bytes) that is committed for the JVM to use.

The statistics available for the operating system in the JVM in J2SE 5.0 are shown in [“JVM Statistics in J2SE 5.0” on page 283](#).

TABLE 16–21 JVM Statistics for J2SE 5.0 - Operating System

Statistic	Data Type	Description
arch	StringStatistic	Operating system architecture.
availableprocessors	CountStatistic	Number of processors available to the JVM.
name	StringStatistic	Operating system name.
version	StringStatistic	Operating system version.

The statistics available for the runtime in the JVM in J2SE 5.0 are shown in [“JVM Statistics in J2SE 5.0” on page 283](#).

TABLE 16–22 JVM Statistics for J2SE 5.0 - Runtime

Statistic	Data Type	Description
name	StringStatistic	Name representing the running JVM
vmname	StringStatistic	JVM implementation name.

TABLE 16–22 JVM Statistics for J2SE 5.0 - Runtime (Continued)

Statistic	Data Type	Description
vmvendor	StringStatistic	JVM implementation vendor.
vmversion	StringStatistic	JVM implementation version.
specname	StringStatistic	JVM specification name.
specvendor	StringStatistic	JVM specification vendor.
specversion	StringStatistic	JVM specification version.
managementspecversion	StringStatistic	Management spec. version implemented by the JVM.
classpath	StringStatistic	Classpath that is used by the system class loader to search for class files.
librarypath	StringStatistic	Java library path.
bootclasspath	StringStatistic	Classpath that is used by the bootstrap class loader to search for class files.
inputarguments	StringStatistic	Input arguments passed to the JVM. Does not include the arguments to the main method.
uptime	CountStatistic	Uptime of the JVM (in milliseconds).

The statistics available for ThreadInfo in the JVM in J2SE 5.0 are shown in [“JVM Statistics in J2SE 5.0” on page 283](#).

TABLE 16–23 JVM Statistics for J2SE 5.0 - Thread Info

Statistic	Data Type	Description
threadid	CountStatistic	Id of the thread.
threadname	StringStatistic	Name of the thread.
threadstate	StringStatistic	State of the thread.
blockedtime	CountStatistic	Time elapsed (in milliseconds) since the thread entered the BLOCKED state. Returns -1 if thread contention monitoring is disabled.
blockedcount	CountStatistic	Total number of times that the thread entered the BLOCKED state.

TABLE 16–23 JVM Statistics for J2SE 5.0 - Thread Info (Continued)

Statistic	Data Type	Description
waitedtime	CountStatistic	Elapsed time (in milliseconds) that the thread has been in a WAITING state. Returns -1 if thread contention monitoring is disabled.
waitedcount	CountStatistic	Total number of times the thread was in WAITING or TIMED_WAITING states.
lockname	StringStatistic	String representation of the monitor lock that the thread is blocked to enter or waiting to be notified through the <code>Object.wait</code> method.
lockownerid	CountStatistic	Id of the thread that holds the monitor lock of an object on which this thread is blocking.
lockownername	StringStatistic	Name of the thread that holds the monitor lock of the object this thread is blocking on.
stacktrace	StringStatistic	Stack trace associated with this thread.

The statistics available for threads in the JVM in J2SE 5.0 are shown in “[JVM Statistics in J2SE 5.0](#)” on page 283.

TABLE 16–24 JVM Statistics for J2SE 5.0 - Threads

Statistic	Data Type	Description
threadcount	CountStatistic	Current number of live daemon and non-daemon threads.
peakthreadcount	CountStatistic	Peak live thread count since the JVM started or the peak was reset.
totalstartedthreadcount	CountStatistic	Total number of threads created and/or started since the JVM started.
daemonthreadcount	CountStatistic	Current number of live daemon threads.
allthreadids	StringStatistic	List of all live thread ids.
currentthreadputime	CountStatistic	CPU time for the current thread (in nanoseconds) if CPU time measurement is enabled. If CPU time measurement is disabled, returns -1.

TABLE 16–24 JVM Statistics for J2SE 5.0 - Threads (Continued)

Statistic	Data Type	Description
monitordeadlockedthreads	StringStatistic	List of thread ids that are monitor deadlocked.

Production Web Container (PWC) Statistics

Statistics are available for the following PWC components and services on the Enterprise Edition (EE) of the Application Server:

- “Production Web Container (PWC) Statistics” on page 288, PWC Virtual Server
- “Production Web Container (PWC) Statistics” on page 288, PWC Request
- “Production Web Container (PWC) Statistics” on page 288, PWC File Cache
- “Production Web Container (PWC) Statistics” on page 288, PWC Keep Alive
- “Production Web Container (PWC) Statistics” on page 288, PWC DNS
- “Admin Console Tasks for Listeners and JMX Connectors” on page 183, PWC Thread Pool
- “Production Web Container (PWC) Statistics” on page 288, PWC Connection Queue
- “Production Web Container (PWC) Statistics” on page 288, PWC HTTP Service

Statistics for PWC virtual servers are listed in “Production Web Container (PWC) Statistics” on page 288.

TABLE 16–25 PWC Virtual Server Statistics (EE only)

Attribute Name	Data Type	Description
id	StringStatistic	The ID of the virtual server.
mode	StringStatistic	The mode the virtual server is in. Options include unknown or active.
hosts	StringStatistic	Name of the hosts serviced by this virtual server.
interfaces	StringStatistic	Type of interfaces (listeners) for which the virtual server is configured.

The statistics available for PWC requests are listed in the following table.

TABLE 16–26 PWC Request Statistics (EE only)

Attribute Name	Data Type	Description
method	StringStatistic	Method used for request.
uri	StringStatistic	Last URI served.
countrequests	CountStatistic	Number of requests served.

TABLE 16-26 PWC Request Statistics (EE only) (Continued)

Attribute Name	Data Type	Description
countbytestransmitted	CountStatistic	Number of bytes transmitted, or 0 if this information is not available
countbytesreceived	CountStatistic	Number of bytes received, or 0 if this information is not available.
ratebytesreceived	CountStatistic	Rate at which data was received over some server-defined interval, or 0 if this information is not available
maxbytestransmissionrate	CountStatistic	Maximum rate at which data was transmitted over some server-defined interval, or 0 if this information is not available.
countopenconnections	CountStatistic	Number of connections that are currently open, or 0 if this information is not available.
maxopenconnections	CountStatistic	Maximum number of concurrently open connections, or 0 if this information is not available.
count2xx	CountStatistic	Total number of responses of code 2XX.
count3xx	CountStatistic	Total number of responses of code 3XX.
count4xx	CountStatistic	Total number of responses of code 4XX.
count5xx	CountStatistic	Total number of responses of code 5XX.
countother	CountStatistic	Total number of responses with other response codes.
count200	CountStatistic	Total number of responses of code 200.
count302	CountStatistic	Total number of responses of code 302.
count304	CountStatistic	Total number of responses of code 304.
count400	CountStatistic	Total number of responses of code 400.

TABLE 16-26 PWC Request Statistics (EE only) (Continued)

Attribute Name	Data Type	Description
count401	CountStatistic	Total number of responses of code 401.
count403	CountStatistic	Total number of responses of code 403.
count404	CountStatistic	Total number of responses of code 404.
count503	CountStatistic	Total number of responses of code 503.

The cache information section provides information on how the file cache is being used. Statistics for PWC file caches are listed in the following table.

TABLE 16-27 PWC File Cache Statistics (EE only)

Attribute Name	Data Type	Description
flagenabled	CountStatistic	Indicates whether the file cache is enabled. Valid values are 0 for no or 1 for yes.
secondsmaxage	CountStatistic	Maximum age of a valid cache entry, in seconds.
countentries	CountStatistic	Number of current cache entries. A single cache entry represents a single URI.
maxentries	CountStatistic	Maximum number of simultaneous cache entries.
countopenentries	CountStatistic	Number of entries associated with an open file.
maxopenentries	CountStatistic	Maximum number of simultaneous cache entries associated with an open file.
sizeheapcache	CountStatistic	Heap space used for cache content.
maxheapcachesize	CountStatistic	Maximum heap space used for cache file content.
sizemmapcache	CountStatistic	Amount of address space used by memory mapped file content.
maxmmapcachesize	CountStatistic	Maximum amount of address space used by the file cache for memory mapped file content.
counthits	CountStatistic	Number of successful cache lookups.
countmisses	CountStatistic	Number of failed cache lookups.
countinfohits	CountStatistic	Number of times a file information lookup succeeded.

TABLE 16-27 PWC File Cache Statistics (EE only) (Continued)

Attribute Name	Data Type	Description
countinfomisses	CountStatistic	Number of misses on cached file information.
countcontenthits	CountStatistic	Number of hits on cached file content.
countcontentmisses	CountStatistic	Number of times a file information lookup failed.

This section provides information about the server's HTTP-level keep-alive system. The statistics available for PWC Keep Alive are listed in the following table.

TABLE 16-28 PWC Keep Alive Statistics (EE only)

Attribute Name	Data Type	Description
countconnections	CountStatistic	Number of connections in keep-alive mode.
maxconnections	CountStatistic	Maximum number of connections simultaneously allowed in keep-alive mode.
counthits	CountStatistic	The total number of times connections in keep-alive mode have subsequently made a valid request.
countflushes	CountStatistic	The number of times keep-alive connections have been closed by the server.
countrefusals	CountStatistic	The number of times the server could not hand off the connection to a keep-alive thread, possibly due to too many persistent connections.
counttimeouts	CountStatistic	The number of times the server terminated keep-alive connections as the client connections timed out without any activity.
secondstimeout	CountStatistic	The time (in seconds) before idle keep-alive connections are closed.

The DNS Cache caches IP addresses and DNS names. The server's DNS cache is disabled by default. A single cache entry represents a singular IP address or DNS name lookup. The statistics available for PWC DNS are listed in the following table.

TABLE 16-29 PWC DNS Statistics (EE only)

Attribute Name	Data Type	Description
flagcacheenabled	CountStatistic	Indicates whether the DNS cache is enabled (on). Either 0 for off or 1 for on.

TABLE 16–29 PWC DNS Statistics (EE only) (Continued)

Attribute Name	Data Type	Description
countcacheentries	CountStatistic	Number of DNS entries presently in the cache.
maxcacheentries	CountStatistic	Maximum number of DNS entries that can be accommodated by the cache.
countcachehits	CountStatistic	Number of times a DNS cache lookup has succeeded.
countcachemisses	CountStatistic	Number of times a DNS cache lookup has failed.
flagasyncenabled	CountStatistic	Indicates whether asynchronous DNS lookups are enabled (on). Either 0 for off or 1 for on.
countasyncnamelookups	CountStatistic	Total number of asynchronous DNS name lookups.
countasynccaddrlookups	CountStatistic	Total number of asynchronous DNS address lookups.
countasynclookupsinprogress	CountStatistic	Number of asynchronous lookups in progress.

Statistics for PWC thread pools are listed in the following table.

TABLE 16–30 PWC Thread Pool Statistics (EE only)

Attribute Name	Data Type	Description
id	StringStatistic	ID of the thread pool.
countthreadside	CountStatistic	Number of request-processing threads currently idle.
countthreads	CountStatistic	Current number of request-processing threads.
maxthreads	CountStatistic	Maximum number of request processing threads that can exist concurrently.
countqueued	CountStatistic	Number of requests queued for processing by this thread pool.
peakqueued	CountStatistic	The largest number of requests in the queue simultaneously.
maxqueued	CountStatistic	Maximum number of requests that can be in the queue at one time.

The Connection Queue is the queue in which requests are held prior to being serviced. Statistics for the connection queue show the number of sessions in the queue and the average delay before the connection is accepted. Statistics for PWC connection queues are listed in the following table.

TABLE 16-31 PWC Connection Queue Statistics (EE only)

Attribute Name	Data Type	Description
id	StringStatistic	ID of the connection queue.
counttotalconnections	CountStatistic	Total number of connections that have been accepted.
countqueued	CountStatistic	Number of connections currently in the queue.
peakqueued	CountStatistic	Largest number of connections that were in the queue simultaneously.
maxqueued	CountStatistic	Maximum size of the connection queue.
countoverflows	CountStatistic	The number of times the queue has been too full to accommodate a connection.
counttotalqueued	CountStatistic	The total number of connections that have been queued. A given connection may be queued multiple times, so counttotalqueued may be greater than or equal to counttotalconnections.
tickstotalqueued	CountStatistic	The total number of ticks that connections have spent in the queue. A tick is a system-dependent unit of time.
countqueued1minuteaverage	CountStatistic	Average number of connections queued in the last 1 minute.
countqueued5minuteaverage	CountStatistic	Average number of connections queued in the last 5 minutes.
countqueued15minuteaverage	CountStatistic	Average number of connections queued in the last 15 minutes.

Statistics for PWC HTTP service are listed in the following table.

TABLE 16-32 PWC HTTP Service Statistics (EE only)

Attribute Name	Data Type	Description
id	StringStatistic	Instance name of the HTTP service.

TABLE 16-32 PWC HTTP Service Statistics (EE only) *(Continued)*

Attribute Name	Data Type	Description
versionserver	StringStatistic	Version number of the HTTP service.
timestarted	StringStatistic	Time the HTTP service was started (GMT).
secondsrunning	CountStatistic	Time (in seconds) since the HTTP service started.
maxthreads	CountStatistic	Maximum number of worker threads in each instance.
maxvirtualservers	CountStatistic	Maximum number of virtual servers that can be configured in each instance.
flagprofilingenabled	CountStatistic	Whether or not HTTP service performance profiling is enabled. Valid values are 0 or 1.
flagvirtualserveroverflow	CountStatistic	Indicates whether more than maxvirtualservers are configured. If this is set to 1, statistics are not being tracked for all virtual servers.
load1minuteaverage	CountStatistic	Average load for requests in the last 1 minute.
load5minuteaverage	CountStatistic	Average load for requests in the last 5 minutes.
load15minuteaverage	CountStatistic	Average load for requests in the last 15 minutes.
ratebytestransmitted	CountStatistic	The rate at which data is transmitted over some server-defined interval. The result is 0 when this information is not available.
ratebytesreceived	CountStatistic	The rate at which data is received over some server-defined interval. The result is 0 when this information is not available.

Admin Console Tasks for Enabling and Disabling Monitoring

- [“To configure monitoring levels using the Admin Console” on page 294](#)
- [“To configure monitoring using the asadmin tool” on page 296](#)

▼ To configure monitoring levels using the Admin Console

- 1 Access the Monitoring Service page. To do this,
 - a. In the tree component, expand the Configurations node.

- b. In the tree, expand the node for the server instance to be configured for monitoring, for example, `server-config`.
 - c. In the tree, select **Monitoring**.
- 2 On the **Monitoring Service** page, choose the appropriate value from the combo box opposite the component(s) or service(s) whose monitoring level is changing.
- By default, monitoring is turned off for all components and services. To turn monitoring on, select **LOW** or **HIGH** from the combo box. To turn monitoring off, select **OFF** from the combo box. It is possible to turn monitoring on or off for the following components and services:
- **JVM** - Set the monitoring level to **LOW** for this option to monitor the Java Virtual Machine.
 - **HTTP Service** - Set the monitoring level to **LOW** for this option to monitor all HTTP listeners and virtual servers.
 - **Transaction Service** - Set the monitoring level to **LOW** for this option to monitor any transaction subsystem.
 - **JMS/Connector Service** - Set the monitoring level to **LOW** for this option to monitor any Java Message Service (JMS).
 - **ORB** - Set the monitoring level to **LOW** for this option to monitor the system ORB used by the Application Server core and its connection managers.
 - **Web Container** - Set the monitoring level to **LOW** for this option to monitor all deployed servlets.
 - **EJB Container** - Set the monitoring level to **LOW** for this option to monitor all deployed EJB components, EJB pools, and EJB caches. Set this method to **HIGH** to also monitor EJB business methods.
 - **JDBC Connection Pool** - Set the monitoring level to **LOW** for this option to monitor all JDBC connection pools.
 - **Thread Pool** - Set the monitoring level to **LOW** for this option to monitor all thread pools.

3 **Click Save.**

There are no Additional Monitoring Service Properties in this release, therefore ignore the Additional Properties table.

More Information Equivalent `asadmin` command

`set`

For example, to turn on monitoring for the HTTP Service, use the following `asadmin` command:

```
asadmin> set --user admin-user
server.monitoring-service.module-monitoring-levels.http-service=LOW
```

▼ To configure monitoring using the asadmin tool

To turn monitoring off, or to set a level for monitoring a component or service, you can use the Admin Console as described in [“To configure monitoring levels using the Admin Console” on page 294](#).

- 1 Use the `get` command to find out what services and components currently have monitoring enabled.

```
asadmin> get --user admin-user server.monitoring-service.module-monitoring-levels.*
```

Returns:

```
server.monitoring-service.module-monitoring-levels.  
connector-connection-pool = OFF  
server.monitoring-service.module-monitoring-levels.  
connector-service = OFF  
server.monitoring-service.module-monitoring-levels.ejb-container = OFF  
server.monitoring-service.module-monitoring-levels.http-service = OFF  
server.monitoring-service.module-monitoring-levels.jdbc-connection-pool = OFF  
server.monitoring-service.module-monitoring-levels.jms-service = OFF  
server.monitoring-service.module-monitoring-levels.jvm = OFF  
server.monitoring-service.module-monitoring-levels.orb = OFF  
server.monitoring-service.module-monitoring-levels.thread-pool = OFF  
server.monitoring-service.module-monitoring-levels.transaction-service = OFF  
server.monitoring-service.module-monitoring-levels.web-container = OFF
```

- 2 Use the `set` command to enable monitoring.

For example, to enable monitoring for the HTTP service:

```
asadmin> set --user admin-user  
server.monitoring-service.module-monitoring-levels.http-service=LOW
```

To disable monitoring, use the `set` command and specify `OFF` for the monitoring level.

Admin Console Tasks for Viewing Monitoring Data

- [“To view monitoring data in the Admin Console” on page 296](#)
- [“Viewing Monitoring Data With the asadmin Tool” on page 299](#)

▼ To view monitoring data in the Admin Console

For more information on the attributes for each component or service, refer to [“About Statistics for Monitored Components and Services” on page 271](#).

- 1 **Access the Monitoring page. To do this:**
 - a. In the tree component, expand the Standalone Instances node.
 - b. Select a stand-alone server instance, for example server (Admin Server), from the list.
 - c. Select the Monitor page.
 - d. Select the Monitor tab on the Monitor page.

- 2 **From the View list, select a component or service that has been deployed onto the server instance and for which monitoring is enabled.**

Monitoring data for the selected component or service displays below the View field. For a description of the monitorable properties, refer to [“About Statistics for Monitored Components and Services”](#) on page 271.

On this page, it is possible to view monitoring data for JVM, Server, Thread Pools, HTTP Service, and Transaction Service if monitoring is enabled for these components and services. A diagram showing how these components and services are organized is shown in [“About the Tree Structure of Monitorable Objects”](#) on page 268.

- 3 **If you do not see the component or service you wish to monitor in this list, select the Configure Monitoring link to enable and disable monitoring for selected components and services.**

Select OFF to disable monitoring for a component or service. Select LOW or HIGH to enable monitoring for a component or service.

For more information on enabling and disabling monitoring, refer to [“To configure monitoring levels using the Admin Console”](#) on page 294 or [“To configure monitoring using the asadmin tool”](#) on page 296.

- 4 **Select the *Applications* tab of the Monitor page to view monitoring data for application components that are deployed onto the server instance and for which monitoring is enabled. Select the application from the Application list. Select the specific component from the Component list.**

If no monitoring data appears for the application or component, select the Configure Monitoring link to enable or disable monitoring for a component or service. To monitor applications, turn on the container in which they execute. For example, select Low or High for the Web Container for web applications and/or the EJB Container for EJB applications.

If no monitoring data is displayed for applications, it is most likely that the application does not exist or is not exercised. Applications monitoring data is available only when the application exists, monitoring is enabled for the application, and the application is being exercised. Once the application is executed, it is registered in the monitoring registry and the monitoring data displays.

Use the Admin Console to monitor remote applications and instances. For this to occur, the remote instance must be running and the configuration set.

Monitoring data for the selected component displays below the selected component. For a description of the monitorable properties, refer to [“About Statistics for Monitored Components and Services” on page 271](#). A diagram showing how these components and services are organized for applications can be viewed in [“About the Tree Structure of Monitorable Objects” on page 268](#).

- 5 Select the *Resources* page to view monitoring data for resources that are deployed onto the server instance and for which monitoring has been enabled. Select the resource from the View list.**

If the resource for which you wish to view monitoring data does not appear, select the [Configure Monitoring](#) link to enable or disable monitoring for a resource.

If no monitoring data is displayed for resources, it is most likely that the resource does not exist or is not exercised. Resources monitoring data is available only when the resources exist, monitoring is enabled for the resource at a level of HIGH, and the resource is being exercised. For example, if you have created a JDBC connector service, but applications that use that connector service have not yet requested a connector from the service, that service has not yet been created, therefore, no service exists and no monitoring data is available. Once the JDBC application is executed and requests a connector from a service, the service is registered in the monitoring registry and monitoring data appears.

Monitoring data for the selected component or service displays below the View field. For a description of the monitorable properties, refer to [“About Statistics for Monitored Components and Services” on page 271](#). A diagram showing how these components and services are organized for resources can be viewed in [“About the Tree Structure of Monitorable Objects” on page 268](#).

- 6 Select the *Transactions* page to freeze the transaction subsystem in order to roll back transactions and determine the transactions that are in process at the time of the freeze.**
- 7 To enable monitoring for the Transaction Service, select the [Configure Monitoring](#) link and make sure that Transaction Service is set to LOW.**

To freeze the Transaction Service in order to roll back transactions, select Freeze. To rollback a transaction, select the checkbox beside the transaction and click Rollback.

More Information Equivalent `asadmin` command

```
get --monitor
```

For example, to view monitoring data for the JVM, use the following `asadmin` command:

```
asadmin> get --user adminuser --monitor server.jvm.*
```

Viewing Monitoring Data With the `asadmin` Tool

- “To use the `asadmin` tool to view monitoring data” on page 299
- “Understanding and Specifying Dotted Names” on page 300
- “Examples of the `list` and `get` Commands” on page 301
- “Examples for the `list --user admin-user --monitor` Command” on page 301
- “Examples for the `get --user admin-user --monitor` Command” on page 302
- “To use the PetStore example” on page 304
- “Expected Output for `list` and `get` Commands at All Levels” on page 307

▼ To use the `asadmin` tool to view monitoring data

To view monitoring data using the `asadmin` tool, use the `asadmin list` and `asadmin get` commands followed by the dotted name of a monitorable object, as follows.

1 To view the names of the objects that can be monitored, use the `asadmin list` command.

For example, to view a list of application components and subsystems that have monitoring enabled for the server instance, type the following command in a terminal window:

```
asadmin> list --user adminuser --monitor server
```

The preceding command returns a list of application components and subsystems that have monitoring enabled, for example:

```
server.resources
server.connector-service
server.orb
server.jms-service
server.jvm
server.applications
server.http-service
server.thread-pools
```

For further examples using the `list` command, refer to [“Examples of the `list` and `get` Commands” on page 301](#). For further information on the dotted names you can use with the `list` command, refer to [“Understanding and Specifying Dotted Names” on page 300](#).

2 To display monitoring statistics for an application component or subsystem for which monitoring has been enabled, use the `asadmin get` command.

To get the statistics, type the `asadmin get` command in a terminal window, specifying a name displayed by the `list` command in the preceding step. The following example attempts to get all attributes from a subsystem for a specific object:

```
asadmin> get --user adminuser --monitor server.jvm.*
```

The command returns the following attributes and data:

```
server.jvm.dotted-name = server.jvm
server.jvm.heapsize-current = 21241856
server.jvm.heapsize-description = Provides statistical information about
    the JVM's memory heap size.
server.jvm.heapsize-highwatermark = 21241856
server.jvm.heapsize-lastsampletime = 1080232913938
server.jvm.heapsize-lowerbound = 0
server.jvm.heapsize-lowwatermark = 0
server.jvm.heapsize-name = JvmHeapSize
server.jvm.heapsize-starttime = 1080234457308
server.jvm.heapsize-unit = bytes
server.jvm.heapsize-upperbound = 518979584
server.jvm.uptime-count = 1080234457308
server.jvm.uptime-description = Provides the amount of time the JVM has
    been running.
server.jvm.uptime-lastsampletime = 1080234457308
server.jvm.uptime-name = JvmUpTime
server.jvm.uptime-starttime = 1080232913928
server.jvm.uptime-unit = milliseconds
```

For further examples using the `get` command, refer to [“Examples of the list and get Commands” on page 301](#). For further information on the dotted names you can use with the `get` command, refer to [“Understanding and Specifying Dotted Names” on page 300](#).

Understanding and Specifying Dotted Names

In the `admin list` and `get` commands, specify the dotted name of monitorable objects. All child objects are addressed using the dot (`.`) character as separator, thus these are referred to as *dotted names*. If a child node is of singleton type, then only the monitoring object type is needed to address the object, otherwise a name of the form `.name` is needed to address the object.

For example, `http-service` is one of the valid monitorable object types and is a singleton. To address a singleton child node representing the `http-service` of instance `server`, the dotted name is:

```
server.http-service
```

Another example, `application`, is a valid monitorable object type and is not a singleton. To address a non-singleton child node representing, for example, the application `PetStore`, the dotted name is:

```
server.applications.petstore
```

The dotted names can also address specific attributes in monitorable objects. For example, `http-service` has a monitorable attribute called `bytes received-lastsampletime`. The following name addresses the `bytes received` attribute:

```
server.http-service.server.http-listener-1.
bytesreceived-laststampletime
```

The administrator is not expected to know the valid dotted names for `asadmin list` and `get` commands. The `list` command displays available monitorable objects, while the `get` command used with a wildcard parameter allows the inspection of all available attributes on any monitorable object.

The underlying assumptions for using the `list` and `get` commands with dotted names are:

- Any `list` command that has a dotted name that is **not** followed by a wildcard (*) gets as its result the current node's immediate children. For example, `list --user adminuser --monitor server` lists all immediate children belonging to the `server` node.
- Any `list` command that has a dotted name followed by a wildcard of the form `.*` gets as its result a hierarchical tree of children nodes from the current node. For example, `list --user adminuser --monitor server.applications.*` lists all children of `applications` and their subsequent child nodes and so on.
- Any `list` command that has a dotted name preceded or followed by a wildcard of the form `*dottedname` or `dotted * name` or **dotted name *** gets as its result all nodes and their children matching the regular expression created by the provided matching pattern.
- A `get` command followed by a `.*` or a `*` gets as its result the set of attributes and their values belonging to the current node to be matched.

For more information, read [“Expected Output for list and get Commands at All Levels”](#) on page 307.

Examples of the list and get Commands

This section contains the following topics:

- [“Examples for the list --user admin-user --monitor Command”](#) on page 301
- [“Examples for the get --user admin-user --monitor Command”](#) on page 302

Examples for the list --user admin-user --monitor Command

The `list` command provides information about the application components and subsystems currently being monitored for the specified server instance name. Using this command, you can see the monitorable components and subcomponents for a server instance. For a more complete listing of `list` examples, see [“Expected Output for list and get Commands at All Levels”](#) on page 307.

Example 1

```
asadmin> list --user admin-user --monitor server
```

The preceding command returns a list of application components and subsystems that have monitoring enabled, for example:

```
server.resources
server.orb
server.jvm
server.jms-service
server.connector-service
server.applications
server.http-service
server.thread-pools
```

It is also possible to list applications that are currently monitored in the specified server instance. This is useful when particular monitoring statistics are sought from an application using the `get` command.

Example 2

```
asadmin> list --user admin-user --monitor server.applications
```

Returns:

```
server.applications.adminapp
  server.applications.admingui
server.applications.myApp
```

For a more comprehensive example, see [“To use the PetStore example” on page 304](#).

Examples for the `get --user admin-user --monitor` Command

This command retrieves the following monitored information:

- All attribute(s) monitored within a component or subsystem
- Specific attribute monitored within a component or subsystem

When an attribute is requested that does not exist for a particular component or subsystem, an error is returned. Similarly, when a specific attribute is requested that is not active for a component or subsystem, an error is returned.

Refer to [“Expected Output for list and get Commands at All Levels” on page 307](#) for more information on the use of the `get` command.

Example 1

Attempt to get all attributes from a subsystem for a specific object:

```
asadmin> get --user admin-user --monitor server.jvm.*
```

Returns:

```
server.jvm.dotted-name= server.jvm
server.jvm.heapsize-current = 21241856
server.jvm.heapsize-description = Provides statistical information about
```

```

    the JVM's memory heap size.
server.jvm.heapsize-highwatermark = 21241856
server.jvm.heapsize-lastsampletime = 1080232913938
server.jvm.heapsize-lowerbound = 0
server.jvm.heapsize-lowwatermark = 0
server.jvm.heapsize-name = JvmHeapSize
server.jvm.heapsize-starttime = 1080234457308
server.jvm.heapsize-unit = bytes
server.jvm.heapsize-upperbound = 518979584
server.jvm.uptime-count = 1080234457308
server.jvm.uptime-description = Provides the amount of time the JVM has
    been running.
server.jvm.uptime-lastsampletime = 1080234457308
server.jvm.uptime-name = JvmUpTime
server.jvm.uptime-starttime = 1080232913928
server.jvm.uptime-unit = milliseconds

```

Example 2

Attempt to get all attributes from a J2EE application:

```
asadmin> get --user admin-user --monitor server.applications.myJ2eeApp.*
```

Returns:

```

No matches resulted from the wildcard expression.
CLI137 Command get failed.

```

There are no monitorable attributes exposed at the J2EE-application level, therefore this reply displays.

Example 3

Attempt to get a specific attribute from a subsystem:

```
asadmin> get --user admin-user --monitor server.jvm.uptime-lastsampletime
```

Returns:

```
server.jvm.uptime-lastsampletime = 1093215374813
```

Example 4

Attempt to get an unknown attribute from within a subsystem attribute:

```
asadmin> get --user admin-user --monitor server.jvm.badname
```

Returns:

```
No such attribute found from reflecting the corresponding Stats
interface: [badname]
CLI137 Command get failed.
```

▼ To use the PetStore example

The following example illustrates how the `asadmin` tool might be used for monitoring purposes.

A user wants to inspect the number of calls made to a method in the sample `PetStore` application after it has been deployed onto the Application Server. The instance onto which it has been deployed is named `server`. A combination of the `list` and `get` commands are used to access desired statistics on a method.

1 Start the Application Server and the `asadmin` tool.

2 Set some useful environment variables to avoid entering them for every command:

```
asadmin> export AS_ADMIN_USER=admin AS_ADMIN_PASSWORD=admin123
asadmin>export AS_ADMIN_HOST=localhost AS_ADMIN_PORT=4849
```

3 List monitorable components for instance `server`:

```
asadmin> list --user adminuser --monitor server*
```

Returns (output will be similar to):

```
server
server.applications
server.applications.CometEJB
server.applications.ConverterApp
server.applications.petstore
server.http-service
server.resources
server.thread-pools
```

The list of monitorable components includes `thread-pools`, `http-service`, `resources`, and all deployed (and enabled) applications.

4 List the monitorable subcomponents in the `PetStore` application (`-m` can be used instead of `--monitor`):

```
asadmin> list -m server.applications.petstore
```

Returns:

```
server.applications.petstore.signon-ejb_jar
server.applications.petstore.catalog-ejb_jar
server.applications.petstore.uidgen-ejb_jar
server.applications.petstore.customer-ejb_jar
server.applications.petstore.petstore-ejb_jar
server.applications.petstore.petstore\war
```



```
server.applications.petstore.AsyncSenderJAR_jar
server.applications.petstore.cart-ejb_jar
```

5 List the monitorable subcomponents in the EJB module `signon-ejb_jar` of the PetStore application:

```
asadmin> list -m server.applications.petstore.signon-ejb_jar
```

Returns:

```
server.applications.petstore.signon-ejb_jar.SignOnEJB
server.applications.petstore.signon-ejb_jar.UserEJB
```

6 List the monitorable subcomponents in the entity bean `UserEJB` for the EJB module `signon-ejb_jar` of the PetStore application:

```
asadmin> list -m server.applications.petstore.signon-ejb_jar.UserEJB
```

Returns (with dotted name removed for space considerations):

```
server.applications.petstore.signon-ejb_jar.UserEJB.bean-cache
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods
server.applications.petstore.signon-ejb_jar.UserEJB.bean-pool
```

7 List the monitorable subcomponents in the method `getUserName` for the entity bean `UserEJB` in the EJB module `signon-ejb_jar` of the PetStore application:

```
asadmin> list -m
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.getUserName
```

Returns:

```
Nothing to list at server.applications.petstore.signon-ejb_jar.
UserEJB.bean-methods.getUserName. To get the valid names beginning with a
string, use the wildcard "*" character. For example, to list all names
that begin with "server", use "list server*".
```

8 There are no monitorable subcomponents for methods. Get all monitorable statistics for the method `getUserName`.

```
asadmin> get -m
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.getUserName.*
```

Returns:

```
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-count = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-description = Provides the time in milliseconds
spent during the last successful/unsuccessful attempt to execute the
operation.
```

```
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-lastsampletime = 1079981809259
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-name = ExecutionTime
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-starttime = 1079980593137
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-unit = count
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-count = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-description = Provides the number of times an
    operation was called, the total time that was spent during the
    invocation and so on.
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-lastsampletime = 1079980593137
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-maxtime = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-mintime = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-name = ExecutionTime
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-starttime = 1079980593137
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-totaltime = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.methodstatistic-unit =
    server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumerrors-count = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumerrors-description = Provides the total number of errors
    that occurred during invocation or execution of an operation.
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumerrors-lastsampletime = 1079981809273
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumerrors-name = TotalNumErrors
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumerrors-starttime = 1079980593137
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumerrors-unit = count
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumsuccess-count = 0
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumsuccess-description = Provides the total number of
    successful invocations of the method.
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumsuccess-lastsampletime = 1079981809255
```

```
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumsuccess-name = TotalNumSuccess
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumsuccess-starttime = 1079980593137
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.totalnumsuccess-unit = count
```

9 To also get a specific statistic, such as execution time, use a command such as the following:

```
asadmin> get -m server.applications.petstore.signon-ejb_jar.
UserEJB.bean-methods.getUserName.executiontime-count
```

Returns:

```
server.applications.petstore.signon-ejb_jar.UserEJB.bean-methods.
getUserName.executiontime-count = 1
```

Expected Output for list and get Commands at All Levels

The following tables show the command, dotted name, and corresponding output at each level of the tree.

TABLE 16-33 Top Level

Command	Dotted Name	Output
list -m	server	server.applicationsserver.thread-poolserver.resourcesserver.http-serviceserver.transaction-serviceserver.orb.connection-managersserver.orb.connection-managers.orb\Connections\Inbound\AcceptedConnectionsserver.jvm
list -m	server.*	Hierarchy of child nodes below this node.
get -m	server.*	No output except a message saying there are no attributes at this node.

The following table shows the command, dotted name, and corresponding output for the applications level.

TABLE 16-34 Applications Level

Command	Dotted Name	Output
list -m	server.applications or *applications	appl1app2web-module1_warejb-module2_jar...

TABLE 16-34 Applications Level (Continued)

Command	Dotted Name	Output
list -m	server.applications.* or *applications.*	Hierarchy of child nodes below this node.
get -m	server.applications.* or *applications.*	No output except message saying there are no attributes at this node.

The following table shows the command, dotted name, and corresponding output for stand-alone modules and enterprise applications at the applications level.

TABLE 16-35 Applications - Enterprise Applications and Standalone Modules

Command	Dotted Name	Output
list -m	server.applications.app1 or *app1 Note: this level is only applicable if an enterprise application has been deployed. It is not applicable if a standalone module is deployed.	ejb-module1_jarweb-module2_war ejb-module3_jarweb-module3_war...
list -m	server.applications.app1.* or *app1.*	Hierarchy of child nodes below this node.
get -m	server.applications.app1.* or *app1.*	No output except message saying there are no attributes at this node.
list -m	server.applications.app1.ejb-module1_jar or *ejb-module1_jar or server.applications.ejb-module1_jar	bean1bean2bean3...

TABLE 16-35 Applications - Enterprise Applications and Standalone Modules (Continued)

Command	Dotted Name	Output
list -m	server.applications.app1.ejb-module1_jar or *ejb-module1_jar or server.applications.ejb-module1_jar	Hierarchy of child nodes below this node.
get -m	server.applications.app1.ejb-module1_jar.* or *ejb-module1_jar.* or server.applications.ejb-module1_jar.*	No output except message saying there are no attributes at this node.
list -m	server.applications.app1. ejb-module1_jar.bean1 Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	List of child nodes: bean-poolbean-cachebean-method
list -m	server.applications.app1. ejb-module1_jar.bean1 Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	Hierarchy of child nodes and a list of all attributes for this node and for any subsequent child nodes.

TABLE 16-35 Applications - Enterprise Applications and Standalone Modules (Continued)

Command	Dotted Name	Output
get -m	server.applications.app1. ejb-module1_jar.bean1.* Note: In standalone modules, the node containing the application name (app1 in this example) does not appear.	The following attributes and their associated values: CreateCount_CountCreateCount_ DescriptionCreateCount_ LastSampleTimeCreateCount_ NameCreateCount_ StartTimeCreateCount_ UnitMethodReadyCount_ CurrentMethodReadyCount_ DescriptionMethodReadyCount_ HighWaterMarkMethodReadyCount_ LastSampleTimeMethodReadyCount_ LowWaterMarkMethodReadyCount_ NameMethodReadyCount_ StartTimeMethodReadyCount_ UnitRemoveCount_CountRemoveCount_ DescriptionRemoveCount_ LastSampleTimeRemoveCount_ NameRemoveCount_StartTimeAttribute RemoveCount_Unit
list -m	server.applications.app1. ejb-module1_jar.bean1.bean-pool Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	No attributes, but a message saying: Nothing to list at server.applications.app1. ejb-module1_jar.bean1-cache. To get the valid names beginning with a string, use the wildcard (*) character. For example, to list all names that begin with server, use list server*.
get -m	server.applications.app1. ejb-module1_jar.bean1.bean-pool.* Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	List of attributes and values corresponding to EJB Pool attributes as described in Table 1-4.
list -m	server.applications.app1. ejb-module1_jar.bean1.bean-cache Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”
get -m	server.applications.app1. ejb-module1_jar.bean1.bean-cache.* Note: In standalone modules, the node containing the application name (app1 in this example) does not appear.	List of attributes and values corresponding to EJB Cache attributes as described in Table 1-5.

TABLE 16-35 Applications - Enterprise Applications and Standalone Modules (Continued)

Command	Dotted Name	Output
list -m	server.applications.app1. ejb-module1_jar.bean1.bean-method.method1 Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	No attributes, but a message saying "Use get command with the --monitor option to view this node's attributes and values."
get -m	server.applications.app1. ejb-module1_jar.bean1.bean-method.method1.* Note: In standalone modules, the node containing the application name (app1 in this example) will not appear.	List of attributes and values corresponding to EJB Methods attributes as described in Table 1-2.
list -m	server.applications.app1.web-module1_war	Displays the virtual server(s) assigned to the module.
get -m	server.applications.app1.web-module1_war.*	No output except a message saying there are no attributes at this node.
list -m	server.applications.app1.web-module1_war. virtual_server	Displays list of servlets registered.
get -m	server.applications.app1.web-module1_war. virtual_server.*	No output except a message saying there are no attributes at this node.
list -m	server.applications.app1.web-module1_war. virtual_server.servlet1	No attributes, but a message saying "Use get command with the --monitor option to view this node's attributes and values."
get -m	server.applications.app1.web-module1_war. virtual_server.servlet1.*	List of attributes and values corresponding to web container (Servlet) attributes as described in Table 1-7.

The following table shows the command, dotted name, and corresponding output for the HTTP Service level.

TABLE 16-36 HTTP-Service Level

Command	Dotted Name	Output
list -m	server.http-service	List of virtual servers.
get -m	server.http-service.*	No output except message saying there are no attributes at this node.
list -m	server.http-service.server	List of HTTP Listeners.
get -m	server.http-service.server.*	No output except message saying there are no attributes at this node.

TABLE 16-36 HTTP-Service Level (Continued)

Command	Dotted Name	Output
list -m	server.http-service.server. http-listener1	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”
get -m	server.http-service.server.*	List of attributes and values corresponding to HTTP Service attributes as described in Table 1-9.

The following table shows the command, dotted name, and corresponding output for the thread pools level.

TABLE 16-37 Thread-Pools Level

Command	Dotted Name	Output
list -m	server.thread-pools	List of thread-pool names.
get -m	server.thread-pools.*	No output except message saying there are no attributes at this node.
list -m	server.thread-pools.orb\ .threadpool\.thread-pool-1	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”
get -m	server.thread-pools.orb\ .threadpool\.thread-pool-1.*	List of attributes and values corresponding to Thread Pool attributes as described in Table 1-14.

The following table shows the command, dotted name, and corresponding output for the resources level.

TABLE 16-38 Resources Level

Command	Dotted Name	Output
list -m	server.resources	List of pool names.
get -m	server.resources.*	No output except message saying there are no attributes at this node.

TABLE 16-38 Resources Level (Continued)

Command	Dotted Name	Output
list -m	server.resources.jdbc-connection-pool-pool.connection-pool1	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”
get -m	server.resources.jdbc-connection-pool-pool.connection-pool1.*	List of attributes and values corresponding to Connection Pool attributes as described in Table 1-10.

The following table shows the command, dotted name, and corresponding output for the transaction service level.

TABLE 16-39 Transaction-Service Level

Command	Dotted Name	Output
list -m	server.transaction-service	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”
get -m	server.transaction-service.*	List of attributes and values corresponding to Transaction Service attributes as described in Table 1-15.

The following table shows the command, dotted name, and corresponding output for the ORB level.

TABLE 16-40 ORB Level

Command	Dotted Name	Output
list -m	server.orb	server-orb.connection-managers
get -m	server.orb.*	No output except message saying there are no attributes at this node.
list -m	server.orb.connection-managers	Name(s) of ORB connection managers.
get -m	server.orb.connection-managers.*	No output except message saying there are no attributes at this node.
list -m	server.orb.connection-managers.orb\Connections\Inbound\AcceptedConnections	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”

TABLE 16-40 ORB Level (Continued)

Command	Dotted Name	Output
get -m	server.orb.connection-managers. orb\.Connections\.Inbound\ .AcceptedConnections.*	List of attributes and values corresponding to ORB Connection Manager attributes as described in Table 1-13.

The following table shows the command, dotted name, and corresponding output for the JVM level.

TABLE 16-41 JVM Level

Command	Dotted Name	Output
list -m	server.jvm	No attributes, but a message saying “Use get command with the --monitor option to view this node’s attributes and values.”
get -m	server.jvm.*	List of attributes and values corresponding to JVM attributes as described in Table 1-16.

Using JConsole

For JConsole to work with the Application Server, security has to be disabled for the JMX connector. The current version of the Application Server (SE/EE edition) has security enabled by default.

▼ To disable security for the JMX Connector

- 1 Use the Admin Console to disable security for the JMX Connector. To do this from the Admin Console,
 - a. Expand the Node Agent.
 - b. Select the Node Agent.
 - c. Deselect SSL3 and TLS in JMX tab.
 - d. Select Save.

2 Use `asadmin` to disable security for the JMX Connector. To do this from a terminal window or command prompt,

a. Enter this command:

```
asadmin set server.admin-service.jmx-connector.system.security-enabled=false
```

b. Restart the domain application server (DAS).

For the Platform Edition version, the JMX Connector is disabled by default, therefore there is no need to change any configuration for Platform Edition.

3 Start JConsole and enter the JMX URL, user name, and password on the Advanced tab for logging in. The JMX URL is of the form:

```
service:jmx:rmi:///jndi/rmi://<your machine  
name>:<port>/management/rmi-jmx-connector
```

NOTE: You can get the exact JMX URL from the admin server . log file if you search for message ADM1501.

Java Virtual Machine and Advanced Settings

This chapter explains how to configure the Java Virtual Machine (JVM™) and other advanced settings. It contains the following sections:

- “Admin Console Tasks for JVM Settings” on page 317
- “Admin Console Tasks for Advanced Settings” on page 321

Admin Console Tasks for JVM Settings

- “To configure the JVM general settings” on page 317
- “To configure the JVM classpath settings” on page 318
- “To configure the JVM options” on page 319
- “To disable the security manager” on page 320
- “To configure the JVM profiler settings” on page 321

▼ To configure the JVM general settings

The Java Virtual Machine (JVM) is included in the Java 2 Standard Edition (J2SE™) software, which is required by the Application Server. Because incorrect JVM settings will prevent the server from running, you should take care when changing these settings.

- 1 In tree component, select **Configurations**.
- 2 Click the **JVM Setting** node from the target server instance.
- 3 By default, the **General** link located below the tabs is already selected.

4 On the JVM General Settings page, you may specify the following:

- a. In the Java Home field, enter the name of the installation directory of the Java 2 Standard Edition (J2SE) software.**

The Application Server relies on the J2SE software. To verify that the J2SE version you specify is supported in this release, refer to the Release Notes. (See the link in the [“Further Information” on page 55](#) section.)

Note – If you enter a nonexistent directory name or the installation directory name of an unsupported version of the J2SE software, then the Application Server will not start.

- b. In the Javac field, type the command-line options for the Java programming language compiler.**

The Application Server runs the compiler when EJB components are deployed.

- c. To set up debugging with the JPDA (Java Platform Debugger Architecture), you select the Debug Enabled checkbox and specify options in the Debug Options field.**

JPDA is used by application developers. For more information, see the Debugging J2EE Applications chapter of the Application Server Developer’s Guide. (For a link to the guide, see [“Further Information” on page 55](#).)

- d. In the RMI Compile Options field, type the command-line options for the `rmic` compiler.**

The Application Server runs the `rmic` compiler when EJB components are deployed.

- e. In the Bytecode Preprocessor field, type a comma separated list of class names.**

Each class must implement the `com.sun.appserv.BytecodePreprocessor` interface. The classes are called in the order specified.

Tools such as profilers might require entries in the Bytecode Preprocessor field. Profilers generate information used to analyze server performance. For more information about profiling, see the Debugging J2EE Applications chapter of the *Application Server Developer’s Guide*.

5 Click Save.**6 Restart the server.****▼ To configure the JVM classpath settings**

The classpath is the list of JAR files that the Java runtime environment searches for classes and other resource files.

- 1 In tree component, select the Application Server node.
- 2 Click the JVM Settings tab.
- 3 Select the Path Settings link below the tabs.
- 4 On the JVM Classpath Settings page, you may specify the following:
 - a. In the Environment Classpath checkbox, retain the default selection to ignore the CLASSPATH environment variable.

The CLASSPATH environment variable is convenient for basic tutorials in programming, but is not recommended for enterprise environments.
 - b. To view the Application Server's classpath, examine the read-only contents of the Server Classpath field.
 - c. To insert a JAR file into the beginning of the server's classpath, enter the full path name of the file in the Classpath Prefix field.
 - d. To add a JAR file to the end of the server's classpath, enter the full path name of the file in the Classpath Suffix field.

For example, you would specify the JAR file of a database driver. See [“To integrate a JDBC driver” on page 83](#).
 - e. In the Native Library Path Prefix and Suffix fields, you may prepend or append entries to the native library path.

The native library path is a concatenation of the server's relative path for its native shared libraries, the standard JRE native library path, the shell environment setting (LD_LIBRARY_PATH on UNIX), and any path specified on the JVM Profiler Settings page.
- 5 Click Save.
- 6 Restart the server.

▼ To configure the JVM options

On the JVM Options page, you may specify the options of the Java application launcher (java tool) that runs the Application Server. The -D options designate properties that are specific to the Application Server.

- 1 In tree component, select the Application Server node.
- 2 Click the JVM Settings tab.

- 3 Select the **JVM Options** link below the tabs.
- 4 On the **JVM Options** page, to modify an option you edit the **Value** field.
- 5 To add an option:
 - a. Click **Add JVM Option**.
 - b. In the blank row that appears, type the information in the **Value** field.
- 6 To remove an option:
 - a. Select the checkbox next to the option.
 - b. Click **Delete**.
- 7 Click **Save**.
- 8 **Restart the server.**

For more information about JVM options, see <http://java.sun.com/j2se/1.4.2/docs/tooldocs/tools.html> and <http://java.sun.com/docs/hotspot/VMOptions.html>.

▼ **To disable the security manager**

Disabling the Application Server's security manager may improve performance for some types of applications. The J2EE authorization and authentication features will still work even if the security manager has been disabled. You may disable the security manager in a development environment, but you should not disable it in a production environment.

- 1 **Go to the JVM Options page of the Admin Console.**

For instructions, see “[To configure the JVM options](#)” on page 319.
- 2 **On the JVM Options page, remove this option:**
`-Djava.security.policy`
- 3 **Click Save.**
- 4 **Restart the server.**

▼ To configure the JVM profiler settings

A profiler tool generates data that is used to analyze performance and identify potential bottlenecks.

- 1 In tree component, select the Application Server node.
- 2 Click the JVM Settings tab.
- 3 Select the Profiler link below the tabs.
- 4 The information that you specify on the JVM Profiler Settings page depends on which profiler product you're using.
For examples and instructions, see the Debugging J2EE Applications chapter of the *Application Server Developer's Guide*. (For a link to the guide, see [“Further Information” on page 55](#).)
- 5 Click Save.
- 6 Restart the server.

Admin Console Tasks for Advanced Settings

- [“To set the advanced domain attributes” on page 321](#)

▼ To set the advanced domain attributes

- 1 In the tree component, expand the StandAlone instance and select the server instance node.
- 2 Select the Advanced tab.
- 3 On the Domain Attributes page, you may do the following:
 - a. In the Application Root field, identify the full directory path where the applications will be deployed.
 - b. In the Log Root field, specify where the server instance log files are kept.

c. Typically, you will leave the Locale field blank to use the default locale of the host.

A locale is an identifier that specifies a particular combination of language and region. For example, the locale for American English is en_US, and for Japanese it is ja_JP. In order to use a non-English locale, the Application Server must be localized, a process that includes translating English into another language.

4 Click Save.

5 Restart the server.

Compiling and Configuring Apache Web Server

This appendix describes how to compile the Apache source code and configure an installation of Apache Web Server to use the Sun Java System Application Server load balancer plug-in.

Download the appropriate Apache source code. For information about the versions and platforms of Apache Web Server supported for the Application Server, see the *Application Server Release Notes*.

This appendix contains the following topics:

- [“Installing Apache” on page 323](#)

Installing Apache

This section contains the following topics:

- [“Minimum Requirements for Apache 1.3” on page 323](#)
- [“Minimum Requirements for Apache 2” on page 324](#)
- [“To install SSL-aware Apache” on page 325](#)

Minimum Requirements for Apache 1.3

This section describes the minimum requirements to successfully compile Apache 1.3 web server to run the load balancer plug-in. The Apache source must be compiled and built to run with SSL.

Requirements for UNIX and Linux platforms:

- `openssl-0.9.7d` (source)
- `mod_ssl-2.8.16-1.3.29` (source)
- `apache_1.3.29` (sources)
- `gcc-3.3-sol9-sparc-local` packages (for Solaris 9 SPARC/x86).

- flex-2.5.4a-sol9-sparc-local packages (for Solaris 9 SPARC)
- flex-2.5.4a-sol9-intel-local packages (for Solaris 9 x86)

In addition, before compiling Apache:

- On Linux, install the Sun Java System Application Server on the same machine.
- On Solaris 8, ensure that gcc and make are in the PATH.
- On Solaris 9, ensure that gcc version 3.3 and make are in the PATH, and flex is installed.
- If you are using gcc on Red Hat Enterprise Linux Advanced Server 2.1, the version must be later than gcc 3.0.

Note –

- To use another C compiler, set the path of the C compiler and make utility in the PATH environment variable. For example:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:install-dir/lib
```

- These software sources are available at <http://www.sunfreeware.com>.
-

Minimum Requirements for Apache 2

This section describes the minimum requirements to successfully compile Apache 2 web server to run the load balancer plug-in. The Apache source must be compiled and built to run with SSL.

Requirements for UNIX and Linux platforms:

- openssl-0.9.7e (source)
- httpd-2.0.49 (source)
- gcc-3.3-sol9-sparc-local packages (for Solaris SPARC).
- gcc-3.3-sol9-intel-local packages (for Solaris x86)
- flex-2.5.4a-sol9-sparc-local packages (for Solaris SPARC)
- flex-2.5.4a-sol9-intel-local packages (for Solaris x86)

In addition, before compiling Apache:

- On the Linux platform, install the Sun Java System Application Server on the same machine.
- On the Solaris operating system, ensure that gcc version 3.3 and make are in the PATH, and flex is installed.
- On the Solaris 10 operating system, before running make for OpenSSL, run /usr/local/lib/gcc-lib/sparc-sun-solaris2.9/3.3/install-tools/mkheaders (Solaris SPARC) or /usr/local/lib/gcc-lib/i386-pc-solaris2.9/3.3/install-tools/mkheaders (Solaris x86).

- If you are using gcc on Red Hat Enterprise Linux Advanced Server 2.1, the version must be later than gcc 3.0.

Note –

- To use another C compiler, set the path of the C compiler and make utility in the PATH environment variable. For example: `export LD_LIBRARY_PATH=install-dir/lib:$LD_LIBRARY_PATH`. This example is for sh.
 - These software sources are available at <http://www.sunfreeware.com>.
-

▼ To install SSL-aware Apache

Before You Begin You must have already downloaded and uncompressed the Apache software.

- 1 Compile and build OpenSSL. For more information about OpenSSL, see <http://www.openssl.org/>.**

This step is not required on the Linux platform if the version of OpenSSL installed with Linux is 0.9.7.e.

Download and unpack the OpenSSL source.

- `cd openssl-0.9.7e`
- `make`
- `make install`

- 2 For Apache 1.3, configure Apache with `mod_ssl`. You do not need to complete this step for Apache 2. For more information about `mod_ssl`, see <http://www.modssl.org/>.**

Unpack the `mod_ssl` source and follow these steps.

- `cd mod_ssl-2.8.14-1.3.x`
- Run** `./configure --with-apache=../apache_1.3.x --with-ssl=../openssl-0.9.7e --prefix=install-dir --enable-module=ssl --enable-shared=ssl --enable-rule=SHARED_CORE --enable-module=so`

The directory specified in the above command example is a variable. The *prefix* argument indicates where to install Apache. The *x* in the version number represents your actual version.

- 3 For Apache 2.0, configure the source tree:**

- cd to the `http-2.0_x` directory.**

- b. Run `./configure --with-ssl=open-ssl-install-path --prefix=install-dir --enable-ssl --enable-so`

The directory specified in the above command example is a variable. The *prefix* argument indicates where to install Apache. The *x* in the version number represents your actual version.

4 For Apache on Linux 2.1, before compiling:

- a. Open `src/Makefile` and find the end of the automatically generated section.
- b. Add the following lines after the first four lines after the automatically generated section:

```
LIBS+= -licuuc -licui18n -lnspr4 -lpthread -lxerces-c -lsupport -lnsprwrap
-lns-httpd40
```

```
LDLDFLAGS+= -L/install-dir/lib -L/opt/sun/private/lib
```

Note – The `-L/opt/sun/private/lib` part is only required if you installed Application Server as part of a Java Enterprise System installation.

For example:

```
## (End of automatically generated section)##
```

```
CFLAGS=$(OPTIM) $(CFLAGS1) $(EXTRA_CFLAGS)LIBS=$(EXTRA_LIBS)
```

```
$(LIBS1)INCLUDES=$(INCLUDES1) $(INCLUDES0)
```

```
$(EXTRA_INCLUDES)LDLDFLAGS=$(LDLDFLAGS1) $(EXTRA_LDFLAGS)"
```

```
LIBS+= -licuuc -licui18n -lnspr4 -lpthread -lxerces-c -lsupport -lnsprwrap
-lns-httpd40LDLDFLAGS+= -L/install-dir/lib -L/opt/sun/private/lib
```

- c. Create an environment variable `LD_LIBRARY_PATH` equal to: `install-dir/lib` (for all installations) and `install-dir/lib:opt/sun/private/lib` (for Application Server installed as part of a Java Enterprise System installation).
- #### 5 Compile Apache as described in the installation instructions for the version you are using. Full documentation is at <http://httpd.apache.org/>.

In general the steps are:

- a. `make`
- b. `make certificate` (Apache 1.3 only)
- c. `make install`

The command `make certificate` asks for a secure password. Remember this password as it is required for starting secure Apache.

6 Configure Apache for your environment.

Once you have completed the Apache installation, see “Modifications to Apache Web Server” after plug-in installation.

Automatically Restarting a Domain or Node Agent

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

This Appendix contains the following topics:

- “Restarting Automatically on UNIX Platforms” on page 329
- “Restarting Automatically on the Microsoft Windows Platform” on page 330
- “Security for Automatic Restarts” on page 331

Restarting Automatically on UNIX Platforms

To restart your domain on a UNIX platform, add a line of text to the `/etc/inittab` file.

For example, to restart `domain1` for an Application Server installed in the `opt/SUNWappserver` directory, using a password file called `password.txt`:

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

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

To restart a node agent, the syntax is similar. For example, to restart `agent1` for an Application Server installed in the `opt/SUNWappserver` directory, using a password file called `password.txt`:

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

Restarting Automatically on the Microsoft Windows Platform

To restart automatically on Microsoft Windows, create a Windows Service. Use the `appservService.exe` and `appservAgentService.exe` executable files shipped with the Sun Java System Application Server in conjunction with the Service Control command (`sc.exe`) provided by Microsoft.

The `sc.exe` command comes with Windows XP and is either located in the `C:\windows\system32` directory or `C:\winnt\system32` directory. As of this writing, the Windows 2000 `sc.exe` is available for download at <ftp://ftp.microsoft.com/reskit/win2000/sc.zip>. For more information on using `sc.exe`, see http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnllpro/html/msdn_scmslite.asp.

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

```
C:\winnt\system32\sc.exe create service-name binPath= "fully-qualified-path-to-appservService.exe
fully-qualified-path-to-asadmin.bat start-command"
"fully-qualified-path-to-asadmin.bat stop-command"
start= auto DisplayName= "display-name"
```

For example, to create a service called `SunJavaSystemAppServer DOMAIN1` that starts and stops the domain `domain1`, using a password file `C:\Sun\AppServer\password.txt`:

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

To create a service that starts and stops the node agent `agent1`:

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

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

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

Security for Automatic Restarts

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

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

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

If the “interact with desktop” option is not set, the service stays in a “start-pending” state and appears to hang. Kill the service process to recover from this state.
- On Windows or UNIX, create a domain using the `--savemasterpassword=true` option and create a password file to store the admin password. When starting the component, use the `--passwordfile` option to point to the file that contains the password.

For example:

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

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

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

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

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

```
AS_ADMIN_password=password
```

For example, for a password adminadmin:

```
AS_ADMIN_password=adminadmin
```

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

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

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

```
AS_ADMIN_password=password
```

For example, for a password adminadmin:

```
AS_ADMIN_password=adminadmin
```

Dotted Name Attributes for domain.xml

This appendix describes the dotted name attributes that can be used to address the MBean and its attributes. Every element in the `domain.xml` file has a corresponding MBean. Because the syntax for using these names involves separating names between periods, these names are called *dotted names*.

This appendix contains the following topics:

- [“Top Level Elements” on page 333](#)
- [“Elements Not Aliased” on page 335](#)

Top Level Elements

The following conditions must be adhered to for all top level elements in the `domain.xml` file:

- Each server, configuration, cluster, or node agent must have a unique name.
- Servers, configurations, clusters, or node agents cannot be named `domain`.
- Server instances cannot be named `agent`.

The following table identifies the top level elements and the corresponding dotted name prefix.

Element Name	Dotted Name Prefix
<code>applications</code>	<code>domain.applications</code>
<code>resources</code>	<code>domain.resources</code>
<code>configurations</code>	<code>domain.configs</code>
<code>servers</code>	<code>domain.servers</code> Every server contained in this element is accessible as <i>server-name</i> . Where <i>server-name</i> is the value of the name attribute for the server subelement.

Element Name	Dotted Name Prefix
clusters	domain.clusters Every cluster contained in this element is accessible as <i>cluster-name</i> . Where <i>cluster-name</i> is the value of the name attribute for the cluster subelement.
node-agents	domain.node-agents
lb-configs	domain.lb-configs
system-property	domain.system-property

Two levels of aliasing are available:

1. The first level of alias allows access to attributes of server instances or clusters without going through the `domain.servers` or `domain.clusters` prefix. So, for example, a dotted name of the form `server1` maps to the dotted name `domain.servers.server1` (where `server1` is a server instance).
2. The second level of alias is used to refer to configurations, applications, and resources of a cluster or a standalone server instance (`target`).

The following table identifies dotted names beginning with the server name, or cluster name, that are aliased to top level names under the domain:

Dotted Name	Aliased to	Comments
<code>target.applications.*</code>	<code>domain.applications.*</code>	The alias resolves to applications referenced by the <i>target</i> only.
<code>target.resources.*</code>	<code>domain.resources.*</code>	The alias resolves to all <code>jdbc-connection-pool</code> , <code>connector-connection-pool</code> , <code>resource-adapter-config</code> , and all other resources referenced by the <i>target</i> .

The following table identifies dotted names beginning with the server name, or cluster name, that are aliased to top level names within the configuration referenced by the server or cluster.

Dotted Name	Aliased to
<code>target.http-service</code>	<code>config-name.http-service</code>
<code>target.iiop-service</code>	<code>config-name.iiop-service</code>
<code>target.admin-service</code>	<code>config-name.admin-service</code>
<code>target.web-container</code>	<code>config-name.web-container</code>
<code>target.ejb-container</code>	<code>config-name.ejb-container</code>
<code>target.mdb-container</code>	<code>config-name.mdb-container</code>

Dotted Name	Aliased to
<i>target.jms-service</i>	<i>config-name.jms-service</i>
<i>target.log-service</i>	<i>config-name.log-service</i>
<i>target.security-service</i>	<i>config-name.security-service</i>
<i>target.transaction-service</i>	<i>config-name.transaction-service</i>
<i>target.monitoring-service</i>	<i>config-name.monitoring-service</i>
<i>target.java-config</i>	<i>config-name.java-config</i>
<i>target.availability-service</i>	<i>config-name.availability-service</i>
<i>target.thread-pools</i>	<i>config-name.thread-pools</i>

Elements Not Aliased

A clustered instance should not be aliased. To get a system property for a clustered instance, the dotted name attribute you should use is as follows:

domain.servers.clustered-instance-name.system-property, not
clustered-instance-name.system-property.

Index

A

ACC

See containers

application client, 135

acceptor threads, in HTTP listeners, 231

access log, HTTP service, 235-236

accesslog property, virtual servers, 241

AddressList property, 98

AddressListBehavior property, 99

AddressListIterations property, 99

Admin Console, 33

allowLinking property, virtual servers, 242

append-version property, 108

applets, 135

application client JAR files, 59

application client modules, deploying, 70

Application Server

restart, 142

shut down, 142

Application Server domains, 34

applications

auto deploy, 77

deploying on virtual servers, 75

deployment plan, 79

directory deployment, 78

disabling, 74

enabling, 74

listing deployed, 72

listing subcomponents, 73

module descriptors, 73

naming conventions, 59

performance, 139

applications (*Continued*)

redeploying, 58,75

undeploying, 73-74

asadmin command, 257,258

create-threadpool, 257

delete-threadpool, 258

asadmin utility, 33

asadmintruststore file, 36

auto-deploying applications, 77

B

bean-cache, monitoring attribute names, 275

C

cache-hits, 275

cache-misses, 275

caching

cleanup, 141

disabling, 139

Enterprise JavaBeans, 139

timeouts, 141

central repository, applications deployed to, 57

chunkedRequestBufferSize property, 234

chunkedRequestTimeoutSeconds property, 235

classpath, in lifecycle modules, 69

client access, 31

ClientID property, 98

clustering, 29

clusters, defining, 41

- connection factories, JMS
 - AddressList property, 98
 - AddressListBehavior property, 99
 - AddressListIterations property, 99
 - ClientID property, 98
 - creating, 97-100
 - deleting, 101
 - editing, 100
 - MessageServiceAddressList property, 99
 - overview, 95-97
 - Password property, 99
 - ReconnectAttempts property, 99
 - ReconnectEnabled property, 99
 - ReconnectInterval property, 99
 - transaction support, 98
 - UserName property, 99
- connection pool, HTTP service, 237-238
- connector, 32
- connector connection pools, JMS resources and, 97
- connector modules, deploying, 67
- connector resources, JMS resources and, 97
- connectors, modules, 255
- container, 31
- containers
 - applet, 135
 - application client, 135
 - Enterprise JavaBeans, 135, 136, 139-141
 - configuring, 139-141
 - J2EE, 135
 - servlet
 - See containers, 135
 - web, 135
 - web, 135
- CORBA, 249
 - threads, 255
- create-domain command, 35
- create-jndi-resource command, 123
- custom resources
 - creating, 120
 - deleting, 121
 - listing, 121
 - using, 119
- D**
- databases
 - JNDI names, 118
 - Oracle, 143
 - PointBase, 143
 - resource references, 118
- defining clusters, 41
- delete-domain command, 36
- deployment plan, 79
- Description property, JMS destination resources, 102
- destinations, JMS
 - creating destination resources, 101-102
 - creating physical destinations, 103-105
 - deleting destination resources, 103
 - deleting physical destinations, 105
 - Description property, 102
 - editing destination resources, 102-103
 - maxNumActiveConsumers property, 104
 - Name property, 102
 - overview, 95-97
- directory deployment, 78
- disabling applications, 74
- dnsCacheEnabled property, 235
- docroot property, virtual servers, 241
- documentation, overview, 24-26
- domain, deploying applications to, 58
- domains, 35
 - creating, 35
- E**
- EAR files, 59
- EJB JAR files, 59
- EJB modules, deploying, 65
- enabling applications, 74
- enterprise applications, 59
 - deploying, 60
- Enterprise Java Beans, threads, 255
- Enterprise JavaBeans
 - activating, 136
 - active, 140
 - authorization, 136
 - caching, 136, 139, 140-141
 - creating, 136

Enterprise JavaBeans (Continued)

- entity, 136, 139, 140-141
 - idle, 139, 140
 - message-driven, 136, 141-142
 - passivating, 136, 139, 140
 - persistent, 136
 - pooling, 139-140, 141-142
 - removing from cache, 141
 - removing idle, 142
 - session, 136
 - stateful session, 140-141, 142
 - stateless session, 139
 - timer service, 142-143
- entity beans
- See Enterprise JavaBeans
 - entity, 139
- execution-time-millis, 273
- external repositories, accessing, 122
- external resource, deleting, 124
- external resources
- creating, 122-123
 - editing, 123-124

G

- get command, monitoring data, 302

H

- high availability, 30
- HTTP file cache, HTTP service, 239-240
- HTTP listeners
 - acceptor threads, 231
 - creating, 244-246
 - default virtual server, 231
 - deleting, 247
 - editing, 246-247
 - overview, 231-233
- HTTP ports, changing, 52-53
- HTTP protocol, HTTP service, 238-239
- HTTP service
 - access log, 235-236
 - chunkedRequestBufferSize property, 234

HTTP service (Continued)

- chunkedRequestTimeoutSeconds property, 235
- HTTP Service, configuring, 233-235
- HTTP service
 - connection pool, 237-238
 - dnsCacheEnabled property, 235
 - HTTP file cache, 239-240
 - HTTP listeners, 231-233
 - HTTP protocol, 238-239
 - Keep-Alive subsystem, 232, 237
 - keepAliveQueryMaxSleepTime property, 234
 - keepAliveQueryMeanTime property, 234
 - monitoringCacheEnabled property, 234
 - monitoringCacheRefreshInMillis property, 234
 - overview, 229
 - request processing threads, 232, 236-237
 - ssl3SessionTimeout property, 234
 - sslCacheEntries property, 234
 - sslClientAuthDataLimit property, 234
 - sslClientAuthTimeout property, 234
 - sslSessionTimeout property, 234
 - stackSize property, 234
 - statsProfilingEnabled property, 234
 - traceEnabled property, 233
 - virtual servers, 229-230
- HTTP sessions, 136

I

- IIOP listeners, 250
 - creating, 251-252
 - deleting, 253-254
 - editing, 253
- IIOP ports, changing, 53
- instance-name property, 108
- instance-name-suffix property, 108
- instances, 41

J

- J2EE group, 173
- J2SE software, 55
- Java Message Service (JMS), See JMS resources, 95

Java Naming and Directory Service, See JNDI, 136

Javadocs, 25

JavaMail, 32

JavaMail API, overview, 113

JavaMail sessions

 creating, 114-115

 deleting, 116

 editing, 115-116

JavaServer Pages, 135

JCE provider

 configuring, 209

JDBC, 32

 drivers, 224

 resources, 143

JMS hosts

 creating, 109-110

 deleting, 111-112

 editing, 111

jms-max-messages-load, 275

JMS provider, 95

 append-version property, 108

 configuring, 105-109

 instance-name property, 108

 instance-name-suffix property, 108

 JMS hosts, 109-110, 111

JMS resources

 connection factory resources, 95-97, 97-100, 100, 101

 destination resources, 95-97, 101-102, 102-103, 103

 overview, 95-97

 physical destinations, 95-97, 103-105, 105

 queues, 95-97

 topics, 95-97

jmsra system resource adapter, 97

JNDI, 136

 custom resources, creating, 120

 custom resources, deleting, 121

 custom resources, using, 119

 external repositories, 122

 external resource, deleting, 124

 external resources, creating, 122-123

 external resources, editing, 123-124

 lookup names for EJB components, 60

 lookups and associated references, 119

JNDI (*Continued*)

 names, 118, 143

JSP, See JavaServer Pages, 135

K

Keep-Alive subsystem

 HTTP service, 232, 237

keepAliveQueryMaxSleepTime property, 234

keepAliveQueryMeanTime property, 234

kestore.jks file, 191-192

keypoint intervals, 227

keypoint operations, 227

L

lifecycle modules

 classpath, 69

 creating, 68

 load order, 69

list command, monitoring, 301

list-custom-resources command, 121

list-domains command, 36

list-jndi-resource command, 124

load order, in lifecycle modules., 69

log levels, configuring, 263-264

log records, 259-260

logging

 configuring general settings, 262-263

 configuring levels, 263-264

 logger namespaces, 260-261

 overview, 259-260

 transactions, 226-227

 viewing the server log, 264-266

M

man pages, 33

max_history_files property, 262-263

maxNumActiveConsumers property, JMS physical destinations, 104

Message Queue software, 95

MessageServiceAddressList property, 99
 Messaging, 32
 module descriptors, viewing, 73
 monitoring

- bean-cache attributes, 275
- container subsystems, 268-269
- ORB service, 281-282
- transaction service, 282-283
- using get command, 302
- using list command, 301

 monitoringCacheEnabled property, 234
 monitoringCacheRefreshInMillis property, 234

N

Name property, JMS destination resources, 102
 naming, JNDI and resource reference, 119
 naming and directory service, 32
 naming conventions, for applications, 59
 naming service, 32
 numbeansinpool, 274
 numexpiredsessionsremoved, 275
 numpassivationerrors, 275
 numpassivations, 275
 numpassivationsuccess, 275
 numthreadswaiting, 274

O

Oasis Web Services Security, *See* WSS
 object request broker, threads, 255
 Object Request Broker (ORB), 249

- configuring, 250-251
- overview, 250

 online help, 55
 Oracle, 143
 ORB, 249

- configuring, 250-251
- IIOP listeners, 250
- overview, 250
- See object request broker, 255
- service, monitoring, 281-282

P

Password property, 99
 performance

- increasing, 139
- problems, 139
- thread pools, 255

 PointBase, 143
 pooling

- Enterprise JavaBeans, 139-140, 141-142

 Port listeners, 51
 port numbers, changing, 52
 port numbers, viewing, 52

Q

queues

- work
 - See thread pools, 256

 queues, JMS, 95-97

R

RAR files, 59
 realms, certificate, 157
 reap interval, 137, 138
 ReconnectAttempts property, 99
 ReconnectEnabled property, 99
 ReconnectInterval property, 99
 redeploying applications, 58, 75
 request processing threads

- HTTP service, 232, 236-237

 resource adapters, 224

- deploying, 67
- jmsra, 97

 resource managers, 224
 Resource RAR files, 59
 resource references, 118
 restart server, 37
 rollback

- See transactions
 - rolling back, 223

 RSA encryption, 209

S

- security, 32
- server administration, 32
- server log, viewing, 264-266
- services, timer, 142-143
- services for applications, 32
- servlets, 135
- session manager, 137
- sessions
 - configuring, 136-139
 - custom IDs, 138
 - deleting, 138
 - deleting data, 137
 - file name, 138
 - HTTP, 136, 139
 - IDs, 138
 - inactive, 137, 138
 - managing, 137
 - storing, 139
 - storing data, 138
 - timeouts, 137
- single sign-on, virtual server properties, 242
- Solaris
 - patches, 26
 - support, 26
- ssl3SessionTimeout property, 234
- sslCacheEntries property, 234
- sslClientAuthDataLimit property, 234
- sslClientAuthTimeout property, 234
- sslSessionTimeout property, 234
- sso-enabled property, virtual servers, 242
- sso-max-inactive-seconds property, virtual servers, 242
- sso-reap-interval-seconds property, virtual servers, 242
- stackSize property, 234
- start-domain command, 36, 121, 124
- stateful session beans, See Enterprise JavaBeans, 140-141
- stateless session beans, See Enterprise JavaBeans, 139
- statsProfilingEnabled property, 234
- stop-domain command, 37
- subcomponents of applications, listing, 73
- Sun Java System Message Queue software, 95

T

- targets
 - managing application, 74
 - of deployed applications, 58
- thread pools, 255
 - creating, 256-257
 - deleting, 258
 - editing, 257
 - idle, 256, 257
 - naming, 256-257
 - performance, 255
 - thread starvation, 255
 - timeouts, 256, 257
 - work queues, 256, 257
- threads
 - removing, 256, 257
 - See thread pools, 255
- timeouts, 140, 141, 142
 - thread pools, 256, 257
- timer service
 - See Enterprise JavaBeans timer service, 142-143
- timers
 - See Enterprise JavaBeans timer service, 142-143
- topics, JMS, 95-97
- total-beans-created, 275
- total-beans-destroyed, 275
- total-num-errors, 273
- total-num-success, 273
- traceEnabled property, 233
- transaction management, 32
- Transaction Manager
 - See transactions managers, 224
- transaction service, monitoring, 282-283
- transactions, 223
 - associating, 224
 - attributes, 224
 - committing, 223
 - completing, 224
 - demarcations, 224
 - distributed, 224
 - Enterprise JavaBeans, 139

transactions (*Continued*)

- JMS connection factories, 98
 - logging, 226-227
 - managers, 224
 - recovering, 224, 225
 - rolling back, 223
 - timeouts, 226
- truststore.jks file, 191-192

U

- undeploying applications, 73-74
- UserName property, 99

V

- virtual servers
 - accesslog property, 241
 - allowLinking property, 242
 - creating, 240-242
 - deleting, 243-244
 - deploying applications to additional, 75
 - docroot property, 241
 - editing, 243
 - overview, 229-230
 - sso-enabled property, 242
 - sso-max-inactive-seconds property, 242
 - sso-reap-interval-seconds property, 242

W

- WAR files, 59
- web applications, 59
 - deploying, 63
 - launching, 65
- web services, 32
- web sessions, See HTTP sessions, 136
- work queues, See thread pools, 256

