

# The Java EE 6 Tutorial, Volume II

Advanced Topics



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

Part No: 820-7628-10  
December 2009

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

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

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

This distribution may include materials developed by third parties.

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

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Enterprise JavaBeans, EJB, GlassFish, J2EE, J2SE, Java Naming and Directory Interface, JavaBeans, Javadoc, JDBC, JDK, JavaScript, JavaServer, JavaServer Pages, JMX, JRE, JSP, JVM, MySQL, NetBeans, OpenSolaris, SunSolve, Sun GlassFish, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

---

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

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

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Enterprise JavaBeans, EJB, GlassFish, J2EE, J2SE, Java Naming and Directory Interface, JavaBeans, Javadoc, JDBC, JDK, JavaScript, JavaServer, JavaServer Pages, JMX, JRE, JSP, JVM, MySQL, NetBeans, OpenSolaris, SunSolve, Sun GlassFish, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, aux États-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 États-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 États-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

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

# Contents

---

<b>Preface</b> .....	23
<b>Part I Introduction</b> .....	31
<b>1 Overview</b> .....	33
Java EE 6 Highlights .....	34
Java EE Application Model .....	34
Distributed Multitiered Applications .....	35
Security .....	36
Java EE Components .....	37
Java EE Clients .....	37
Web Components .....	39
Business Components .....	40
Enterprise Information System Tier .....	41
Java EE Containers .....	41
Container Services .....	42
Container Types .....	42
Web Services Support .....	44
XML .....	44
SOAP Transport Protocol .....	45
WSDL Standard Format .....	45
Java EE Application Assembly and Deployment .....	45
Packaging Applications .....	46
Development Roles .....	47
Java EE Product Provider .....	48
Tool Provider .....	48
Application Component Provider .....	48

Application Assembler .....	49
Application Deployer and Administrator .....	49
Java EE 6 APIs .....	50
Enterprise JavaBeans Technology .....	50
Java Servlet Technology .....	50
JavaServer Faces Technology .....	51
JavaServer Pages Technology .....	51
JavaServer Pages Standard Tag Library .....	51
Java Persistence API .....	52
Java Transaction API .....	52
Java API for RESTful Web Services (JAX-RS) .....	52
Java Message Service API .....	52
Java EE Connector Architecture .....	53
JavaMail API .....	53
Java Authorization Service Provider Contract for Containers (Java ACC) .....	53
Java Authentication Service Provider Interface for Containers (JASPIC) .....	53
Java API for XML Registries .....	54
Simplified Systems Integration .....	54
Java EE 6 APIs Included in the Java Platform, Standard Edition 6.0 (Java SE 6) .....	55
Java Database Connectivity API .....	55
Java Naming and Directory Interface .....	55
JavaBeans Activation Framework .....	56
Java API for XML Processing .....	56
Java Architecture for XML Binding (JAXB) .....	56
SOAP with Attachments API for Java .....	56
Java API for XML Web Services (JAX-WS) .....	57
Java Authentication and Authorization Service .....	57
Sun GlassFish Enterprise Server v3 .....	57
Tools .....	57
<b>2 Using the Tutorial Examples .....</b>	<b>59</b>
Required Software .....	59
Java Platform, Standard Edition .....	59
Java EE 6 Software Development Kit (SDK) .....	60
Apache Ant .....	60

Java EE 6 Tutorial Component .....	61
NetBeans IDE .....	62
Starting and Stopping the Enterprise Server .....	63
Starting the Administration Console .....	64
▼ To Start the Administration Console in NetBeans IDE .....	64
Starting and Stopping the Java DB Database Server .....	64
Building the Examples .....	65
Tutorial Example Directory Structure .....	65
Getting the Latest Updates to the Tutorial .....	66
▼ To Update the Tutorial through the Update Center .....	66
Debugging Java EE Applications .....	66
Using the Server Log .....	66
Using a Debugger .....	67
<b>Part II The Web Tier .....</b>	<b>69</b>
<b>3 Getting Started with Web Applications .....</b>	<b>71</b>
Web Applications .....	71
Web Application Life Cycle .....	73
Web Modules .....	75
Packaging Web Modules .....	76
Deploying a WAR File .....	77
Testing Deployed Web Modules .....	79
Listing Deployed Web Modules .....	79
Updating Web Modules .....	79
Undeploying Web Modules .....	80
Configuring Web Applications .....	81
Mapping URLs to Web Components .....	81
Declaring Welcome Files .....	83
Setting Initialization Parameters .....	83
Mapping Errors to Error Screens .....	85
Declaring Resource References .....	86
Further Information about Web Applications .....	88

<b>4</b>	<b>JavaServer Faces Technology</b> .....	89
	What Is a JavaServer Faces Application? .....	90
	JavaServer Faces Technology Benefits .....	91
	Creating a Simple JavaServer Faces Application .....	92
	Developing Backing Beans .....	92
	Creating the Web Page .....	93
	Mapping the Faces Servlet Instance .....	94
	The Lifecycle of the helloWorld Application .....	94
	▼ Running the Application in NetBeans IDE .....	95
	Further Information about JavaServer Faces Technology .....	96
<b>5</b>	<b>Introduction to Facelets</b> .....	97
	Advantages of Facelets .....	97
	What's Facelets ? .....	98
	Web Pages .....	98
	Tag Library Support .....	98
	Unified Expression Language Support .....	99
	Developing a Simple Facelets Application .....	99
	Creating a Facelets Application .....	100
	Configuring the Application .....	103
	Building, Packaging, Deploying and Running the Application .....	105
	Templating .....	107
	Composite Components .....	109
	Resources .....	112
<b>6</b>	<b>Unified Expression Language</b> .....	113
	Overview of EL .....	113
	Immediate and Deferred Evaluation Syntax .....	114
	Value and Method Expressions .....	116
	Defining a Tag Attribute Type .....	122
	Literal Expressions .....	123
	Operators .....	124
	Reserved Words .....	125
	Examples of EL Expressions .....	125

---

<b>7 Using JavaServer Faces Technology in Web Pages</b> .....	127
Setting Up a Page .....	127
Adding Components to a Page Using HTML Tags .....	128
Common Component Tag Attributes .....	131
Adding HTML Head and Body Tags .....	133
Adding a Form Component .....	133
Using Text Components .....	134
Using Command Components for Performing Actions and Navigation .....	139
Adding Graphics and Images With the <code>h:graphicImage</code> Tag .....	141
Laying Out Components With the <code>Panel</code> Component .....	141
Displaying Components for Selecting One Value .....	143
Rendering Components for Selecting Multiple Values .....	145
Using The <code>SelectItem</code> and <code>SelectItems</code> Components .....	147
Using Data-Bound Table Components .....	148
Displaying Error Messages With the <code>h:message</code> and <code>h:messages</code> Tags .....	152
Creating Bookmarkable URLs with <code>h:button</code> and <code>h:link</code> Tags .....	153
Resource Relocation using <code>h:output</code> Tags .....	154
Using Core Tags .....	157
<b>8 Using Converters, Listeners and Validators</b> .....	159
Using the Standard Converters .....	159
Converting a Component's Value .....	160
Using <code>DateTimeConverter</code> .....	161
Using <code>NumberConverter</code> .....	163
Registering Listeners on Components .....	164
Registering a Value-Change Listener on a Component .....	165
Registering an Action Listener on a Component .....	166
Using the Standard Validators .....	166
Validating a Component's Value .....	167
Using the <code>LongRangeValidator</code> .....	167
Referencing a Backing Bean Method .....	168
Referencing a Method That Performs Navigation .....	169
Referencing a Method That Handles an Action Event .....	169
Referencing a Method That Performs Validation .....	169
Referencing a Method That Handles a Value-Change Event .....	170

<b>9</b>	<b>Developing With JavaServer Faces Technology</b> .....	171
	Backing Beans .....	171
	Creating a Backing Bean .....	171
	Writing Bean Properties .....	174
	Writing Properties Bound to Component Values .....	175
	Writing Properties Bound to Component Instances .....	181
	Writing Properties Bound to Converters, Listeners, or Validators .....	182
	Writing Backing Bean Methods .....	183
	Writing a Method to Handle Navigation .....	183
	Writing a Method to Handle an Action Event .....	185
	Writing a Method to Perform Validation .....	185
	Writing a Method to Handle a Value-Change Event .....	186
	Bean Validation .....	187
<b>10</b>	<b>Configuring JavaServer Faces Applications</b> .....	189
	Application Configuration Resource File .....	189
	Ordering of Application Configuration Resources .....	190
	Configuring Project Stage .....	192
	Configuring Beans .....	192
	Using the managed - bean Element .....	193
	Initializing Properties Using the managed - property Element .....	195
	Initializing Maps and Lists .....	201
	Using Annotations .....	201
	Registering Custom Error Messages as a Resource Bundle .....	202
	Registering Custom Localized Static Text .....	203
	Default Validator .....	203
	Configuring Navigation Rules .....	204
	Implicit Navigation Rules .....	207
	Basic Requirements of a JavaServer Faces Application .....	207
	Configuring an Application With a Deployment Descriptor .....	208
	Including the Classes, Pages, and Other Resources .....	212
<b>11</b>	<b>Java Servlet Technology</b> .....	213
	What Is a Servlet? .....	213
	Servlet Life Cycle .....	214

Handling Servlet Life-Cycle Events .....	214
Handling Servlet Errors .....	215
Sharing Information .....	216
Using Scope Objects .....	216
Controlling Concurrent Access to Shared Resources .....	216
Initializing a Servlet .....	217
Writing Service Methods .....	217
Getting Information from Requests .....	218
Constructing Responses .....	219
Filtering Requests and Responses .....	220
Programming Filters .....	220
Programming Customized Requests and Responses .....	221
Specifying Filter Mappings .....	222
Invoking Other Web Resources .....	224
Including Other Resources in the Response .....	225
Transferring Control to Another Web Component .....	225
Accessing the Web Context .....	226
Maintaining Client State .....	226
Accessing a Session .....	226
Associating Objects with a Session .....	226
Session Management .....	227
Session Tracking .....	227
Finalizing a Servlet .....	228
Tracking Service Requests .....	228
Notifying Methods to Shut Down .....	229
Creating Polite Long-Running Methods .....	230
Further Information about Java Servlet Technology .....	230
<b>Part III Web Services .....</b>	<b>231</b>
<b>12 Introduction to Web Services .....</b>	<b>233</b>
What Are Web Services? .....	233
Types of Web Services .....	233
Deciding Which Type of Web Service to Use .....	236
When Should I Use JAX-WS? .....	236

When Should I Use JAX-RS? .....	236
<b>13 Building Web Services with JAX-WS .....</b>	<b>237</b>
Setting the Port .....	238
Creating a Simple Web Service and Client with JAX-WS .....	238
Requirements of a JAX-WS Endpoint .....	239
Coding the Service Endpoint Implementation Class .....	240
Building, Packaging, and Deploying the Service .....	240
Testing the Service without a Client .....	242
A Simple JAX-WS Client .....	242
Types Supported by JAX-WS .....	244
Web Services Interoperability and JAX-WS .....	245
Further Information about JAX-WS .....	245
<b>14 Building RESTful Web Services with JAX-RS and Jersey .....</b>	<b>247</b>
What are RESTful Web Services? .....	247
Where Does Jersey Fit In? .....	248
Creating a RESTful Root Resource Class .....	248
Developing RESTful Web Services with JAX-RS and Jersey .....	249
Overview of a Jersey-Annotated Application .....	250
The @Path Annotation and URI Path Templates .....	251
Responding to HTTP Resources .....	254
Using @Consumes and @Produces to Customize Requests and Responses .....	257
Extracting Request Parameters .....	259
Overview of JAX-RS and Jersey: Further Information .....	262
Example Applications for JAX-RS and Jersey .....	263
Creating a RESTful Web Service .....	263
Example: Creating a Simple Hello World Application Using JAX-RS and Jersey .....	270
Example: Adding on to the Simple Hello World RESTful Web Service .....	273
JAX-RS in the First Cup Example .....	275
Real World Examples .....	275
Further Information .....	275

---

<b>15</b>	<b>Advanced Topics in JAX-RS</b> .....	277
	Advanced Topics in JAX-RS .....	277
<b>Part IV</b>	<b>Enterprise Beans</b> .....	279
<b>16</b>	<b>Enterprise Beans</b> .....	281
	What Is an Enterprise Bean? .....	281
	Benefits of Enterprise Beans .....	281
	When to Use Enterprise Beans .....	282
	Types of Enterprise Beans .....	282
	What Is a Session Bean? .....	283
	Types of Session Beans .....	283
	When to Use Session Beans .....	284
	What Is a Message-Driven Bean? .....	285
	What Makes Message-Driven Beans Different from Session Beans? .....	285
	When to Use Message-Driven Beans .....	286
	Accessing Enterprise Beans .....	286
	Using Enterprise Beans in Clients .....	287
	Deciding on Remote or Local Access .....	288
	Local Clients .....	289
	Remote Clients .....	291
	Web Service Clients .....	292
	Method Parameters and Access .....	293
	The Contents of an Enterprise Bean .....	294
	Packaging Enterprise Beans In EJB JAR Modules .....	294
	Packaging Enterprise Beans in WAR Modules .....	295
	Naming Conventions for Enterprise Beans .....	296
	The Life Cycles of Enterprise Beans .....	297
	The Life Cycle of a Stateful Session Bean .....	297
	The Lifecycle of a Stateless Session Bean .....	298
	The Lifecycle of a Singleton Session Bean .....	298
	The Lifecycle of a Message-Driven Bean .....	299
	Further Information about Enterprise Beans .....	300

<b>17</b>	<b>Getting Started with Enterprise Beans</b> .....	301
	Creating the Enterprise Bean .....	301
	Coding the Enterprise Bean .....	302
	Creating the converter Web Client .....	302
	Compiling, Packaging, and Running the converter Example .....	303
	Modifying the Java EE Application .....	305
	Modifying a Class File .....	305
<b>18</b>	<b>Running the Enterprise Bean Examples</b> .....	307
	The cart Example .....	307
	The Business Interface .....	308
	Session Bean Class .....	308
	The Remove Method .....	312
	Helper Classes .....	312
	Building, Packaging, Deploying, and Running the cart Example .....	312
	Undeploying the cart Example .....	315
	A Singleton Session Bean Example: counter .....	315
	Creating a Singleton Session Bean .....	315
	The Architecture of the counter Example .....	320
	Building, Deploying, and Running the counter Example .....	323
	A Web Service Example: helloservice .....	324
	The Web Service Endpoint Implementation Class .....	325
	Stateless Session Bean Implementation Class .....	325
	Building, Packaging, Deploying, and Testing the helloservice Example .....	326
	Using the Timer Service .....	328
	Creating Calendar-Based Timer Expressions .....	328
	Programmatic Timers .....	331
	Automatic Timers .....	333
	Canceling and Saving Timers .....	334
	Getting Timer Information .....	334
	Transactions and Timers .....	335
	The timersession Example .....	335
	Building, Packaging, Deploying, and Running the timersession Example .....	338
	Handling Exceptions .....	339

---

<b>19</b>	<b>A Message-Driven Bean Example</b> .....	341
	simpleMessage Example Application Overview .....	341
	The simpleMessage Application Client .....	342
	The Message-Driven Bean Class .....	343
	The onMessage Method .....	344
	Packaging, Deploying, and Running the simpleMessage Example .....	345
	Creating the Administered Objects for the simpleMessage Example .....	345
	Building, Deploying, and Running the simpleMessage Application Using NetBeans IDE .....	346
	Building, Deploying, and Running the simpleMessage Application Using Ant .....	347
	Removing the Administered Objects for the simpleMessage Example .....	348
<b>Part V</b>	<b>Contexts and Dependency Injection for the Java EE Platform</b> .....	349
<b>20</b>	<b>Introduction to Contexts and Dependency Injection for the Java EE Platform</b> .....	351
	Overview of Contexts and Dependency Injection for the Java EE Platform .....	352
	About Beans .....	353
	About Managed Beans .....	353
	Beans as Injectable Objects .....	354
	Using Qualifiers .....	355
	Injecting Beans .....	356
	Using Scopes .....	356
	Giving Beans EL Names .....	358
	Adding Setter and Getter Methods .....	358
	Using a Managed Bean in a Facelets Page .....	359
	Injecting Objects by Using Producer Methods .....	360
	Configuring a CDI Application .....	361
	Further Information .....	361
<b>21</b>	<b>Running the Basic Contexts and Dependency Injection Examples</b> .....	363
	The simplegreeting Example .....	363
	The simplegreeting Source Files .....	364
	The Facelets Template and Page .....	364
	Configuration Files .....	366
	Building, Packaging, Deploying, and Running the simplegreeting Example .....	366

The guessnumber Example .....	368
The guessnumber Source Files .....	369
The Facelets Page .....	373
Building, Packaging, Deploying, and Running the guessnumber Example .....	375
<b>Part VI Persistence .....</b>	<b>379</b>
<b>22 Introduction to the Java Persistence API .....</b>	<b>381</b>
Entities .....	381
Requirements for Entity Classes .....	381
Persistent Fields and Properties in Entity Classes .....	382
Primary Keys in Entities .....	385
Multiplicity in Entity Relationships .....	387
Direction in Entity Relationships .....	388
Embeddable Classes in Entities .....	390
Entity Inheritance .....	391
Managing Entities .....	395
The Persistence Context .....	396
The EntityManager Interface .....	396
Persistence Units .....	400
Querying Entities .....	401
<b>23 Running the Persistence Examples .....</b>	<b>403</b>
The order Application .....	403
Entity Relationships in the order Application .....	403
Primary Keys in the order Application .....	406
Entity Mapped to More Than One Database Table .....	409
Cascade Operations in the order Application .....	410
BLOB and CLOB Database Types in the order Application .....	410
Temporal Types in the order Application .....	411
Managing the order Application's Entities .....	411
Building and Running the order Application .....	414
The roster Application .....	415
Relationships in the roster Application .....	415

---

Entity Inheritance in the roster Application .....	416
Criteria Queries in the roster Application .....	418
Automatic Table Generation in the roster Application .....	420
Building and Running the roster Application .....	420
<b>24 The Java Persistence Query Language .....</b>	<b>423</b>
Query Language Terminology .....	423
Creating Queries Using the Java Persistence Query Language .....	424
Named Parameters in Queries .....	424
Positional Parameters in Queries .....	425
Simplified Query Language Syntax .....	425
Select Statements .....	425
Update and Delete Statements .....	426
Example Queries .....	426
Simple Queries .....	426
Queries That Navigate to Related Entities .....	428
Queries with Other Conditional Expressions .....	429
Bulk Updates and Deletes .....	431
Full Query Language Syntax .....	432
BNF Symbols .....	432
BNF Grammar of the Java Persistence Query Language .....	432
FROM Clause .....	437
Path Expressions .....	441
WHERE Clause .....	442
SELECT Clause .....	452
ORDER BY Clause .....	454
The GROUP BY Clause .....	454
<b>25 Creating Queries Using the Criteria API .....</b>	<b>457</b>
Overview of the Criteria and Metamodel APIs .....	457
Modeling Entity Classes with the Metamodel API .....	459
Using Metamodel Classes .....	460
Basic Type-Safe Queries Using the Criteria API and Metamodel API .....	461
Creating a Criteria Query .....	461
Query Roots .....	462

Querying Relationships Using Joins .....	463
Path Navigation in Criteria Queries .....	463
Restricting Criteria Query Results .....	464
Managing Criteria Query Results .....	467
Executing Queries .....	468
<b>Part VII Security .....</b>	<b>471</b>
<b>26 Introduction to Security in the Java EE Platform .....</b>	<b>473</b>
Overview of Java EE Security .....	474
A Simple Security Example .....	474
Security Functions .....	477
Characteristics of Application Security .....	478
Security Implementation Mechanisms .....	479
Java SE Security Implementation Mechanisms .....	479
Java EE Security Implementation Mechanisms .....	480
Securing Containers .....	482
Using Deployment Descriptors for Declarative Security .....	482
Using Annotations .....	483
Using Programmatic Security .....	484
Securing the Enterprise Server .....	484
Working with Realms, Users, Groups, and Roles .....	486
What Are Realms, Users, Groups, and Roles? .....	486
Managing Users and Groups on the Enterprise Server .....	489
Setting Up Security Roles .....	491
Mapping Roles to Users and Groups .....	493
Establishing a Secure Connection Using SSL .....	494
Installing and Configuring SSL Support .....	494
Specifying a Secure Connection in Your Application Deployment Descriptor .....	495
Verifying SSL Support .....	496
Working with Digital Certificates .....	497
Further Information about Security .....	500

---

<b>27</b>	<b>Getting Started Securing Enterprise Applications</b> .....	503
	Responsibility for Administering Security .....	503
	Securing Enterprise Beans .....	504
	Securing an Enterprise Bean Using Declarative Security and Annotations .....	507
	Securing an Enterprise Bean Programmatically .....	515
	Propagating a Security Identity (Run-As) .....	518
	Deploying Secure Enterprise Beans .....	520
	Securing Application Clients .....	521
	Using Login Modules .....	521
	Using Programmatic Login .....	522
	Securing Enterprise Information Systems (EIS) Applications .....	523
	Container-Managed Sign-On .....	523
	Component-Managed Sign-On .....	523
	Configuring Resource Adapter Security .....	524
	Mapping an Application Principal to EIS Principals .....	526
<b>28</b>	<b>Getting Started Securing Web Applications</b> .....	527
	Overview of Web Application Security .....	528
	Using Deployment Descriptors to Secure Web Applications .....	530
	Introduction to Web Application Deployment Descriptors .....	530
	Specifying Security Constraints .....	533
	Specifying an Authentication Mechanism .....	537
	Working with Security Roles .....	544
	Using Programmatic Security with Web Applications .....	548
	Authenticating Users Programmatically .....	549
	Checking Caller Identity Programmatically .....	551
	Example Code for Programmatic Security .....	551
	Declaring and Linking Role References .....	553
	Using Message Security with Web Applications .....	555
	Examples: Securing Web Applications .....	555
	Setting Up Your System for Running the Security Examples .....	556
	Example: Basic Authentication with a Servlet .....	556
	Example: Basic Authentication with JAX-WS .....	562
	Example: Form-Based Authentication with a Servlet .....	566

<b>29</b>	<b>Advanced Concepts in Java EE Security</b> .....	573
	Advanced Concepts in Java EE Security .....	573
<b>Part VIII</b>	<b>Java EE Supporting Technologies</b> .....	575
<b>30</b>	<b>Introduction to Java EE Supporting Technologies</b> .....	577
	Transactions .....	577
	Resources .....	578
	The Java EE Connector Architecture and Resource Adapters .....	578
	Java Message Service .....	579
	Java DataBase Connectivity (JDBC) Software .....	579
<b>31</b>	<b>Transactions</b> .....	581
	What Is a Transaction? .....	581
	Container-Managed Transactions .....	582
	Transaction Attributes .....	583
	Rolling Back a Container-Managed Transaction .....	586
	Synchronizing a Session Bean's Instance Variables .....	587
	Methods Not Allowed in Container-Managed Transactions .....	587
	Bean-Managed Transactions .....	587
	JTA Transactions .....	588
	Returning without Committing .....	589
	Methods Not Allowed in Bean-Managed Transactions .....	589
	Transaction Timeouts .....	589
	Updating Multiple Databases .....	590
	Transactions in Web Components .....	591
<b>32</b>	<b>Resource Connections</b> .....	593
	Resources and JNDI Naming .....	593
	DataSource Objects and Connection Pools .....	594
	Resource Injection .....	595
	Field-Based Injection .....	596
	Method-Based Injection .....	597
	Class-Based Injection .....	598

Resource Adapters .....	598
Resource Adapter Contracts .....	599
Metadata Annotations .....	603
Replacing Deployment Descriptors With Metadata Annotations .....	604
Example 1: @Connector Annotation .....	604
Example 2: @ConnectionFactory Annotation .....	606
Example 3: @Activation Annotation .....	607
Common Client Interface .....	608
Further Information about Resources .....	609
<b>33 Java Message Service Concepts .....</b>	<b>611</b>
Overview of the JMS API .....	611
What Is Messaging? .....	611
What Is the JMS API? .....	612
When Can You Use the JMS API? .....	612
How Does the JMS API Work with the Java EE Platform? .....	613
Basic JMS API Concepts .....	614
JMS API Architecture .....	615
Messaging Domains .....	616
Message Consumption .....	618
The JMS API Programming Model .....	618
JMS Administered Objects .....	619
JMS Connections .....	621
JMS Sessions .....	621
JMS Message Producers .....	622
JMS Message Consumers .....	623
JMS Messages .....	625
JMS Queue Browsers .....	627
JMS Exception Handling .....	628
Creating Robust JMS Applications .....	628
Using Basic Reliability Mechanisms .....	629
Using Advanced Reliability Mechanisms .....	633
Using the JMS API in Java EE Applications .....	638
Using @Resource Annotations in Enterprise Bean or Web Components .....	638
Using Session Beans to Produce and to Synchronously Receive Messages .....	639

Using Message-Driven Beans to Receive Messages Asynchronously .....	640
Managing Distributed Transactions .....	642
Using the JMS API with Application Clients and Web Components .....	644
Further Information about JMS .....	644
<b>34 Java Message Service Examples .....</b>	<b>645</b>
Writing Simple JMS Applications .....	646
A Simple Example of Synchronous Message Receives .....	646
A Simple Example of Asynchronous Message Consumption .....	655
A Simple Example of Browsing Messages in a Queue .....	660
Running JMS Clients on Multiple Systems .....	664
Writing Robust JMS Applications .....	670
A Message Acknowledgment Example .....	670
A Durable Subscription Example .....	672
A Local Transaction Example .....	674
An Application That Uses the JMS API with a Session Bean .....	680
Writing the Application Components for the <code>clientsessionmdb</code> Example .....	680
Creating Resources for the <code>clientsessionmdb</code> Example .....	683
Building, Deploying, and Running the <code>clientsessionmdb</code> Example Using NetBeans IDE .....	683
Building, Deploying, and Running the <code>clientsessionmdb</code> Example Using Ant .....	685
An Application That Uses the JMS API with an Entity .....	686
Overview of the <code>clientmdbentity</code> Example Application .....	686
Writing the Application Components for the <code>clientmdbentity</code> Example .....	688
Creating Resources for the <code>clientmdbentity</code> Example .....	690
Building, Deploying, and Running the <code>clientmdbentity</code> Example Using NetBeans IDE .....	691
Building, Deploying, and Running the <code>clientmdbentity</code> Example Using Ant .....	693
An Application Example That Consumes Messages from a Remote Server .....	695
Overview of the <code>consumerremote</code> Example Modules .....	695
Writing the Module Components for the <code>consumerremote</code> Example .....	696
Creating Resources for the <code>consumerremote</code> Example .....	697
Using Two Application Servers for the <code>consumerremote</code> Example .....	697
Building, Deploying, and Running the <code>consumerremoteModules</code> Using NetBeans IDE ...	697
Building, Deploying, and Running the <code>consumerremote</code> Modules Using Ant .....	699
An Application Example That Deploys a Message-Driven Bean on Two Servers .....	701

---

Overview of the sendremote Example Modules .....	702
Writing the Module Components for the sendremote Example .....	703
Creating Resources for the sendremote Example .....	704
Using Two Application Servers for the sendremote Example .....	705
Building, Deploying, and Running the sendremote Modules Using NetBeans IDE .....	705
Building, Deploying, and Running the sendremote Modules Using Ant .....	708
<b>Index</b> .....	<b>711</b>



# Preface

---

This tutorial is a guide to developing enterprise applications for the Java™ Platform, Enterprise Edition 6 (Java EE 6).

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

Enterprise Server v3 is developed through the GlassFish project open-source community at <https://glassfish.dev.java.net/>. The GlassFish project provides a structured process for developing the Enterprise Server platform that makes the new features of the Java EE platform available faster, while maintaining the most important feature of Java EE: compatibility. It enables Java developers to access the Enterprise Server source code and to contribute to the development of the Enterprise Server. The GlassFish project is designed to encourage communication between Sun engineers and the community.

The following topics are addressed here:

- “Before You Read This Book” on page 23
- “Enterprise Server Documentation Set” on page 24
- “Related Documentation” on page 25
- “Symbol Conventions” on page 26
- “Typographic Conventions” on page 27
- “Default Paths and File Names” on page 27
- “Documentation, Support, and Training” on page 28
- “Searching Sun Product Documentation” on page 28
- “Third-Party Web Site References” on page 29
- “Sun Welcomes Your Comments” on page 29

## Before You Read This Book

Before proceeding with this tutorial, you should have a good knowledge of the Java programming language. A good way to get to that point is to work through *The Java Tutorial, Fourth Edition*, Sharon Zakhour et al. (Addison-Wesley, 2006). You should also be familiar with the Java DataBase Connectivity (JDBC™) and relational database features described in *JDBC API Tutorial and Reference, Third Edition*, Maydene Fisher et al. (Addison-Wesley, 2003).

# Enterprise Server Documentation Set

The Enterprise Server documentation set describes deployment planning and system installation. The Uniform Resource Locator (URL) for Enterprise Server documentation is <http://docs.sun.com/coll/1343.9>. For an introduction to Enterprise Server, refer to the books in the order in which they are listed in the following table.

TABLE P-1 Books in the Enterprise Server Documentation Set

Book Title	Description
<i>Release Notes</i>	Provides late-breaking information about the software and the documentation. Includes a comprehensive, table-based summary of the supported hardware, operating system, Java Development Kit (JDK™), and database drivers.
<i>Quick Start Guide</i>	Explains how to get started with the Enterprise Server product.
<i>Installation Guide</i>	Explains how to install the software and its components.
<i>Upgrade Guide</i>	Explains how to upgrade to the latest version of Enterprise Server. This guide also describes differences between adjacent product releases and configuration options that can result in incompatibility with the product specifications.
<i>Administration Guide</i>	Explains how to configure, monitor, and manage Enterprise Server subsystems and components from the command line by using the <code>asadmin(1M)</code> utility. Instructions for performing these tasks from the Administration Console are provided in the Administration Console online help.
<i>Application Deployment Guide</i>	Explains how to assemble and deploy applications to the Enterprise Server and provides information about deployment descriptors.
<i>Your First Cup: An Introduction to the Java EE Platform</i>	Provides a short tutorial for beginning Java EE programmers that explains the entire process for developing a simple enterprise application. The sample application is a web application that consists of a component that is based on the Enterprise JavaBeans™ specification, a JAX-RS web service, and a JavaServer™ Faces component for the web front end.
<i>Application Development Guide</i>	Explains how to create and implement Java Platform, Enterprise Edition (Java EE platform) applications that are intended to run on the Enterprise Server. These applications follow the open Java standards model for Java EE components and APIs. This guide provides information about developer tools, security, and debugging.
<i>Add-On Component Development Guide</i>	Explains how to use published interfaces of Enterprise Server to develop add-on components for Enterprise Server. This document explains how to perform <i>only</i> those tasks that ensure that the add-on component is suitable for Enterprise Server.

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

Book Title	Description
<i>Embedded Server Guide</i>	Explains how to run applications in embedded Enterprise Server and to develop applications in which Enterprise Server is embedded.
<i>Scripting Framework Guide</i>	Explains how to develop scripting applications in languages such as Ruby on Rails and Groovy on Grails for deployment to Enterprise Server.
<i>Troubleshooting Guide</i>	Describes common problems that you might encounter when using Enterprise Server and how to solve them.
<i>Error Message Reference</i>	Describes error messages that you might encounter when using Enterprise Server.
<i>Reference Manual</i>	Provides reference information in man page format for Enterprise Server administration commands, utility commands, and related concepts.
<i>Domain File Format Reference</i>	Describes the format of the Enterprise Server configuration file, <code>domain.xml</code> .
<i>Java EE 6 Tutorial, Volume I</i>	Explains how to use Java EE 6 platform technologies and APIs to develop Java EE applications.
<i>Message Queue Release Notes</i>	Describes new features, compatibility issues, and existing bugs for Sun GlassFish Message Queue.
<i>Message Queue Administration Guide</i>	Explains how to set up and manage a Sun GlassFish Message Queue messaging system.
<i>Message Queue Developer's Guide for JMX Clients</i>	Describes the application programming interface in Sun GlassFish Message Queue for programmatically configuring and monitoring Message Queue resources in conformance with the Java Management Extensions (JMX).
<i>System Virtualization Support in Sun Java System Products</i>	Summarizes Sun support for Sun Java System products when used in conjunction with system virtualization products and features.

## Related Documentation

*The Java EE 6 Tutorial, Volume II* ([https://www.sun.com/offers/details/java\\_ee6\\_tutorial.xml](https://www.sun.com/offers/details/java_ee6_tutorial.xml)) contains all the topics in *Java EE 6 Tutorial, Volume I* and adds advanced topics, additional technologies, and case studies. The document is available to registered users of Enterprise Server.

Javadoc™ tool reference documentation for packages that are provided with Enterprise Server is available as follows:

- The API specification for version 6 of Java EE is located at <http://java.sun.com/javase/6/docs/api/>.
- API documentation for packages that are specific to the Enterprise Server product is located at: <http://javadoc.glassfish.org/v3/apidoc/>.

Additionally, the following resources might be useful:

- The Java EE Specifications (<http://java.sun.com/javae/technologies/index.jsp>)
- The Java EE Blueprints (<http://java.sun.com/reference/blueprints/index.html>)

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

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

The sample applications demonstrate a broad range of Java EE technologies. The samples are bundled with the Java EE Software Development Kit (SDK).

## Symbol Conventions

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

TABLE P-2 Symbol Conventions

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

## Typographic Conventions

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

TABLE P-3 Typographic Conventions

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

## Default Paths and File Names

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

TABLE P-4 Default Paths and File Names

Placeholder	Description	Default Value
<i>as-install</i>	Represents the base installation directory for the Enterprise Server or the Software Development Kit (SDK) of which the Enterprise Server is a part.	Installations on the Solaris™ operating system, Linux operating system, and Mac operating system: <i>user's-home-directory/glassfishv3/glassfish</i> Windows, all installations: <i>SystemDrive:\glassfishv3\glassfish</i>
<i>as-install-parent</i>	Represents the parent of the base installation directory for Enterprise Server.	Installations on the Solaris™ operating system, Linux operating system, and Mac operating system: <i>user's-home-directory/glassfishv3</i> Windows, all installations: <i>SystemDrive:\glassfishv3</i>

TABLE P-4 Default Paths and File Names (Continued)

Placeholder	Description	Default Value
<i>tut-install</i>	Represents the base installation directory for the Java EE Tutorial after you install the Enterprise Server or the SDK and run the Update Tool.	<i>as-install/docs/javaee-tutorial</i>
<i>domain-root-dir</i>	Represents the directory in which a domain is created by default.	<i>as-install/domains/</i>
<i>domain-dir</i>	Represents the directory in which a domain's configuration is stored.  In configuration files, <i>domain-dir</i> is represented as follows:  \${com.sun.aas.instanceRoot}	<i>domain-root-dir/domain-name</i>

## Documentation, Support, and Training

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

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

## Searching Sun Product Documentation

Besides searching Sun product documentation from the docs.sun.com<sup>SM</sup> web site, you can use a search engine by typing the following syntax in the search field:

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

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

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

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

## Third-Party Web Site References

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

---

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

---

## Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click Send Comments. In the online form, provide the full document title and part number. The part number is a 7-digit or 9-digit number that can be found on the book's title page or in the document's URL. For example, the part number of this book is 820-7627.





**PART I**

# Introduction

Part One introduces the platform, the tutorial, and the examples.



# Overview

---

Developers today increasingly recognize the need for distributed, transactional, and portable applications that leverage the speed, security, and reliability of server-side technology. In the world of information technology, enterprise applications must be designed, built, and produced for less money, with greater speed, and with fewer resources.

With the Java™ Platform, Enterprise Edition (Java EE), development of Java enterprise applications has never been easier or faster. The aim of the Java EE platform is to provide developers with a powerful set of APIs while reducing development time, reducing application complexity, and improving application performance.

The Java EE platform uses a simplified programming model. XML deployment descriptors are optional. Instead, a developer can simply enter the information as an annotation directly into a Java source file, and the Java EE server will configure the component at deployment and runtime. These annotations are generally used to embed in a program data that would otherwise be furnished in a deployment descriptor. With annotations, the specification information is put directly in your code next to the program element that it affects.

In the Java EE platform, dependency injection can be applied to all resources that a component needs, effectively hiding the creation and lookup of resources from application code. Dependency injection can be used in EJB containers, web containers, and application clients. Dependency injection allows the Java EE container to automatically insert references to other required components or resources using annotations.

This tutorial uses examples to describe the features and functionalities available in the Java EE platform for developing enterprise applications. Whether you are a new or experienced Enterprise developer, you should find the examples and accompanying text a valuable and accessible knowledge base for creating your own solutions.

If you are new to Java EE enterprise application development, this chapter is a good place to start. Here you will review development basics, learn about the Java EE architecture and APIs, become acquainted with important terms and concepts, and find out how to approach Java EE application programming, assembly, and deployment.

The following topics are addressed here:

- [“Java EE 6 Highlights” on page 34](#)
- [“Java EE Application Model” on page 34](#)
- [“Distributed Multitiered Applications” on page 35](#)
- [“Java EE Containers” on page 41](#)
- [“Web Services Support” on page 44](#)
- [“Java EE Application Assembly and Deployment” on page 45](#)
- [“Packaging Applications” on page 46](#)
- [“Development Roles” on page 47](#)
- [“Java EE 6 APIs” on page 50](#)
- [“Java EE 6 APIs Included in the Java Platform, Standard Edition 6.0 \(Java SE 6\)” on page 55](#)
- [“Sun GlassFish Enterprise Server v3” on page 57](#)

## Java EE 6 Highlights

The Java EE 6 platform includes the following new features:

- Profiles, configurations of the Java EE platform targeted at specific classes of applications. Specifically, the Java EE 6 platform introduces a Web Profile targeted at web applications, as well as a Full Profile that contains all Java EE technologies.
- New technologies, including the following:
  - Java API for RESTful Web Services (JAX-RS)
  - Contexts and Dependency Injection for the Java EE Platform (JSR-299), informally known as Web Beans
  - Java Authentication Service Provider Interface for Containers (JASPIC)
- New features for Enterprise JavaBeans™ (EJB™) components (see [“Enterprise JavaBeans Technology” on page 50](#) for details)
- New features for servlets (see [“Java Servlet Technology” on page 50](#) for details)
- New features for JavaServer™ Faces components (see [“JavaServer Faces Technology” on page 51](#) for details)

## Java EE Application Model

The Java EE application model begins with the Java programming language and the Java virtual machine. The proven portability, security, and developer productivity they provide forms the basis of the application model. Java EE is designed to support applications that implement enterprise services for customers, employees, suppliers, partners, and others who make demands on or contributions to the enterprise. Such applications are inherently complex, potentially accessing data from a variety of sources and distributing applications to a variety of clients.

To better control and manage these applications, the business functions to support these various users are conducted in the middle tier. The middle tier represents an environment that is closely controlled by an enterprise's information technology department. The middle tier is typically run on dedicated server hardware and has access to the full services of the enterprise.

The Java EE application model defines an architecture for implementing services as multitier applications that deliver the scalability, accessibility, and manageability needed by enterprise-level applications. This model partitions the work needed to implement a multitier service into two parts: the business and presentation logic to be implemented by the developer, and the standard system services provided by the Java EE platform. The developer can rely on the platform to provide solutions for the hard systems-level problems of developing a multitier service.

## Distributed Multitiered Applications

The Java EE platform uses a distributed multitiered application model for enterprise applications. Application logic is divided into components according to function, and the various application components that make up a Java EE application are installed on different machines depending on the tier in the multitiered Java EE environment to which the application component belongs.

[Figure 1–1](#) shows two multitiered Java EE applications divided into the tiers described in the following list. The Java EE application parts shown in [Figure 1–1](#) are presented in “[Java EE Components](#)” on [page 37](#).

- Client-tier components run on the client machine.
- Web-tier components run on the Java EE server.
- Business-tier components run on the Java EE server.
- Enterprise information system (EIS)-tier software runs on the EIS server.

Although a Java EE application can consist of the three or four tiers shown in [Figure 1–1](#), Java EE multitiered applications are generally considered to be three-tiered applications because they are distributed over three locations: client machines, the Java EE server machine, and the database or legacy machines at the back end. Three-tiered applications that run in this way extend the standard two-tiered client and server model by placing a multithreaded application server between the client application and back-end storage.

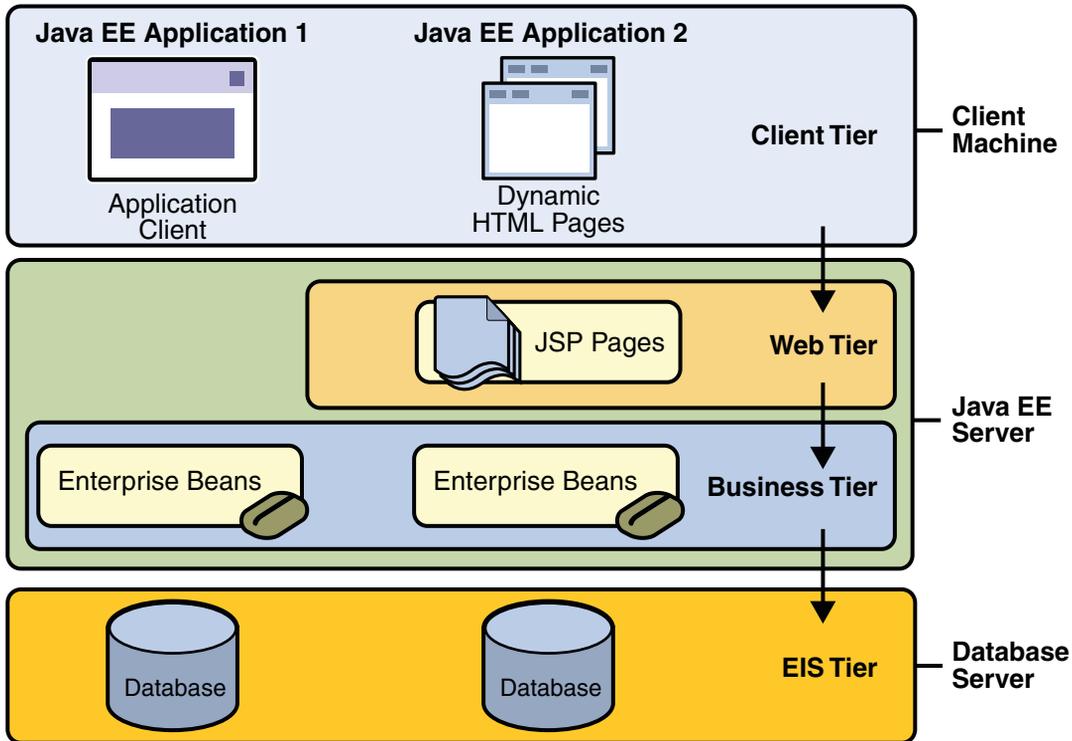


FIGURE 1-1 Multitiered Applications

## Security

While other enterprise application models require platform-specific security measures in each application, the Java EE security environment enables security constraints to be defined at deployment time. The Java EE platform makes applications portable to a wide variety of security implementations by shielding application developers from the complexity of implementing security features.

The Java EE platform provides standard declarative access control rules that are defined by the developer and interpreted when the application is deployed on the server. Java EE also provides standard login mechanisms so application developers do not have to implement these mechanisms in their applications. The same application works in a variety of different security environments without changing the source code.

## Java EE Components

Java EE applications are made up of components. A *Java EE component* is a self-contained functional software unit that is assembled into a Java EE application with its related classes and files and that communicates with other components.

The Java EE specification defines the following Java EE components:

- Application clients and applets are components that run on the client.
- Java Servlet, JavaServer Faces, and JavaServer Pages™ (JSP™) technology components are web components that run on the server.
- Enterprise JavaBeans (EJB) components (enterprise beans) are business components that run on the server.

Java EE components are written in the Java programming language and are compiled in the same way as any program in the language. The difference between Java EE components and “standard” Java classes is that Java EE components are assembled into a Java EE application, are verified to be well formed and in compliance with the Java EE specification, and are deployed to production, where they are run and managed by the Java EE server.

## Java EE Clients

A Java EE client can be a web client or an application client.

### Web Clients

A *web client* consists of two parts: (1) dynamic web pages containing various types of markup language (HTML, XML, and so on), which are generated by web components running in the web tier, and (2) a web browser, which renders the pages received from the server.

A web client is sometimes called a *thin client*. Thin clients usually do not query databases, execute complex business rules, or connect to legacy applications. When you use a thin client, such heavyweight operations are off-loaded to enterprise beans executing on the Java EE server, where they can leverage the security, speed, services, and reliability of Java EE server-side technologies.

### Applets

A web page received from the web tier can include an embedded applet. An *applet* is a small client application written in the Java programming language that executes in the Java virtual machine installed in the web browser. However, client systems will likely need the Java Plug-in and possibly a security policy file for the applet to successfully execute in the web browser.

Web components are the preferred API for creating a web client program because no plug-ins or security policy files are needed on the client systems. Also, web components enable cleaner and more modular application design because they provide a way to separate applications programming from web page design. Personnel involved in web page design thus do not need to understand Java programming language syntax to do their jobs.

## Application Clients

An *application client* runs on a client machine and provides a way for users to handle tasks that require a richer user interface than can be provided by a markup language. It typically has a graphical user interface (GUI) created from the Swing or the Abstract Window Toolkit (AWT) API, but a command-line interface is certainly possible.

Application clients directly access enterprise beans running in the business tier. However, if application requirements warrant it, an application client can open an HTTP connection to establish communication with a servlet running in the web tier. Application clients written in languages other than Java can interact with Java EE servers, enabling the Java EE platform to interoperate with legacy systems, clients, and non-Java languages.

## The JavaBeans™ Component Architecture

The server and client tiers might also include components based on the JavaBeans component architecture (JavaBeans components) to manage the data flow between an application client or applet and components running on the Java EE server, or between server components and a database. JavaBeans components are not considered Java EE components by the Java EE specification.

JavaBeans components have properties and have `get` and `set` methods for accessing the properties. JavaBeans components used in this way are typically simple in design and implementation but should conform to the naming and design conventions outlined in the JavaBeans component architecture.

## Java EE Server Communications

Figure 1–2 shows the various elements that can make up the client tier. The client communicates with the business tier running on the Java EE server either directly or, as in the case of a client running in a browser, by going through web pages or servlets running in the web tier.

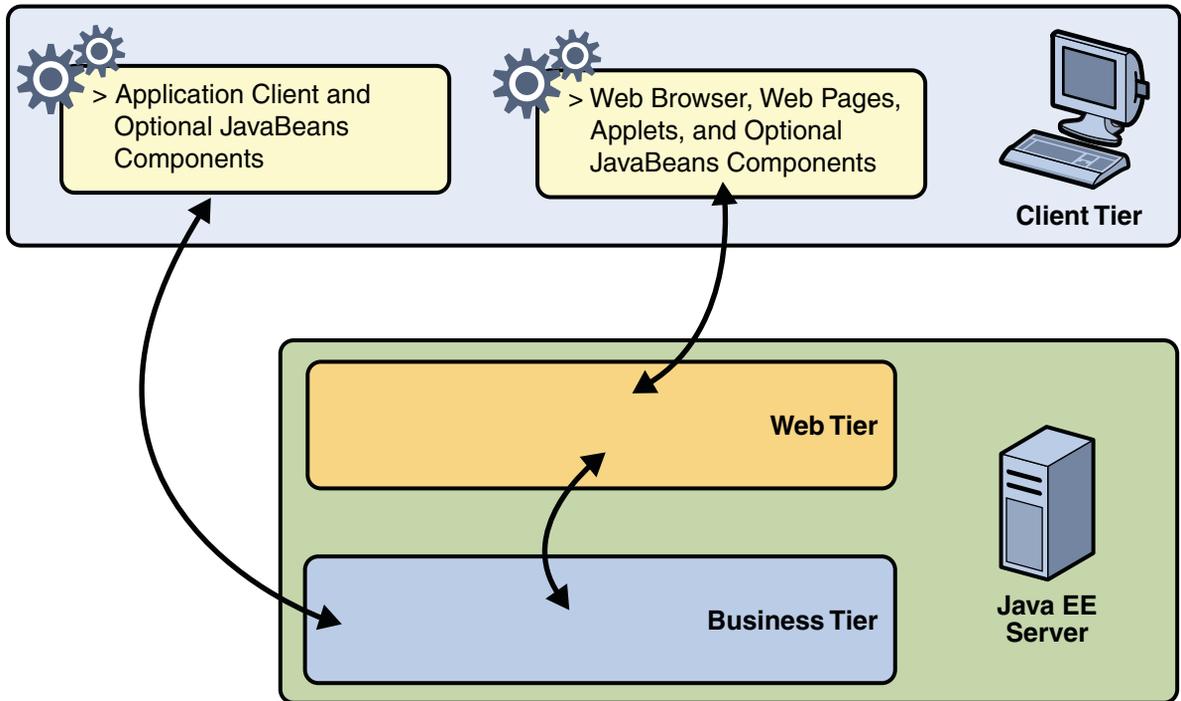


FIGURE 1-2 Server Communication

## Web Components

Java EE web components are either servlets or web pages created using JavaServer Faces technology and/or JSP technology (JSP pages). *Servlets* are Java programming language classes that dynamically process requests and construct responses. *JSP pages* are text-based documents that execute as servlets but allow a more natural approach to creating static content. *JavaServer Faces* technology builds on servlets and JSP technology and provides a user interface component framework for web applications.

Static HTML pages and applets are bundled with web components during application assembly but are not considered web components by the Java EE specification. Server-side utility classes can also be bundled with web components and, like HTML pages, are not considered web components.

As shown in [Figure 1-3](#), the web tier, like the client tier, might include a JavaBeans component to manage the user input and send that input to enterprise beans running in the business tier for processing.

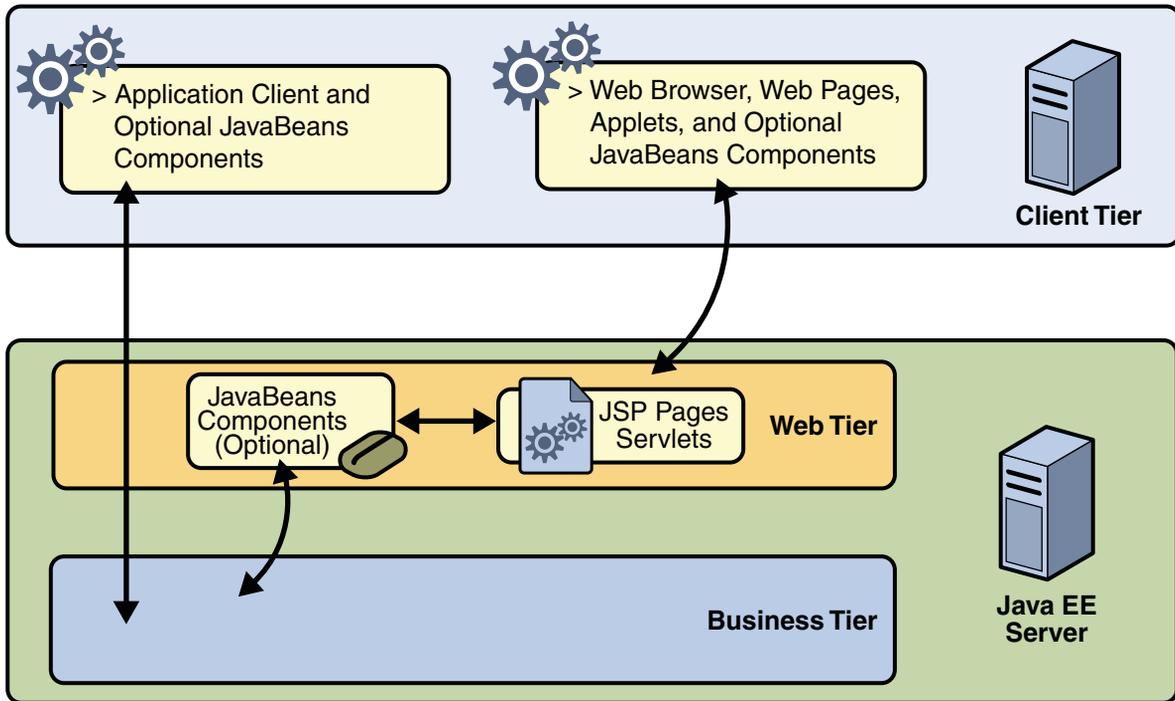


FIGURE 1-3 Web Tier and Java EE Applications

## Business Components

Business code, which is logic that solves or meets the needs of a particular business domain such as banking, retail, or finance, is handled by enterprise beans running in either the business tier or the web tier. Figure 1-4 shows how an enterprise bean receives data from client programs, processes it (if necessary), and sends it to the enterprise information system tier for storage. An enterprise bean also retrieves data from storage, processes it (if necessary), and sends it back to the client program.

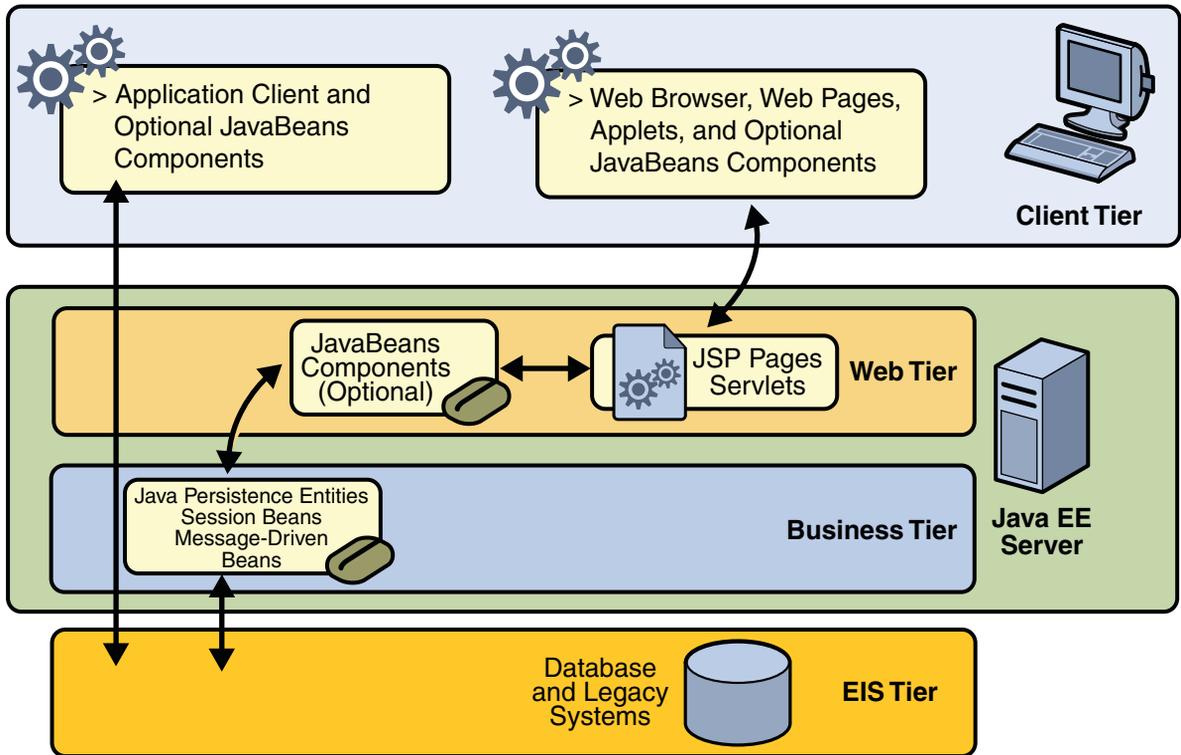


FIGURE 1-4 Business and EIS Tiers

## Enterprise Information System Tier

The enterprise information system tier handles EIS software and includes enterprise infrastructure systems such as enterprise resource planning (ERP), mainframe transaction processing, database systems, and other legacy information systems. For example, Java EE application components might need access to enterprise information systems for database connectivity.

## Java EE Containers

Normally, thin-client multitiered applications are hard to write because they involve many lines of intricate code to handle transaction and state management, multithreading, resource pooling, and other complex low-level details. The component-based and platform-independent Java EE architecture makes Java EE applications easy to write because business logic is organized into reusable components. In addition, the Java EE server provides underlying

services in the form of a container for every component type. Because you do not have to develop these services yourself, you are free to concentrate on solving the business problem at hand.

## Container Services

*Containers* are the interface between a component and the low-level platform-specific functionality that supports the component. Before a web, enterprise bean, or application client component can be executed, it must be assembled into a Java EE module and deployed into its container.

The assembly process involves specifying container settings for each component in the Java EE application and for the Java EE application itself. Container settings customize the underlying support provided by the Java EE server, including services such as security, transaction management, Java Naming and Directory Interface™ (JNDI) lookups, and remote connectivity. Here are some of the highlights:

- The Java EE security model lets you configure a web component or enterprise bean so that system resources are accessed only by authorized users.
- The Java EE transaction model lets you specify relationships among methods that make up a single transaction so that all methods in one transaction are treated as a single unit.
- JNDI lookup services provide a unified interface to multiple naming and directory services in the enterprise so that application components can access these services.
- The Java EE remote connectivity model manages low-level communications between clients and enterprise beans. After an enterprise bean is created, a client invokes methods on it as if it were in the same virtual machine.

Because the Java EE architecture provides configurable services, application components within the same Java EE application can behave differently based on where they are deployed. For example, an enterprise bean can have security settings that allow it a certain level of access to database data in one production environment and another level of database access in another production environment.

The container also manages nonconfigurable services such as enterprise bean and servlet life cycles, database connection resource pooling, data persistence, and access to the Java EE platform APIs (see “[Java EE 6 APIs](#)” on page 50).

## Container Types

The deployment process installs Java EE application components in the Java EE containers as illustrated in [Figure 1-5](#).

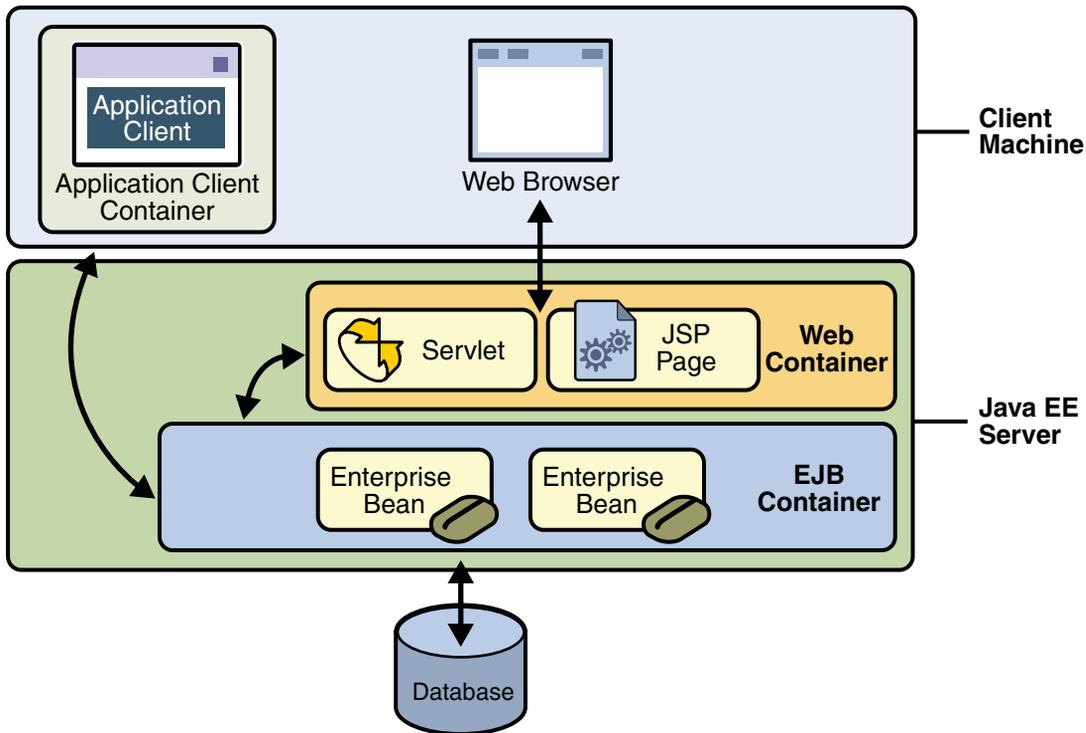


FIGURE 1-5 Java EE Server and Containers

- **Java EE server:** The runtime portion of a Java EE product. A Java EE server provides EJB and web containers.
- **Enterprise JavaBeans (EJB) container:** Manages the execution of enterprise beans for Java EE applications. Enterprise beans and their container run on the Java EE server.
- **Web container:** Manages the execution of web pages, servlets, and some EJB components for Java EE applications. Web components and their container run on the Java EE server.
- **Application client container:** Manages the execution of application client components. Application clients and their container run on the client.
- **Applet container:** Manages the execution of applets. Consists of a web browser and Java Plug-in running on the client together.

# Web Services Support

Web services are web-based enterprise applications that use open, XML-based standards and transport protocols to exchange data with calling clients. The Java EE platform provides the XML APIs and tools you need to quickly design, develop, test, and deploy web services and clients that fully interoperate with other web services and clients running on Java-based or non-Java-based platforms.

To write web services and clients with the Java EE XML APIs, all you do is pass parameter data to the method calls and process the data returned; or for document-oriented web services, you send documents containing the service data back and forth. No low-level programming is needed because the XML API implementations do the work of translating the application data to and from an XML-based data stream that is sent over the standardized XML-based transport protocols. These XML-based standards and protocols are introduced in the following sections.

The translation of data to a standardized XML-based data stream is what makes web services and clients written with the Java EE XML APIs fully interoperable. This does not necessarily mean that the data being transported includes XML tags because the transported data can itself be plain text, XML data, or any kind of binary data such as audio, video, maps, program files, computer-aided design (CAD) documents and the like. The next section introduces XML and explains how parties doing business can use XML tags and schemas to exchange data in a meaningful way.

## XML

XML is a cross-platform, extensible, text-based standard for representing data. When XML data is exchanged between parties, the parties are free to create their own tags to describe the data, set up schemas to specify which tags can be used in a particular kind of XML document, and use XML stylesheets to manage the display and handling of the data.

For example, a web service can use XML and a schema to produce price lists, and companies that receive the price lists and schema can have their own stylesheets to handle the data in a way that best suits their needs. Here are examples:

- One company might put XML pricing information through a program to translate the XML to HTML so that it can post the price lists to its intranet.
- A partner company might put the XML pricing information through a tool to create a marketing presentation.
- Another company might read the XML pricing information into an application for processing.

## SOAP Transport Protocol

Client requests and web service responses are transmitted as Simple Object Access Protocol (SOAP) messages over HTTP to enable a completely interoperable exchange between clients and web services, all running on different platforms and at various locations on the Internet. HTTP is a familiar request-and-response standard for sending messages over the Internet, and SOAP is an XML-based protocol that follows the HTTP request-and-response model.

The SOAP portion of a transported message handles the following:

- Defines an XML-based envelope to describe what is in the message and how to process the message
- Includes XML-based encoding rules to express instances of application-defined data types within the message
- Defines an XML-based convention for representing the request to the remote service and the resulting response

## WSDL Standard Format

The Web Services Description Language (WSDL) is a standardized XML format for describing network services. The description includes the name of the service, the location of the service, and ways to communicate with the service. WSDL service descriptions can be published on the Web. The Sun GlassFish™ Enterprise Server provides a tool for generating the WSDL specification of a web service that uses remote procedure calls to communicate with clients.

# Java EE Application Assembly and Deployment

A Java EE application is packaged into one or more standard units for deployment to any Java EE platform-compliant system. Each unit contains:

- A functional component or components (such as an enterprise bean, web page, servlet, or applet)
- An optional deployment descriptor that describes its content

Once a Java EE unit has been produced, it is ready to be deployed. Deployment typically involves using a platform's deployment tool to specify location-specific information, such as a list of local users that can access it and the name of the local database. Once deployed on a local platform, the application is ready to run.

## Packaging Applications

A Java EE application is delivered in either a Java Archive (JAR) file, a Web Archive (WAR) file, or an Enterprise Archive (EAR) file. A WAR or EAR file is a standard JAR (.jar) file with a .war or .ear extension. Using JAR, WAR, and EAR files and modules makes it possible to assemble a number of different Java EE applications using some of the same components. No extra coding is needed; it is only a matter of assembling (or packaging) various Java EE modules into Java EE JAR, WAR, or EAR files.

An EAR file (see [Figure 1–6](#)) contains Java EE modules and, optionally, deployment descriptors. A *deployment descriptor* is an XML document with an .xml extension that describes the deployment settings of an application, a module, or a component. Because deployment descriptor information is declarative, it can be changed without the need to modify the source code. At runtime, the Java EE server reads the deployment descriptor and acts upon the application, module, or component accordingly.

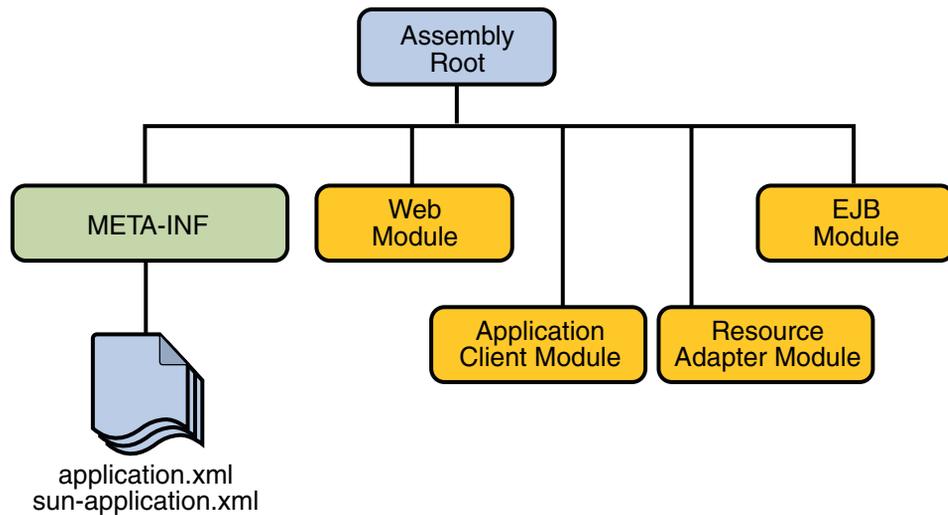


FIGURE 1–6 EAR File Structure

There are two types of deployment descriptors: Java EE and runtime. A *Java EE deployment descriptor* is defined by a Java EE specification and can be used to configure deployment settings on any Java EE-compliant implementation. A *runtime deployment descriptor* is used to configure Java EE implementation-specific parameters. For example, the Sun GlassFish Enterprise Server runtime deployment descriptor contains information such as the context root of a web application, and Enterprise Server implementation-specific parameters, such as

caching directives. The Enterprise Server runtime deployment descriptors are named `sun-moduleType.xml` and are located in the same META-INF directory as the Java EE deployment descriptor.

A *Java EE module* consists of one or more Java EE components for the same container type and, optionally, one component deployment descriptor of that type. An enterprise bean module deployment descriptor, for example, declares transaction attributes and security authorizations for an enterprise bean. A Java EE module can be deployed as a *stand-alone* module.

The four types of Java EE modules are as follows:

- EJB modules, which contain class files for enterprise beans and an EJB deployment descriptor. EJB modules are packaged as JAR files with a `.jar` extension.
- Web modules, which contain servlet class files, web files, supporting class files, GIF and HTML files, and a web application deployment descriptor. Web modules are packaged as JAR files with a `.war` (Web ARchive) extension.
- Application client modules, which contain class files and an application client deployment descriptor. Application client modules are packaged as JAR files with a `.jar` extension.
- Resource adapter modules, which contain all Java interfaces, classes, native libraries, and other documentation, along with the resource adapter deployment descriptor. Together, these implement the Connector architecture (see [“Java EE Connector Architecture” on page 53](#)) for a particular EIS. Resource adapter modules are packaged as JAR files with an `.rar` (resource adapter archive) extension.

## Development Roles

Reusable modules make it possible to divide the application development and deployment process into distinct roles so that different people or companies can perform different parts of the process.

The first two roles involve purchasing and installing the Java EE product and tools. After software is purchased and installed, Java EE components can be developed by application component providers, assembled by application assemblers, and deployed by application deployers. In a large organization, each of these roles might be executed by different individuals or teams. This division of labor works because each of the earlier roles outputs a portable file that is the input for a subsequent role. For example, in the application component development phase, an enterprise bean software developer delivers EJB JAR files. In the application assembly role, another developer may combine these EJB JAR files into a Java EE application and save it in an EAR file. In the application deployment role, a system administrator at the customer site uses the EAR file to install the Java EE application into a Java EE server.

The different roles are not always executed by different people. If you work for a small company, for example, or if you are prototyping a sample application, you might perform the tasks in every phase.

## Java EE Product Provider

The Java EE product provider is the company that designs and makes available for purchase the Java EE platform APIs, and other features defined in the Java EE specification. Product providers are typically application server vendors who implement the Java EE platform according to the Java EE 6 Platform specification.

## Tool Provider

The tool provider is the company or person who creates development, assembly, and packaging tools used by component providers, assemblers, and deployers.

## Application Component Provider

The application component provider is the company or person who creates web components, enterprise beans, applets, or application clients for use in Java EE applications.

## Enterprise Bean Developer

An enterprise bean developer performs the following tasks to deliver an EJB JAR file that contains one or more enterprise beans:

- Writes and compiles the source code
- Specifies the deployment descriptor (optional)
- Packages the `.class` files and deployment descriptor into the EJB JAR file

## Web Component Developer

A web component developer performs the following tasks to deliver a WAR file containing one or more web components:

- Writes and compiles servlet source code
- Writes JavaServer Faces, JSP, and HTML files
- Specifies the deployment descriptor (optional)
- Packages the `.class`, `.jsp`, and `.html` files and deployment descriptor into the WAR file

## Application Client Developer

An application client developer performs the following tasks to deliver a JAR file containing the application client:

- Writes and compiles the source code
- Specifies the deployment descriptor for the client (optional)

- Packages the `.class` files and deployment descriptor into the JAR file

## Application Assembler

The application assembler is the company or person who receives application modules from component providers and may assemble them into a Java EE application EAR file. The assembler or deployer can edit the deployment descriptor directly or can use tools that correctly add XML tags according to interactive selections.

A software developer performs the following tasks to deliver an EAR file containing the Java EE application:

- Assembles EJB JAR and WAR files created in the previous phases into a Java EE application (EAR) file
- Specifies the deployment descriptor for the Java EE application (optional)
- Verifies that the contents of the EAR file are well formed and comply with the Java EE specification

## Application Deployer and Administrator

The application deployer and administrator is the company or person who configures and deploys the Java EE application, administers the computing and networking infrastructure where Java EE applications run, and oversees the runtime environment. Duties include such things as setting transaction controls and security attributes and specifying connections to databases.

During configuration, the deployer follows instructions supplied by the application component provider to resolve external dependencies, specify security settings, and assign transaction attributes. During installation, the deployer moves the application components to the server and generates the container-specific classes and interfaces.

A deployer or system administrator performs the following tasks to install and configure a Java EE application:

- Configures the Java EE application for the operational environment
- Verifies that the contents of the EAR file are well formed and comply with the Java EE specification
- Deploys (installs) the Java EE application EAR file into the Java EE server

## Java EE 6 APIs

The following sections give a brief summary of the technologies required by the Java EE platform, and the APIs used in Java EE applications.

### Enterprise JavaBeans Technology

An Enterprise JavaBeans (EJB) component, or *enterprise bean*, is a body of code having fields and methods to implement modules of business logic. You can think of an enterprise bean as a building block that can be used alone or with other enterprise beans to execute business logic on the Java EE server.

There are two kinds of enterprise beans: session beans and message-driven beans.

A *session bean* represents a transient conversation with a client. When the client finishes executing, the session bean and its data are gone.

A *message-driven bean* combines features of a session bean and a message listener, allowing a business component to receive messages asynchronously. Commonly, these are Java Message Service (JMS) messages.

In the Java EE 6 platform, new enterprise bean features include the following:

- The ability to package local enterprise beans in a WAR file
- Singleton session beans, which provide easy access to shared state
- A lightweight subset of Enterprise JavaBeans functionality that can be provided within Java EE Profiles such as the Java EE Web Profile.

### Java Servlet Technology

Java Servlet technology lets you define HTTP-specific servlet classes. A servlet class extends the capabilities of servers that host applications that are accessed by way of a request-response programming model. Although servlets can respond to any type of request, they are commonly used to extend the applications hosted by web servers.

In the Java EE 6 platform, new Java Servlet technology features include the following:

- Annotation support
- Asynchronous support
- Ease of configuration
- Enhancements to existing APIs
- Pluggability

## JavaServer Faces Technology

JavaServer Faces technology is a user interface framework for building web applications. The main components of JavaServer Faces technology are as follows:

- A GUI component framework.
- A flexible model for rendering components in different kinds of HTML or different markup languages and technologies. A `Renderer` object generates the markup to render the component and converts the data stored in a model object to types that can be represented in a view.
- A standard `RenderKit` for generating HTML/4.01 markup.

The following features support the GUI components:

- Input validation
- Event handling
- Data conversion between model objects and components
- Managed model object creation
- Page navigation configuration

All this functionality is available using standard Java APIs and XML-based configuration files.

In the Java EE 6 platform, new features of JavaServer Faces include the following:

- The ability to use annotations instead of a configuration file to specify managed beans
- Facelets, a display technology that replaces JavaServer Pages (JSP) technology using XHTML files
- Ajax support
- Composite components
- Implicit navigation

## JavaServer Pages™ Technology

JavaServer Pages (JSP) technology lets you put snippets of servlet code directly into a text-based document. A JSP page is a text-based document that contains two types of text: static data (which can be expressed in any text-based format such as HTML, WML, and XML) and JSP elements, which determine how the page constructs dynamic content.

## JavaServer Pages Standard Tag Library

The JavaServer Pages Standard Tag Library (JSTL) encapsulates core functionality common to many JSP applications. Instead of mixing tags from numerous vendors in your JSP applications,

you employ a single, standard set of tags. This standardization allows you to deploy your applications on any JSP container that supports JSTL and makes it more likely that the implementation of the tags is optimized.

JSTL has iterator and conditional tags for handling flow control, tags for manipulating XML documents, internationalization tags, tags for accessing databases using SQL, and commonly used functions.

## Java Persistence API

The Java Persistence API is a Java standards-based solution for persistence. Persistence uses an object-relational mapping approach to bridge the gap between an object oriented model and a relational database. The Java Persistence API can also be used in Java SE applications, outside of the Java EE environment. Java Persistence consists of three areas:

- The Java Persistence API
- The query language
- Object/relational mapping metadata

## Java Transaction API

The Java Transaction API (JTA) provides a standard interface for demarcating transactions. The Java EE architecture provides a default auto commit to handle transaction commits and rollbacks. An *auto commit* means that any other applications that are viewing data will see the updated data after each database read or write operation. However, if your application performs two separate database access operations that depend on each other, you will want to use the JTA API to demarcate where the entire transaction, including both operations, begins, rolls back, and commits.

## Java API for RESTful Web Services (JAX-RS)

The Java API for RESTful Web Services (JAX-RS) defines APIs for the development of Web services built according to the Representational State Transfer (REST) architectural style. A JAX-RS application is a web application that consists of classes that are packaged as a servlet in a WAR file along with required libraries.

The JAX-RS API is new to the Java EE 6 platform.

## Java Message Service API

The Java Message Service (JMS) API is a messaging standard that allows Java EE application components to create, send, receive, and read messages. It enables distributed communication that is loosely coupled, reliable, and asynchronous.

## Java EE Connector Architecture

The Java EE Connector architecture is used by tools vendors and system integrators to create resource adapters that support access to enterprise information systems that can be plugged in to any Java EE product. A *resource adapter* is a software component that allows Java EE application components to access and interact with the underlying resource manager of the EIS. Because a resource adapter is specific to its resource manager, typically there is a different resource adapter for each type of database or enterprise information system.

The Java EE Connector architecture also provides a performance-oriented, secure, scalable, and message-based transactional integration of Java EE-based web services with existing EISs that can be either synchronous or asynchronous. Existing applications and EISs integrated through the Java EE Connector architecture into the Java EE platform can be exposed as XML-based web services by using JAX-WS and Java EE component models. Thus JAX-WS and the Java EE Connector architecture are complementary technologies for enterprise application integration (EAI) and end-to-end business integration.

## JavaMail™ API

Java EE applications use the JavaMail API to send email notifications. The JavaMail API has two parts: an application-level interface used by the application components to send mail, and a service provider interface. The Java EE platform includes the JavaMail API with a service provider that allows application components to send Internet mail.

## Java Authorization Service Provider Contract for Containers (Java ACC)

The Java ACC specification defines a contract between a Java EE application server and an authorization policy provider. All Java EE containers support this contract.

The Java ACC specification defines `java.security.Permission` classes that satisfy the Java EE authorization model. The specification defines the binding of container access decisions to operations on instances of these permission classes. It defines the semantics of policy providers that employ the new permission classes to address the authorization requirements of the Java EE platform, including the definition and use of roles.

## Java Authentication Service Provider Interface for Containers (JASPIC)

The Java Authentication Service Provider Interface for Containers (JASPIC) specification defines a service provider interface (SPI) by which authentication providers that implement message authentication mechanisms may be integrated in client or server message processing

containers or runtimes. Authentication providers integrated through this interface operate on network messages provided to them by their calling container. They transform outgoing messages so that the source of the message may be authenticated by the receiving container, and the recipient of the message may be authenticated by the message sender. They authenticate incoming messages and return to their calling container the identity established as a result of the message authentication.

## Java API for XML Registries

The Java API for XML Registries (JAXR) lets you access business and general-purpose registries over the web. JAXR supports the ebXML Registry and Repository standards and the emerging UDDI specifications. By using JAXR, developers can learn a single API and gain access to both of these important registry technologies.

Additionally, businesses can submit material to be shared and search for material that others have submitted. Standards groups have developed schemas for particular kinds of XML documents; two businesses might, for example, agree to use the schema for their industry's standard purchase order form. Because the schema is stored in a standard business registry, both parties can use JAXR to access it.

## Simplified Systems Integration

The Java EE platform is a platform-independent, full systems integration solution that creates an open marketplace in which every vendor can sell to every customer. Such a marketplace encourages vendors to compete, not by trying to lock customers into their technologies but instead by trying to outdo each other in providing products and services that benefit customers, such as better performance, better tools, or better customer support.

The Java EE 6 APIs enable systems and applications integration through the following:

- Unified application model across tiers with enterprise beans
- Simplified request-and-response mechanism with web pages and servlets
- Reliable security model with JAAS
- XML-based data interchange integration with JAXP, SAAJ, and JAX-WS
- Simplified interoperability with the Java EE Connector architecture
- Easy database connectivity with the Java DataBase Connectivity (JDBC™) API
- Enterprise application integration with message-driven beans and JMS, JTA, and JNDI

# Java EE 6 APIs Included in the Java Platform, Standard Edition 6.0 (Java SE 6)

Several APIs that are required by the Java EE 6 platform are included in the Java SE 6 platform and are thus available to Java EE applications.

## Java Database Connectivity API

The Java Database Connectivity (JDBC) API lets you invoke SQL commands from Java programming language methods. You use the JDBC API in an enterprise bean when you have a session bean access the database. You can also use the JDBC API from a servlet or a JSP page to access the database directly without going through an enterprise bean.

The JDBC API has two parts: an application-level interface used by the application components to access a database, and a service provider interface to attach a JDBC driver to the Java EE platform.

## Java Naming and Directory Interface™

The Java Naming and Directory Interface (JNDI) provides naming and directory functionality, enabling applications to access multiple naming and directory services, including existing naming and directory services such as LDAP, NDS, DNS, and NIS. It provides applications with methods for performing standard directory operations, such as associating attributes with objects and searching for objects using their attributes. Using JNDI, a Java EE application can store and retrieve any type of named Java object, allowing Java EE applications to coexist with many legacy applications and systems.

Java EE naming services provide application clients, enterprise beans, and web components with access to a JNDI naming environment. A *naming environment* allows a component to be customized without the need to access or change the component's source code. A container implements the component's environment and provides it to the component as a JNDI *naming context*.

A Java EE component can locate its environment naming context using JNDI interfaces. A component can create a `javax.naming.InitialContext` object and look up the environment naming context in `InitialContext` under the name `java:comp/env`. A component's naming environment is stored directly in the environment naming context or in any of its direct or indirect subcontexts.

A Java EE component can access named system-provided and user-defined objects. The names of system-provided objects, such as JTA `UserTransaction` objects, are stored in the environment naming context, `java:comp/env`. The Java EE platform allows a component to

name user-defined objects, such as enterprise beans, environment entries, JDBC `DataSource` objects, and message connections. An object should be named within a subcontext of the naming environment according to the type of the object. For example, enterprise beans are named within the subcontext `java:comp/env/ejb`, and JDBC `DataSource` references in the subcontext `java:comp/env/jdbc`.

## JavaBeans Activation Framework

The JavaBeans Activation Framework (JAF) is used by the JavaMail API. JAF provides standard services to determine the type of an arbitrary piece of data, encapsulate access to it, discover the operations available on it, and create the appropriate JavaBeans component to perform those operations.

## Java API for XML Processing

The Java API for XML Processing (JAXP), part of the Java SE platform, supports the processing of XML documents using Document Object Model (DOM), Simple API for XML (SAX), and Extensible Stylesheet Language Transformations (XSLT). JAXP enables applications to parse and transform XML documents independent of a particular XML processing implementation.

JAXP also provides namespace support, which lets you work with schemas that might otherwise have naming conflicts. Designed to be flexible, JAXP lets you use any XML-compliant parser or XSL processor from within your application and supports the W3C schema. You can find information on the W3C schema at this URL: <http://www.w3.org/XML/Schema>.

## Java Architecture for XML Binding (JAXB)

The Java Architecture for XML Binding (JAXB) provides a convenient way to bind an XML schema to a representation in Java language programs. JAXB can be used independently or in combination with JAX-WS, where it provides a standard data binding for web service messages. All Java EE application client containers, web containers, and EJB containers support the JAXB API.

## SOAP with Attachments API for Java

The SOAP with Attachments API for Java (SAAJ) is a low-level API on which JAX-WS and JAXR depend. SAAJ enables the production and consumption of messages that conform to the SOAP 1.1 and 1.2 specifications and SOAP with Attachments note. Most developers do not use the SAAJ API, instead using the higher-level JAX-WS API.

## Java API for XML Web Services (JAX-WS)

The JAX-WS specification provides support for web services that use the JAXB API for binding XML data to Java objects. The JAX-WS specification defines client APIs for accessing web services as well as techniques for implementing web service endpoints. The Implementing Enterprise Web Services specification describes the deployment of JAX-WS-based services and clients. The EJB and Java Servlet specifications also describe aspects of such deployment. It must be possible to deploy JAX-WS-based applications using any of these deployment models.

The JAX-WS specification describes the support for message handlers that can process message requests and responses. In general, these message handlers execute in the same container and with the same privileges and execution context as the JAX-WS client or endpoint component with which they are associated. These message handlers have access to the same JNDI `java:comp/env` namespace as their associated component. Custom serializers and deserializers, if supported, are treated in the same way as message handlers.

## Java Authentication and Authorization Service

The Java Authentication and Authorization Service (JAAS) provides a way for a Java EE application to authenticate and authorize a specific user or group of users to run it.

JAAS is a Java programming language version of the standard Pluggable Authentication Module (PAM) framework, which extends the Java Platform security architecture to support user-based authorization.

# Sun GlassFish Enterprise Server v3

Sun GlassFish Enterprise Server v3 is a compliant implementation of the Java EE 6 platform. In addition to supporting all the APIs described in the previous sections, the Enterprise Server includes a number of Java EE tools that are not part of the Java EE 6 platform but are provided as a convenience to the developer.

This section briefly summarizes the tools that make up the Enterprise Server. Instructions for starting and stopping the Enterprise Server, starting the Admin Console, and starting and stopping the Java DB database server are in [Chapter 2, “Using the Tutorial Examples.”](#)

## Tools

The Enterprise Server contains the tools listed in [Table 1–1](#). Basic usage information for many of the tools appears throughout the tutorial. For detailed information, see the online help in the GUI tools.

TABLE 1-1 Enterprise Server Tools

Tool	Description
Admin Console	A web-based GUI Enterprise Server administration utility. Used to stop the Enterprise Server and manage users, resources, and applications.
asadmin	A command-line Enterprise Server administration utility. Used to start and stop the Enterprise Server and manage users, resources, and applications.
appclient	A command-line tool that launches the application client container and invokes the client application packaged in the application client JAR file.
capture-schema	A command-line tool to extract schema information from a database, producing a schema file that the Enterprise Server can use for container-managed persistence.
package-appclient	A command-line tool to package the application client container libraries and JAR files.
Java DB database	A copy of the Java DB database server.
xjc	A command-line tool to transform, or bind, a source XML schema to a set of JAXB content classes in the Java programming language.
schemagen	A command-line tool to create a schema file for each namespace referenced in your Java classes.
wsimport	A command-line tool to generate JAX-WS portable artifacts for a given WSDL file. After generation, these artifacts can be packaged in a WAR file with the WSDL and schema documents along with the endpoint implementation and then deployed.
wsgen	A command-line tool to read a web service endpoint class and generate all the required JAX-WS portable artifacts for web service deployment and invocation.

## Using the Tutorial Examples

---

This chapter tells you everything you need to know to install, build, and run the examples. It covers the following topics:

- “Required Software” on page 59
- “Starting and Stopping the Enterprise Server” on page 63
- “Starting the Administration Console” on page 64
- “Starting and Stopping the Java DB Database Server” on page 64
- “Building the Examples” on page 65
- “Tutorial Example Directory Structure” on page 65
- “Getting the Latest Updates to the Tutorial” on page 66
- “Debugging Java EE Applications” on page 66

### Required Software

The following software is required to run the examples.

- “Java Platform, Standard Edition” on page 59
- “Java EE 6 Software Development Kit (SDK)” on page 60
- “Apache Ant” on page 60
- “Java EE 6 Tutorial Component” on page 61
- “NetBeans IDE” on page 62

### Java™ Platform, Standard Edition

To build, deploy, and run the examples, you need a copy of the Java Platform, Standard Edition 6.0 Software Development Kit (JDK 6). You can download the JDK 6 software from <http://java.sun.com/javase/downloads/index.jsp>.

Download the current JDK update that does not include any other software (such as NetBeans or Java EE).

## Java EE 6 Software Development Kit (SDK)

Sun GlassFish™ Enterprise Server v3 is targeted as the build and runtime environment for the tutorial examples. To build, deploy, and run the examples, you need a copy of the Enterprise Server and, optionally, NetBeans IDE. To obtain the Enterprise Server, you must install the Java EE 6 Software Development Kit (SDK), which you can download from <http://java.sun.com/javasee/downloads/>. Make sure you download the Java EE 6 SDK, not the Java EE 6 Web Profile SDK.

### SDK Installation Tips

During the installation of the SDK:

- Configure the Enterprise Server administration username and password as anonymous. This is the default setting.
- Accept the default port values for the Admin Port (4848) and the HTTP Port (8080).
- Allow the installer to download and configure the Update Tool. If you access the Internet through a firewall, provide the proxy host and port.

This tutorial refers to the directory where you install the Enterprise Server as *as-install-parent*. For example, the default installation directory on Microsoft Windows is `C:\glassfishv3`, so *as-install-parent* is `C:\glassfishv3`. The Enterprise Server itself is installed in *as-install*, the `glassfish` directory under *as-install-parent*. So on Microsoft Windows, *as-install* is `C:\glassfishv3\glassfish`.

After you install the Enterprise Server, add the following directories to your PATH to avoid having to specify the full path when you use commands:

```
as-install-parent/bin
as-install/bin
```

## Apache Ant

Ant is a Java technology-based build tool developed by the Apache Software Foundation (<http://ant.apache.org/>), and is used to build, package, and deploy the tutorial examples. To run the tutorial examples, you need Ant 1.7.1. If you do not already have Ant 1.7.1, you can install it from the Update Tool that is part of the Enterprise Server.

### ▼ To Obtain Apache Ant

#### 1 Start the Update Tool.

- From the command line, type the command `updateTool`.

- **On a Windows system, select the following:**  
Start → All Programs → Java EE 6 SDK → Start Update Tool
- 2 Expand the Sun GlassFish Enterprise Server v3 node.**
- 3 Select the Available Add-ons node.**
- 4 From the list, select the Apache Ant Build Tool checkbox.**
- 5 Click Install.**
- 6 Accept the license agreement.**  
After installation, Apache Ant appears in the list of installed components. The tool is installed in the *as-install-parent/ant* directory.

**Next Steps** To use the ant command, add *as-install/ant/bin* to your PATH environment variable.

## Java EE 6 Tutorial Component

The tutorial example source is contained in the tutorial component. To obtain the tutorial component, use the Update Tool.

### ▼ To Obtain the Tutorial Component

- 1 Start the Update Tool.**
  - **From the command line, type the command** `updateTool`.
  - **On a Windows system, select the following:**  
Start → All Programs → Java EE 6 SDK → Start Update Tool
- 2 Expand the Sun GlassFish Enterprise Server v3 node.**
- 3 Select the Available Add-ons node.**
- 4 From the list, select the Java EE 6 Tutorial checkbox.**
- 5 Click Install.**

## 6 Accept the license agreement.

After installation, the Java EE 6 Tutorial appears in the list of installed components. The tool is installed in the *as-install/docs/javaee-tutorial* directory. This directory contains two subdirectories, *docs* and *examples*. The *examples* directory contains subdirectories for each of the technologies discussed in the tutorial.

**Next Steps** Updates to the Java EE 6 Tutorial are published periodically. For details on obtaining these updates, see “[Getting the Latest Updates to the Tutorial](#)” on page 66.

## NetBeans IDE

The NetBeans integrated development environment (IDE) is a free, open-source IDE for developing Java applications, including enterprise applications. NetBeans IDE supports the Java EE platform. You can build, package, deploy, and run the tutorial examples from within NetBeans IDE.

To run the tutorial examples, you need the latest version of NetBeans IDE. You can download NetBeans IDE from <http://www.netbeans.org/downloads/index.html>.

### ▼ To Add Enterprise Server as a Server in NetBeans IDE

To run the tutorial examples in NetBeans IDE, you must register your Enterprise Server installation as a NetBeans Server Instance. Follow these instructions to register the Enterprise Server in NetBeans IDE.

- 1 **Select Tools → Servers to open the Servers dialog.**
- 2 **Click Add Server.**
- 3 **Under Server, select GlassFish v3 and click Next.**
- 4 **Under Server Location, enter the location of your Enterprise Server installation and click Next.**
- 5 **Select Register Local Default Domain.**
- 6 **Click Finish.**

## Starting and Stopping the Enterprise Server

To start the Enterprise Server, open a terminal window or command prompt and execute the following:

```
asadmin start-domain --verbose
```

A *domain* is a set of one or more Enterprise Server instances managed by one administration server. Associated with a domain are the following:

- The Enterprise Server's port number. The default is 8080.
- The administration server's port number. The default is 4848.
- An administration user name and password.

You specify these values when you install the Enterprise Server. The examples in this tutorial assume that you chose the default ports.

With no arguments, the `start-domain` command initiates the default domain, which is `domain1`. The `--verbose` flag causes all logging and debugging output to appear on the terminal window or command prompt (it will also go into the server log, which is located in `domain-dir/logs/server.log`).

Or, on Windows, choose the following:

Start → All Programs → Java EE 6 SDK → Start Application Server

After the server has completed its startup sequence, you will see the following output:

```
Domain domain1 started.
```

To stop the Enterprise Server, open a terminal window or command prompt and execute:

```
asadmin stop-domain domain1
```

Or, on Windows, choose the following:

Start → All Programs → Java EE 6 SDK → Stop Application Server

When the server has stopped you will see the following output:

```
Domain domain1 stopped.
```

## Starting the Administration Console

To administer the Enterprise Server and manage users, resources, and Java EE applications, use the Administration Console tool. The Enterprise Server must be running before you invoke the Administration Console. To start the Administration Console, open a browser at `http://localhost:4848/`.

Or, on Windows, choose the following:

Start → All Programs → Java EE 6 SDK → Administration Console

### ▼ To Start the Administration Console in NetBeans IDE

- 1 Click the Services tab.
- 2 Expand the Servers node.
- 3 Right-click the Enterprise Server instance and select View Admin Console

---

**Note** – You must configure a web browser with NetBeans IDE in order for the Administration Console to be started from within NetBeans IDE.

---

## Starting and Stopping the Java DB Database Server

The Enterprise Server includes the Java DB database.

To start the Java DB database server, open a terminal window or command prompt and execute:

```
asadmin start-database
```

To stop the Java DB server, open a terminal window or command prompt and execute:

```
asadmin stop-database
```

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

To start the database server using NetBeans IDE, follow these steps:

1. Click the Services tab.
2. Expand the Databases node.
3. Right-click Java DB and choose Start Server.

To stop the database using NetBeans IDE, choose Stop Server.

## Building the Examples

The tutorial examples are distributed with a configuration file for either NetBeans IDE or Ant. Directions for building the examples are provided in each chapter. Either NetBeans IDE or Ant may be used to build, package, deploy, and run the examples.

## Tutorial Example Directory Structure

To facilitate iterative development and keep application source separate from compiled files, the tutorial examples use the Java BluePrints application directory structure.

Each application module has the following structure:

- `build.xml`: Ant build file
- `src/java`: Java source files for the module
- `src/conf`: configuration files for the module, with the exception of web applications
- `web`: web pages, style sheets, tag files, and images
- `web/WEB-INF`: configuration files for web applications
- `nbproject`: NetBeans project files

Examples that have multiple application modules packaged into an enterprise application archive (or EAR) have submodule directories that use the following naming conventions:

- `example-name-app-client`: Application clients
- `example-name-ejb`: Enterprise bean JAR files
- `example-name-war`: web applications

The Ant build files (`build.xml`) distributed with the examples contain targets to create a `build` subdirectory and to copy and compile files into that directory; a `dist` subdirectory, which holds the packaged module file; and a `client-jar` directory, which holds the retrieved application client JAR.

## Getting the Latest Updates to the Tutorial

Check for any updates to the tutorial by using the Update Center included with the Java EE 6 SDK.

### ▼ To Update the Tutorial through the Update Center

- 1 Open the Services tab in NetBeans IDE and expand Servers.
- 2 Right-click the GlassFish v3 instance and select View Update Center to display the Update Tool.
- 3 Select Available Updates in the tree to display a list of updated packages.
- 4 Look for updates to the Java EE 6 Tutorial (javaee-tutorial) package.
- 5 If there is an updated version of the Tutorial, select Java EE 6 Tutorial (javaee-tutorial) and click Install.

## Debugging Java EE Applications

This section describes how to determine what is causing an error in your application deployment or execution.

### Using the Server Log

One way to debug applications is to look at the server log in *domain-dir/logs/server.log*. The log contains output from the Enterprise Server and your applications. You can log messages from any Java class in your application with `System.out.println` and the Java Logging APIs (documented at <http://java.sun.com/javase/6/docs/technotes/guides/logging/index.html>) and from web components with the `ServletContext.log` method.

If you start the Enterprise Server with the `--verbose` flag, all logging and debugging output will appear on the terminal window or command prompt and the server log. If you start the Enterprise Server in the background, debugging information is only available in the log. You can view the server log with a text editor or with the Administration Console log viewer.

To use the log viewer:

1. Select the Enterprise Server node.
2. Click the View Log Files button. The log viewer will open and display the last 40 entries.

If you wish to display other entries:

1. Click the Modify Search button.
2. Specify any constraints on the entries you want to see.
3. Click the Search button at the top of the log viewer.

## Using a Debugger

The Enterprise Server supports the Java Platform Debugger Architecture (JPDA). With JPDA, you can configure the Enterprise Server to communicate debugging information using a socket.

To debug an application using a debugger:

1. Enable debugging in the Enterprise Server using the Administration Console:
  - a. Expand the Configuration node.
  - b. Select the JVM Settings node. The default debug options are set to:

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

As you can see, the default debugger socket port is 9009. You can change it to a port not in use by the Enterprise Server or another service.

- c. Select the Debug Enabled check box.
  - d. Click the Save button.
2. Stop the Enterprise Server and then restart it.



( PART II

# The Web Tier

Part Two explores the technologies in the web tier.



# Getting Started with Web Applications

---

A web application is a dynamic extension of a web or application server. There are two types of web applications:

- **Presentation-oriented:** A presentation-oriented web application generates interactive web pages containing various types of markup language (HTML, XHTML, XML, and so on) and dynamic content in response to requests. [Chapter 4, “JavaServer Faces Technology,”](#) through [Chapter 9, “Developing With JavaServer Faces Technology,”](#) cover how to develop presentation-oriented web applications.
- **Service-oriented:** A service-oriented web application implements the endpoint of a web service. Presentation-oriented applications are often clients of service-oriented web applications. [Chapter 13, “Building Web Services with JAX-WS”](#) and [Chapter 14, “Building RESTful Web Services with JAX-RS and Jersey”](#) cover how to develop service-oriented web applications.

The following topics are addressed here:

- “Web Applications” on page 71
- “Web Application Life Cycle” on page 73
- “Web Modules” on page 75
- “Configuring Web Applications” on page 81
- “Further Information about Web Applications” on page 88

## Web Applications

In the Java™ EE platform, *web components* provide the dynamic extension capabilities for a web server. Web components are either Java servlets, web pages, web service endpoints, or JSP pages. The interaction between a web client and a web application is illustrated in [Figure 3–1](#). The client sends an HTTP request to the web server. A web server that implements Java Servlet and JavaServer Pages™ technology converts the request into an `HttpServletRequest` object. This object is delivered to a web component, which can interact with JavaBeans components or a database to generate dynamic content. The web component can then generate an

HTTPServletResponse or it can pass the request to another web component. Eventually a web component generates a HttpServletResponse object. The web server converts this object to an HTTP response and returns it to the client.

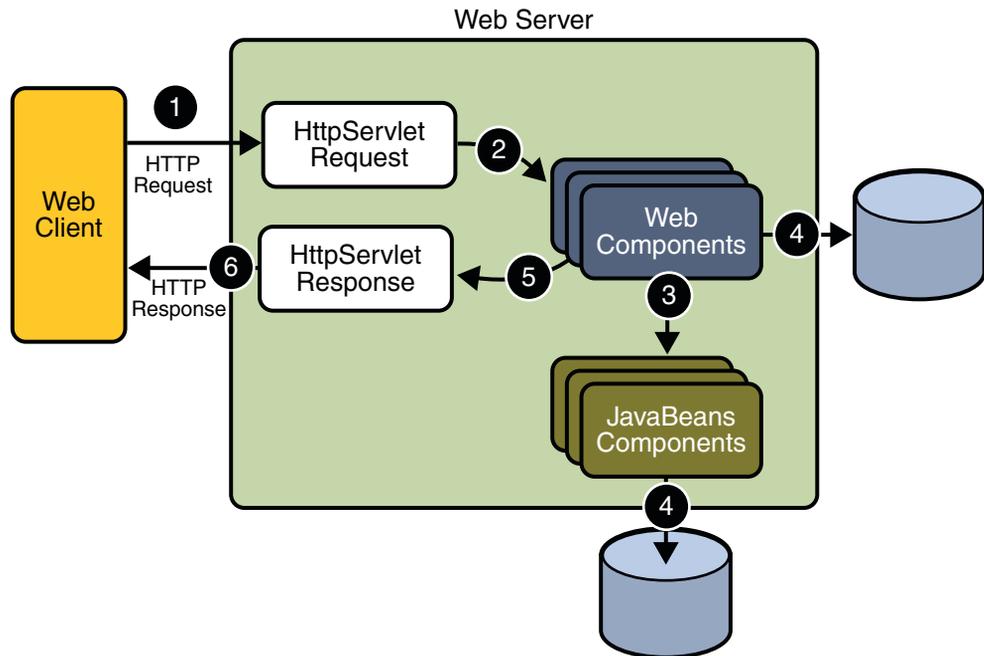


FIGURE 3-1 Java Web Application Request Handling

*Servlets* are Java programming language classes that dynamically process requests and construct responses. Java technologies, such as *JavaServer™ Faces* and *Facelets*, and frameworks are used for building interactive web applications. Although *Servlets* and *Java Server Faces* and *Facelets* pages can be used to accomplish similar things, each has its own strengths. *Servlets* are best suited for service-oriented applications (web service endpoints are implemented as *Servlets*) and the control functions of a presentation-oriented application, such as dispatching requests and handling nontextual data. *Java Server Faces* and *Facelets* pages are more appropriate for generating text-based markup, such as XHTML, and are generally used for presentation-oriented applications.

Web components are supported by the services of a runtime platform called a *web container*. A web container provides services such as request dispatching, security, concurrency, and life-cycle management. It also gives web components access to APIs such as naming, transactions, and email.

Certain aspects of web application behavior can be configured when the application is installed, or *deployed*, to the web container. The configuration information is maintained in a text file in XML format called a *web application deployment descriptor* (DD). A DD must conform to the schema described in the [Java Servlet Specification](#).

This chapter gives a brief overview of the activities involved in developing web applications. First it summarizes the web application life cycle. Then it describes how to package and deploy very simple web applications on the Sun GlassFish Enterprise Server. It moves on to configuring web applications and discusses how to specify the most commonly used configuration parameters.

## Web Application Life Cycle

A web application consists of web components, static resource files such as images, and helper classes and libraries. The web container provides many supporting services that enhance the capabilities of web components and make them easier to develop. However, because a web application must take these services into account, the process for creating and running a web application is different from that of traditional stand-alone Java classes.

The process for creating, deploying, and executing a web application can be summarized as follows:

1. Develop the web component code.
2. Develop the web application deployment descriptor.
3. Compile the web application components and helper classes referenced by the components.
4. Optionally package the application into a deployable unit.
5. Deploy the application into a web container.
6. Access a URL that references the web application.

Developing web component code is covered in the later chapters. Steps 2 through 4 are expanded on in the following sections and illustrated with a Hello, World-style presentation-oriented application. This application allows a user to enter a name into an HTML form ([Figure 3–2](#)) and then displays a greeting after the name is submitted ([Figure 3–3](#)).



FIGURE 3-2 Greeting Form

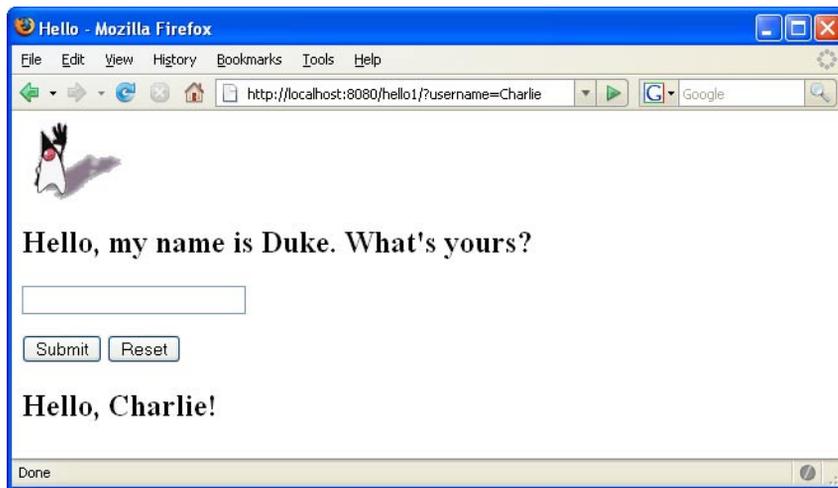


FIGURE 3-3 Response

The Hello application contains two web components that generate the greeting and the response. This chapter discusses `hello2`, a servlet-based web application in which the components are implemented by two servlet classes (*tut-install/examples/web/hello2/src/servlets/GreetingServlet.java* and *tut-install/examples/web/hello2/src/servlets/ResponseServlet.java*). The application is used to illustrate tasks involved in packaging, deploying, configuring, and running an application that contains web components. The section [Chapter 2, “Using the Tutorial](#)

[Examples](#),” explains how to get the code for the example. The source code for the example is in the `tut-install/examples/web/hello2/` directory.

## Web Modules

In the Java EE architecture, web components and static web content files such as images are called *web resources*. A *web module* is the smallest deployable and usable unit of web resources. A Java EE web module corresponds to a *web application* as defined in the Java Servlet specification.

In addition to web components and web resources, a web module can contain other files:

- Server-side utility classes (database beans, shopping carts, and so on). Often these classes conform to the JavaBeans component architecture.
- Client-side classes (applets and utility classes).

A web module has a specific structure. The top-level directory of a web module is the *document root* of the application. The document root is where XHTML pages, *client-side* classes and archives, and static web resources, such as images, are stored.

The document root contains a subdirectory named `WEB-INF`, which contains the following files and directories:

- `web.xml`: The web application deployment descriptor
- Tag library descriptor files
- `classes`: A directory that contains *server-side classes*: servlets, utility classes, and JavaBeans components
- `tags`: A directory that contains tag files, which are implementations of tag libraries
- `lib`: A directory that contains JAR archives of libraries called by server-side classes

If your web module does not contain any servlets, filter, or listener components then it does not need a web application deployment descriptor. In other words, if your web module only contains XHTML pages and static files, you are not required to include a `web.xml` file.

You can also create application-specific subdirectories (that is, package directories) in either the document root or the `WEB-INF/classes/` directory.

A web module can be deployed as an unpacked file structure or can be packaged in a JAR file known as a web archive (WAR) file. Because the contents and use of WAR files differ from those of JAR files, WAR file names use a `.war` extension. The web module just described is portable; you can deploy it into any web container that conforms to the Java Servlet Specification.

To deploy a WAR on the Enterprise Server, the file must also contain a runtime deployment descriptor. The runtime deployment descriptor is an XML file that contains information such as the context root of the web application and the mapping of the portable names of an

application's resources to the Enterprise Server's resources. The Enterprise Server web application runtime DD is named `sun-web.xml` and is located in the `WEB-INF` directory along with the web application DD. The structure of a web module that can be deployed on the Enterprise Server is shown in [Figure 3-4](#).

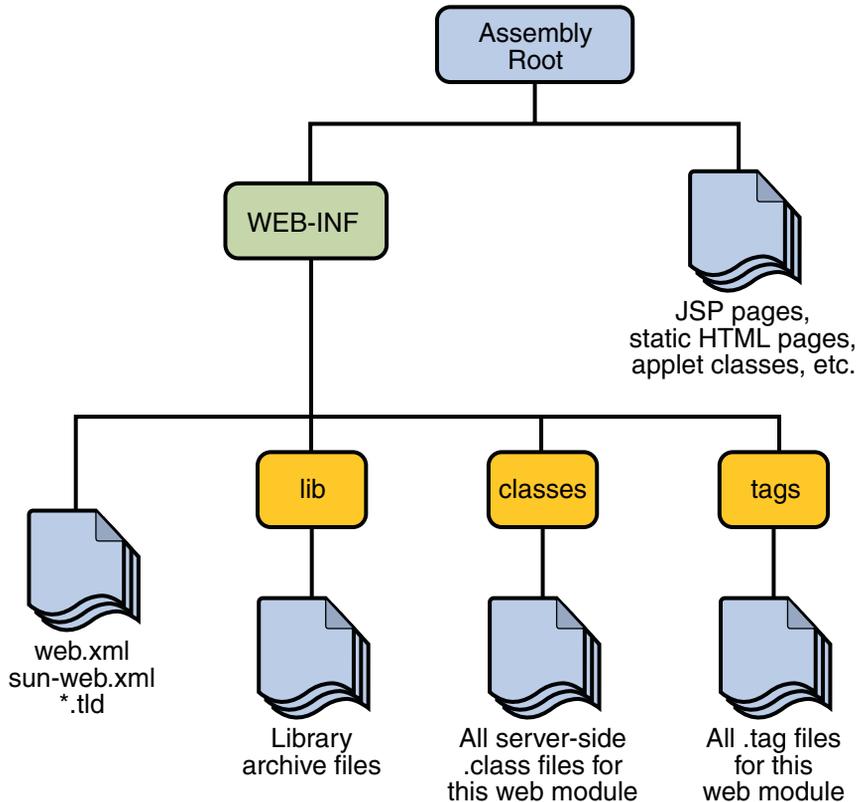


FIGURE 3-4 Web Module Structure

## Packaging Web Modules

A web module must be packaged into a WAR in certain deployment scenarios and whenever you want to distribute the web module. You package a web module into a WAR by executing the `jar` command in a directory laid out in the format of a web module, by using the Ant utility, or by using the IDE tool of your choice. This tutorial shows you how to use NetBeans IDE or Ant to build, package, and deploy the sample applications.

To build the `hello2` application with NetBeans IDE, follow these instructions:

1. Select File→Open Project.
2. In the Open Project dialog, navigate to:  
`tut-install/examples/web/`
3. Select the `hello2` folder.
4. Select the Open as Main Project check box.
5. Click Open Project.
6. In the Projects tab, right-click the `hello2` project and select Build.

To build the `hello2` application using the Ant utility, follow these steps:

1. In a terminal window, go to `tut-install/examples/web/hello2/`.
2. Type **ant**. This command will spawn any necessary compilations, copy files to the `tut-install/examples/web/hello2/build/` directory, create the WAR file, and copy it to the `tut-install/examples/web/hello2/dist/` directory.

## Deploying a WAR File

You can deploy a WAR file to the Enterprise Server in a few ways:

- Copying the WAR into the `domain-dir/autodeploy/` directory.
- Using the Admin Console.
- By running `asadmin` or `ant` to deploy the WAR.
- Using NetBeans IDE.

All these methods are described briefly in this chapter; however, throughout the tutorial, you will use `ant` and NetBeans IDE for packaging and deploying.

## Setting the Context Root

A *context root* identifies a web application in a Java EE server. You specify the context root when you deploy a web module. A context root must start with a forward slash (/) and end with a string.

In a packaged web module for deployment on the Enterprise Server, the context root is stored in `sun-web.xml`.

To edit the context root, do the following:

1. Expand your project tree in the Projects pane of NetBeans IDE.
2. Expand the Web Pages and WEB-INF nodes of your project.
3. Double-click `sun-web.xml`.
4. In the editor pane, click Edit As XML.

5. Edit the context root, which is enclosed by the context - root element.

## Deploying a Packaged Web Module

If you have deployed the `hel1o2` application, before proceeding with this section, undeploy the application by following one of the procedures described in [“Undeploying Web Modules” on page 80](#).

### Deploying with the Admin Console

1. Expand the Applications node.
2. Click the Deploy button.
3. Select the radio button labeled “Package file to be uploaded to the Application Server.”
4. Type the full path to the WAR file (or click on Browse to find it), and then click the OK button.
5. Click Next.
6. Type the application name.
7. Type the context root.
8. Select the Enabled box.
9. Click the Finish button.

### Deploying with `asadmin`

To deploy a WAR with `asadmin`, open a terminal window or command prompt and execute

```
asadmin deploy full-path-to-war-file
```

### Deploying with Ant

To deploy a WAR with the Ant tool, open a terminal window or command prompt in the directory where you built and packaged the WAR, and execute

```
ant deploy
```

### Deploying with NetBeans IDE

To deploy a WAR with NetBeans IDE, do the following:

1. Select File→Open Project.
2. In the Open Project dialog, navigate to your project and open it.
3. In the Projects tab, right-click the project and select Deploy.

## Testing Deployed Web Modules

Now that the web module is deployed, you can view it by opening the application in a web browser. By default, the application is deployed to host `localhost` on port `8080`. The context root of the web application is `hello2`.

To test the application, follow these steps:

1. Open a web browser.
2. Enter the following URL in the web address box:

```
http://localhost:8080/hello2
```

3. Enter your name, and click Submit.

The application should display the name you submitted.

## Listing Deployed Web Modules

The Enterprise Server provides two ways to view the deployed web modules: the Admin Console and the `asadmin` command.

To use the Admin Console:

1. Open the URL `http://localhost:4848/` in a browser.
2. Expand the Applications node.

Use the `asadmin` command as follows:

```
asadmin list-components
```

## Updating Web Modules

A typical iterative development cycle involves deploying a web module and then making changes to the application components. To update a deployed web module, you must do the following:

1. Recompile any modified classes.
2. If you have deployed a packaged web module, update any modified components in the WAR.
3. Redeploy the module.
4. Reload the URL in the client.

## Dynamic Reloading

If dynamic reloading is enabled, you do not have to redeploy an application or module when you change its code or deployment descriptors. All you have to do is copy the changed pages or class files into the deployment directory for the application or module. The deployment directory for a web module named *context-root* is *domain-dir/applications/context-root*. The server checks for changes periodically and redeploys the application, automatically and dynamically, with the changes.

This capability is useful in a development environment, because it allows code changes to be tested quickly. Dynamic reloading is not recommended for a production environment, however, because it may degrade performance. In addition, whenever a reload is done, the sessions at that time become invalid and the client must restart the session.

To enable dynamic reloading, use the Admin Console:

1. Select the Applications Server node.
2. Select the Advanced tab.
3. Check the Reload Enabled box to enable dynamic reloading.
4. Enter a number of seconds in the Reload Poll Interval field to set the interval at which applications and modules are checked for code changes and dynamically reloaded.
5. Click the Save button.

In addition, to load new servlet files, do the following:

1. Create an empty file named `.reload` at the root of the module:  

```
domain-dir/applications/context-root/.reload
```
2. Explicitly update the `.reload` file's time stamp each time you make these changes. On UNIX, execute  

```
touch .reload
```

## Undeploying Web Modules

You can undeploy web modules in four ways: you can use NetBeans IDE, the Admin Console, the `asadmin` command, or the Ant tool.

To use NetBeans IDE:

1. Ensure the Enterprise Server is running.
2. In the Services window, expand the Servers node, Enterprise Server instance and the Applications node.
3. Right-click the application or module and choose Undeploy.

To use the Admin Console:

1. Open the URL `http://localhost:4848/` in a browser.
2. Expand the Applications node.
3. Select the check box next to the module you wish to undeploy.
4. Click the Undeploy button.

Use the `asadmin` command as follows:

```
asadmin undeploy context-root
```

To use the Ant tool, execute the following command in the directory where you built and packaged the WAR:

```
ant undeploy
```

## Configuring Web Applications

Web applications are configured by means of elements contained in the web application deployment descriptor.

The following sections give a brief introduction to the web application features you will usually want to configure.

In the following sections, examples demonstrate procedures for configuring the Hello, World application. If Hello, World does not use a specific configuration feature, the section gives references to other examples that illustrate how to specify the deployment descriptor element.

## Mapping URLs to Web Components

When a request is received by the web container it must determine which web component should handle the request. It does so by mapping the URL path contained in the request to a web application and a web component. A URL path contains the context root and an alias:

```
http://host:port/context-root/alias
```

## Setting the Component Alias

The *alias* identifies the web component that should handle a request. The alias path must start with a forward slash (/) and end with a string or a wildcard expression with an extension (for example, `*.jsp`).

The `hello2` application has two servlets that need to be mapped in the `web.xml` file. You can edit a web application's `web.xml` file in NetBeans IDE by doing the following:

1. Select `File`→`Open Project`.
2. In the `Open Project` dialog, navigate to:  
`tut-install/examples/web/`
3. Select the `hello2` folder.
4. Select the `Open as Main Project` check box.
5. Click `Open Project`.
6. Expand the project tree in the `Projects` pane.
7. Expand the `Web pages` node and then the `WEB-INF` node in the project tree.
8. Double-click the `web.xml` file inside the `WEB-INF` node.

The following steps describe how to make the necessary edits to the `web.xml` file, including how to set the display name and how to map the servlet components. Because the edits are already in the file, you can just use the steps to view the settings.

To set the display name:

1. Click `General` at the top of the editor to open the general view.
2. Enter `hello2` in the `Display Name` field.

To perform the servlet mappings:

1. Click `Servlets` at the top of the editor to open the servlets view.
2. Click `Add Servlet`.
3. In the `Add Servlet` dialog, enter `GreetingServlet` in the `Servlet Name` field.
4. Enter `servlets.GreetingServlet` in the `Servlet Class` field.
5. Enter `/greeting` in the `URL Pattern` field.
6. Click `OK`.
7. Repeat the preceding steps, except enter `ResponseServlet` as the servlet name, `servlets.ResponseServlet` as the servlet class, and `/response` as the URL pattern.

If you are not using NetBeans IDE, you can add these settings using a text editor.

To package the example with NetBeans IDE, do the following:

1. Select `File`→`Open Project`.
2. In the `Open Project` dialog, navigate to:  
`tut-install/examples/web/`
3. Select the `hello2` folder.
4. Select the `Open as Main Project` check box.

5. Click Open Project.
6. In the Projects tab, right-click the `hello2` project and select Build.

To package the example with the Ant utility, do the following:

1. In a terminal window, go to `tut-install/examples/web/hello2/`.
2. Type **ant**. This target will build the WAR file and copy it to the `tut-install/examples/web/hello2/dist/` directory.

To deploy the example using NetBeans IDE, right-click on the project in the Projects pane and select Deploy.

To deploy the example using Ant, type **ant deploy**. The deploy target in this case gives you an incorrect URL to run the application. To run the application, please use the URL shown at the end of this section.

To run the application, first deploy the web module, and then open the URL `http://localhost:8080/hello2/greeting` in a browser.

## Declaring Welcome Files

The *welcome files* mechanism allows you to specify a list of files that the web container will use for appending to a request for a URL (called a *valid partial request*) that is not mapped to a web component.

For example, suppose you define a welcome file `welcome.html`. When a client requests a URL such as `host:port/webapp/directory`, where *directory* is not mapped to a servlet or XHTML page, the file `host:port/webapp/directory/welcome.html` is returned to the client.

If a web container receives a valid partial request, the web container examines the welcome file list and appends to the partial request each welcome file in the order specified and checks whether a static resource or servlet in the WAR is mapped to that request URL. The web container then sends the request to the first resource in the WAR that matches.

If no welcome file is specified, the Enterprise Server will use a file named `index.xhtml` as the default welcome file. If there is no welcome file and no file named `index.xhtml`, the Enterprise Server returns a directory listing.

## Setting Initialization Parameters

The web components in a web module share an object that represents their application context. You can pass initialization parameters to the context or to a web component.

To add a context parameter using NetBeans IDE, do the following:

1. Open the project if you haven't already.
2. Expand the project's node in the Projects pane.
3. Expand the Web Pages node and then the WEB-INF node.
4. Double-click `web.xml`.
5. Click General at the top of the editor pane.
6. Select the Context Parameters node.
7. Click Add.
8. In the Add Context Parameter dialog, do the following:
  - a. Enter the name that specifies the context object in the Param Name field.
  - b. Enter the parameter to pass to the context object in the Param Value field.
  - c. Click OK.

Alternatively, you can edit the XML of the `web.xml` file directly by clicking XML at the top of the editor pane and using the following elements to add a context parameter:

- A `param-name` element that specifies the context object
- A `param-value` element that specifies the parameter to pass to the context object
- A `context-param` element that encloses the previous two elements

To add a web component initialization parameter using NetBeans IDE, do the following:

1. Open the project if you haven't already.
2. Expand the project's node in the Projects pane.
3. Expand the Web Pages node and then the WEB-INF node.
4. Double-click `web.xml`.
5. Click Servlets at the top of the editor pane.
6. After entering the servlet's name, class, and URL pattern, click the Add button under the Initialization Parameters table.
7. In the Add Initialization Parameter dialog:
  - a. Enter the name of the parameter in the Param Name field.
  - b. Enter the parameter's value in the Param Value Field.
  - c. Click OK.

Alternatively, you can edit the XML of the `web.xml` file directly by clicking XML at the top of the editor pane and using the following elements to add a context parameter:

- A `param-name` element that specifies the name of the initialization parameter
- A `param-value` element that specifies the value of the initialization parameter
- An `init-param` element that encloses the previous two elements

## Mapping Errors to Error Screens

When an error occurs during execution of a web application, you can have the application display a specific error screen according to the type of error. In particular, you can specify a mapping between the status code returned in an HTTP response or a Java programming language exception returned by any web component and any type of error screen.

To set up error mappings using NetBeans IDE, do the following:

1. Open the project if you haven't already.
2. Expand the project's node in the Projects pane.
3. Expand the Web Pages node and then the WEB-INF node.
4. Double-click `web.xml`.
5. Click Pages at the top of the editor pane.
6. Expand the Error Pages node.
7. Click Add.
8. In the Add Error Page dialog:
  - a. Click Browse to locate the page that you want to act as the error page.
  - b. Enter the HTTP status code that will cause the error page to be opened in the Error Code field.
  - c. Enter the exception that will cause the error page to load in the Exception Type field.
  - d. Click OK.

Alternatively, you can click XML at the top of the editor pane and enter the error page mapping by hand using the following elements:

- An `exception-type` element specifying either the exception or the HTTP status code that will cause the error page to be opened.
- A `location` element that specifies the name of a web resource to be invoked when the status code or exception is returned. The name should have a leading forward slash (/).
- An `error-page` element that encloses the previous two elements.

You can have multiple `error-page` elements in your deployment descriptor. Each one of the elements identifies a different error that causes an error page to open. This error page can be the same for any number of `error-page` elements.

## Declaring Resource References

If your web component uses objects such as enterprise beans, data sources, or web services, you use Java EE annotations to inject these resources into your application. Annotations eliminate a lot of the boilerplate lookup code and configuration elements that previous versions of Java EE required.

Although resource injection using annotations can be more convenient for the developer, there are some restrictions from using it in web applications. First, you can only inject resources into container-managed objects. This is because a container must have control over the creation of a component so that it can perform the injection into a component. As a result, you cannot inject resources into objects such as simple JavaBeans components. However, JavaServer Faces managed beans are managed by the container; therefore, they can accept resource injections.

Components that can accept resource injections are listed in [Table 3-1](#).

This section describes how to use a couple of the annotations supported by a servlet container to inject resources. [Chapter 23, “Running the Persistence Examples,”](#) describes how web applications use annotations supported by the Java Persistence API. [Chapter 28, “Getting Started Securing Web Applications,”](#) describes how to use annotations to specify information about securing web applications.

**TABLE 3-1** Web Components That Accept Resource Injections

Component	Interface/Class
Servlets	<code>javax.servlet.Servlet</code>
Servlet Filters	<code>javax.servlet.ServletFilter</code>
Event Listeners	<code>javax.servlet.ServletContextListener</code>
	<code>javax.servlet.ServletContextAttributeListener</code>
	<code>javax.servlet.ServletRequestListener</code>
	<code>javax.servlet.ServletRequestAttributeListener</code>
	<code>javax.servlet.http.HttpSessionListener</code>
	<code>javax.servlet.http.HttpSessionAttributeListener</code>
	<code>javax.servlet.http.HttpSessionBindingListener</code>
Taglib Listeners	Same as above
Taglib Tag Handlers	<code>javax.servlet.jsp.tagext.JspTag</code>
Managed Beans	Plain Old Java Objects

## Declaring a Reference to a Resource

The `@Resource` annotation is used to declare a reference to a resource such as a data source, an enterprise bean, or an environment entry.

The `@Resource` annotation is specified on a class, method or field. The container is responsible for injecting references to resources declared by the `@Resource` annotation and mapping it to the proper JNDI resources. In the following example, the `@Resource` annotation is used to inject a data source into a component that needs to make a connection to the data source, as is done when using JDBC technology to access a relational database:

```
@Resource javax.sql.DataSource catalogDS;
public getProductsByCategory() {
    // get a connection and execute the query
    Connection conn = catalogDS.getConnection();
    ..
}
```

The container injects this data source prior to the component being made available to the application. The data source JNDI mapping is inferred from the field name `catalogDS` and the type, `javax.sql.DataSource`.

If you have multiple resources that you need to inject into one component, you need to use the `@Resources` annotation to contain them, as shown by the following example:

```
@Resources ({
    @Resource (name="myDB" type=java.sql.DataSource),
    @Resource(name="myMQ" type=javax.jms.ConnectionFactory)
})
```

The web application examples in this tutorial use the Java Persistence API to access relational databases. This API does not require you to explicitly create a connection to a data source. Therefore, the examples do not use the `@Resource` annotation to inject a data source. However, this API supports the `@PersistenceUnit` and `@PersistenceContext` annotations for injecting `EntityManagerFactory` and `EntityManager` instances, respectively. [Chapter 23, “Running the Persistence Examples,”](#) describes these annotations and the use of the Java Persistence API in web applications.

## Declaring a Reference to a Web Service

The `@WebServiceRef` annotation provides a reference to a web service. The following example shows uses the `@WebServiceRef` annotation to declare a reference to a web service. `WebServiceRef` uses the `wsdlLocation` element to specify the URI of the deployed service’s WSDL file:

```
...
import javax.xml.ws.WebServiceRef;
...
```

```
public class ResponseServlet extends HttpServlet {
    @WebServiceRef(wsdlLocation=
        "http://localhost:8080/helloservice/hello?wsdl")
    static HelloService service;
}
```

## Further Information about Web Applications

For more information on web applications, see:

- The JavaServer Faces 2.0 specification:  
<http://jcp.org/en/jsr/detail?id=314>
- The JavaServer Faces web site:  
<http://java.sun.com/javaee/javaserverfaces/>
- The Java Servlet 3.0 specification:  
<http://jcp.org/en/jsr/detail?id=315>
- The Java Servlet web site:  
<http://java.sun.com/products/servlet>

# JavaServer™ Faces Technology

---

JavaServer Faces technology is a server-side component framework for building Java technology-based web applications.

JavaServer Faces technology consists of the following:

- An [API](#) for representing components and managing their state; handling events, server-side validation, and data conversion; defining page navigation; supporting internationalization and accessibility; and providing extensibility for all these features
- Tag libraries for adding components to web pages and for connecting components to server-side objects

JavaServer Faces technology provides a well-defined programming model and various tag libraries. These features significantly ease the burden of building and maintaining web applications with server-side UIs. With minimal effort, you can complete the following tasks:

- Create a web page
- Drop components onto a web page by adding component tags
- Bind components on a page to server-side data
- Wire component-generated events to server-side application code
- Save and restore application state beyond the life of server requests
- Reuse and extend components through customization

This chapter provides an overview of JavaServer Faces technology. After explaining what a JavaServer Faces application is and going over some of the primary benefits of using JavaServer Faces technology, it describes the process of creating a simple JavaServer Faces application. This chapter also introduces the JavaServer Faces lifecycle by describing the example JavaServer Faces application progressing through the lifecycle stages.

The following topics are addressed here:

- [“What Is a JavaServer Faces Application?”](#) on page 90
- [“JavaServer Faces Technology Benefits”](#) on page 91
- [“Creating a Simple JavaServer Faces Application”](#) on page 92

- “Further Information about JavaServer Faces Technology” on page 96

## What Is a JavaServer Faces Application?

The functionality provided by a JavaServer Faces application is similar to that of any other Java web application. A typical JavaServer Faces application includes the following parts:

- A set of web pages in which components are laid out.
- A set of tags to add components to the web page.
- A set of *backing beans* which are JavaBeans™ components that define properties and functions for components on a page.
- A web deployment descriptor (`web.xml` file).
- Optionally, one or more application configuration resource files such as a `faces-config.xml` file, which can be used to define page navigation rules and configure beans and other custom objects such as custom components.
- Optionally, a set of custom objects created by the application developer. These objects can include custom components, validators, converters, or listeners.
- A set of custom tags for representing custom objects on the page.

Figure 4–1 describes the interaction between client and server in a typical JavaServer Faces application. In response to a client request, a web page is rendered by the web container that implements JavaServer Faces technology.

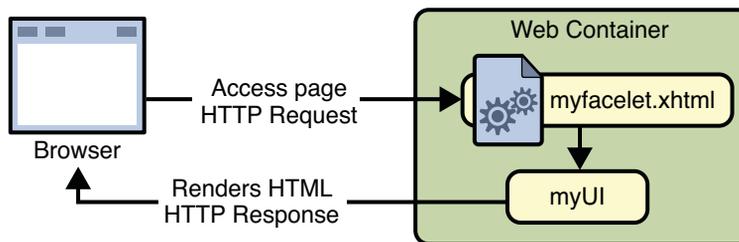


FIGURE 4–1 Responding to a Client Request for a JavaServer Faces Page

The web page, `myfacelet.xhtml`, is built using JavaServer Faces component tags. Component tags are used to add components to the *view* (represented by `myUI` in the diagram), which is the server-side representation of the page. In addition to components, the web page can also reference objects such as the following:

- Any event listeners, validators, and converters that are registered on the components
- The JavaBeans components that capture the data and process the application-specific functionality of the components

On request from the client, the view is rendered as a response. Rendering is the process whereby, based on the server-side view, the web container generates output such as HTML or XHTML that can be read by the browser.

## JavaServer Faces Technology Benefits

One of the greatest advantages of JavaServer Faces technology is that it offers a clean separation between behavior and presentation for web applications.

A JavaServer Faces application can map HTTP requests to component-specific event handling and manage components as stateful objects on the server. JavaServer Faces technology allows you to build web applications that implement the finer-grained separation of behavior and presentation that is traditionally offered by client-side UI architectures.

The separation of logic from presentation also allows each member of a web application development team to focus on a single piece of the development process, and it provides a simple programming model to link the pieces. For example, page authors with no programming expertise can use JavaServer Faces technology tags in a web page to link to server-side objects without writing any scripts.

Another important goal of JavaServer Faces technology is to leverage familiar component and web-tier concepts without limiting you to a particular scripting technology or markup language. JavaServer Faces technology APIs are layered directly on top of the Servlet API, as shown in the following diagram.

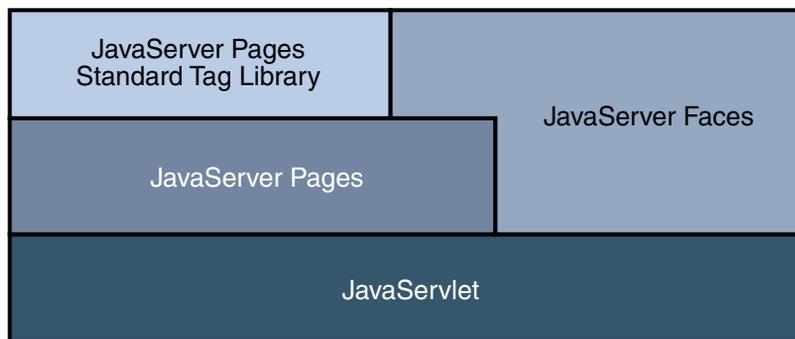


FIGURE 4-2 Java Web Application Technologies

This layering of APIs enables several important application use cases, such as using different presentation technologies, creating your own custom components directly from the component classes, and generating output for various client devices.

Facelets technology, available as part of JavaServer Faces 2.0, is now the preferred presentation technology for building JavaServer Faces based web applications and offers several advantages.

Facelets technology offers the advantages of code reuse and extensibility for components through `Templating` and `Composite Components` features.

When you use the `JavaServer Faces Annotations` feature, you can automatically register the backing bean as a resource available for JavaServer Faces applications. In addition, implicit navigation rules allow the developers to quickly configure page navigation. These features reduce the manual configuration process for applications.

For more information on Facelets technology features, see [Chapter 5, “Introduction to Facelets.”](#)

Most importantly, JavaServer Faces technology provides a rich architecture for managing component state, processing component data, validating user input, and handling events.

## Creating a Simple JavaServer Faces Application

JavaServer Faces technology provides an easy and user-friendly process for creating web applications.

Developing a simple JavaServer Faces application typically requires the following tasks:

- Developing backing beans
- Adding managed bean declarations
- Creating web pages using component tags
- Mapping the `FacesServlet` instance

In this section, the above tasks are described through the process of creating a simple JavaServer Faces Facelets application.

The example is a Hello application which includes a backing bean and a web page. When accessed by a client, the web page prints out a `Hello World` message. The example application is located in `tut-install/examples/web/hello` directory.

The tasks involved in developing this application can be examined by looking at the application in detail.

## Developing Backing Beans

As mentioned earlier in this chapter, a backing bean (a type of managed bean) is a `JavaBean` that is managed by JavaServer Faces. Components in a page are associated with backing beans which provide application logic. The example backing bean, `helloWorld.java`, contains the following code:

```
package Hello;

import javax.faces.bean.ManagedBean;

@ManagedBean
public class Hello{
    final String world = "Hello World!";

    public String getWorld()
    { return world; }

}
```

The example backing bean sets the value of the `world` variable with the string `Hello World!`. The `@ManagedBean` annotation registers the backing bean as a resource with the JavaServer Faces implementation. For more information on managed beans and annotations, see [Developing With JavaServer Faces Technology](#).

## Creating the Web Page

In a typical Facelets application, web pages are created in XHTML. The example web page, `beanhello.xhtml`, is a simple XHTML page. It contains the following content:

```
<html xmlns="http://www.w3.org/1999/xhtml">
  <head>
    <meta http-equiv="Content-Type" content="text/html; charset=UTF-8"/>
    <title>JavaServer Faces Hello World Application</title>
  </head>
  <body>
    #{hello.world}
  </body>
</html>
```

A Facelets XHTML web page can also contain several other elements which are covered later in this tutorial.

The web page connects to the backing bean through the Unified Expression Language (EL) value expression `#{hello.world}`, which retrieves the value of the `world` property from the backing bean `Hello`. Note the use of `hello` to reference the backing bean `Hello`. If no name is specified in the `@ManagedBean` annotation, the backing bean is always accessed with the first letter of the class name in lowercase.

For more information on using EL expressions, see [Chapter 6, “Unified Expression Language.”](#) For more information about Facelets technology, see [Introduction to Facelets](#). For more information about JavaServer Faces programming model and building web pages using JavaServer Faces technology, see [Chapter 7, “Using JavaServer Faces Technology in Web Pages.”](#)

## Mapping the Faces Servlet Instance

The final task requires mapping the Faces Servlet which is done through the web deployment descriptor (`web.xml`). A typical mapping of Faces Servlet is as follows:

```
<servlet>
  <servlet-name>Faces Servlet</servlet-name>
  <servlet-class>javax.faces.webapp.FacesServlet</servlet-class>
  <load-on-startup>1</load-on-startup>
</servlet>
<servlet-mapping>
  <servlet-name>Faces Servlet</servlet-name>
  <url-pattern>/faces/*</url-pattern>
</servlet-mapping>
```

The above file segment represents part of a typical JavaServer Faces web deployment descriptor. The web deployment descriptor can also contain other content relevant to a JavaServer Faces application configuration but that information is not covered here.

Mapping the Faces Servlet is automatically done for you when using a Java EE 6 server such as Sun GlassFish™ Enterprise Server v3.

## The Lifecycle of the `helloWorld` Application

Every web application has a lifecycle. Common tasks such as handling incoming requests, decoding parameters, modifying and saving state, and rendering web pages to the browser are all performed during a web application lifecycle. Some web application frameworks hide the details of the lifecycle from you while others require you to manage them manually.

By default, JavaServer Faces handles most of the lifecycle actions for you automatically. But it does expose the different parts of the request lifecycle, so that you can modify or perform different actions if your application requirements warrant it.

It is not necessary for the beginning user to understand the lifecycle of a JavaServer Faces application, but the information can be useful for creating more complex applications.

The lifecycle of a JavaServer Faces application starts and ends with the following activity: The client makes a request for the web page, and the server responds with the page. The lifecycle consists of two main phases: `execute` and `render`.

During the `execute` phase, several actions can take place: The application view is built or restored, the request parameter values are applied, conversions and validations are performed for component values, backing beans are updated with component values, and application logic is invoked. For a first (`initial`) request, only the view is built. For subsequent (`postback`) requests, some or all of the other actions can take place.

In the render phase, the requested view is rendered as response to the client. Rendering, typically is the process of generating output such as HTML or XHTML that can be read by the client (usually a browser).

The following short description of the example JavaServer Faces application passing through its lifecycle summarizes the activity that takes place behind the scenes.

The `helloWorld` example application goes through the following stages when it is deployed on the Enterprise Server:

1. When the `helloWorld` application is built and deployed on the Enterprise Server, the application is at an uninitiated state.
2. When a client makes a first (`initial`) request for the `beanHello.xhtml` web page, the `helloWorld` Facelets application is compiled.
3. The compiled Facelets application is executed and a new component tree (`UIViewRoot`) is constructed for the `helloWorld` application and is placed in the Faces Context.
4. The component tree is populated with the component and the backing bean property associated with it (represented by the EL expression `hello.world`).
5. A new view is built based on the component tree.
6. The view is rendered to the requesting client as a response.
7. The component tree is destroyed automatically.
8. On subsequent (`postback`) requests, the component tree is rebuilt and the saved state is applied.

For more detailed information on the JavaServer Faces lifecycle, see the JavaServer Faces Specification, Version 2.0 document.

## ▼ Running the Application in NetBeans IDE

To build, package, deploy, and run the JavaServer Faces `helloWorld` example using NetBeans IDE, follow these steps:

- 1 **In NetBeans IDE, select File→Open Project.**
- 2 **In the Open Project dialog box, navigate to the example directory:**  
*tut-install/examples/web*
- 3 **Select the `helloWorld` folder.**
- 4 **Select the Open as Main Project check box.**
- 5 **Click Open Project.**

**6 In the Projects tab, right-click the helloWorld project and select Run.**

This step compiles, assembles, and deploys the application, then brings up a web browser window displaying the following URL:

`http://localhost:8080/helloWorld`

**Example 4-1 Example Output of the helloWorld Application**

Hello World!

## Further Information about JavaServer Faces Technology

For more information on JavaServer Faces technology, see:

- The JavaServer Faces Technology web site:  
<http://java.sun.com/javaee/jaserverfaces/index.jsp>
- The JavaServer Faces 2.0 Technology download web site:  
<http://java.sun.com/javaee/jaserverfaces/download.html>
- Mojarra (JavaServer Faces 2.0) Release Notes:  
<https://jaserverfaces.dev.java.net/nonav/rlnotes/2.0.0/index.html>
- The JavaServer Faces 2.0 API and PDL documentation:  
<http://java.sun.com/javaee/jaserverfaces/reference/api/index.html>

# Introduction to Facelets

---

The term Facelets is used to refer to the JavaServer™ Faces View Definition Framework, which is a page declaration language that was developed for use with JavaServer Faces technology. As of JavaServer Faces 2.0, Facelets is a part of JavaServer Faces specification and also the preferred presentation technology for building JavaServer Faces based applications.

JavaServer Pages™ (JSP™) technology, previously used as the presentation technology for JavaServer Faces, does not support all of the new features available in JavaServer Faces 2.0. JSP is considered as a deprecated presentation technology for JavaServer Faces 2.0.

The following topics are addressed here:

- “Advantages of Facelets” on page 97
- “What's Facelets ?” on page 98
- “Developing a Simple Facelets Application” on page 99
- “Templating” on page 107
- “Composite Components” on page 109
- “Resources” on page 112

## Advantages of Facelets

Reuse of code and ease of development are important considerations for developers to adopt JavaServer Faces as the platform for large scale projects. By supporting these features, Facelets reduces the time and effort on development and deployment.

Facelets advantages include the following:

- Support for code reuse through templating and composite components
- Functional extensibility of components and other server-side objects through customization
- Faster compilation time
- Compile time EL validation

- High performance rendering

## What's Facelets ?

Facelets is a powerful but lightweight page declaration language that is used to build JavaServer Faces views using HTML style templates and to build component trees.

Facelets features include the following:

- Use of XHTML for creating web pages
- Support for Facelets Tag libraries in addition to JavaServer Faces and JSTL tag libraries
- Support for unified expression language
- Templating for components and pages

## Web Pages

Facelets views are usually created as XHTML pages. JavaServer Faces implementations support XHTML pages created in conformance with the XHTML Transitional DTD, as listed at [http://www.w3.org/TR/xhtml1/#a\\_dtd\\_XHTML-1.0-Transitional](http://www.w3.org/TR/xhtml1/#a_dtd_XHTML-1.0-Transitional).

By convention, web pages built with XHTML have an `.xhtml` extension.

## Tag Library Support

JavaServer Faces technology supports different tag libraries to add components to a web page. To support the JavaServer Faces tag library mechanism, Facelets uses XML namespace declarations.

The following table [Table 5–1](#) lists the tag libraries supported by Facelets.

**TABLE 5–1** Tag Libraries Supported by Facelets

Tag Library	URI	Prefix	Example	Contents
JavaServer Faces Facelets Tag Library	<a href="http://java.sun.com/jsf/facelets">http://java.sun.com/jsf/facelets</a>	ui:	ui:component ui:insert	Tags for templating

TABLE 5-1 Tag Libraries Supported by Facelets (Continued)

Tag Library	URI	Prefix	Example	Contents
JavaServer Faces HTML Tag Library	http://java.sun.com/jsf/html	h:	h:head h:body h:outputText h:inputText	JavaServer Faces component tags for all UIComponents.
JavaServer Faces Core Tag Library	http://java.sun.com/jsf/core	f:	f:actionListener f:attribute	Tags for JavaServer Faces custom actions that are independent of any particular RenderKit.
JSTL Core Tag Library	http://java.sun.com/jsp/jstl/core	c:	c:forEach c:catch	JSTL 1.1 Core Tags
JSTL Functions Tag Library	http://java.sun.com/jsp/jstl/functions	fn:	fn:toUpperCase fn:toLowerCase	JSTL 1.1 Functions Tags

In addition, Facelets also supports tags for composite components for which you can declare custom prefixes. For more information on composite components, see [“Composite Components” on page 109](#).

## Unified Expression Language Support

Based on the JavaServer Faces support for unified expression language (EL) syntax defined by JSP 2.1, Facelets uses EL expressions to reference properties and methods of backing beans. EL expressions can be used to bind component objects or values to managed-bean methods or managed-bean properties. For more information on using EL expressions, see [“Using the EL to Reference Backing Beans” on page 173](#).

## Developing a Simple Facelets Application

This section describes the general steps involved in developing a JavaServer Faces application.

Developing a simple JavaServer Faces application, using Facelets technology usually requires these tasks:

- Developing the backing beans
- Creating the pages using the component tags
- Defining page navigation
- Mapping the FacesServlet instance

- Adding managed bean declarations

In the next section some of the above tasks are described in more detail.

## Creating a Facelets Application

The example used in this tutorial is the `guessnumber` application. The application presents you with a page that asks you to guess a number between 0 and 10, validates your input against a random number, and responds with another page that informs you, if you guessed the number correctly or incorrectly.

### Developing a Backing Bean

In a typical JavaServer Faces application each page in the application connects to a backing bean (a type of managed bean). The backing bean defines the methods and properties that are associated with the components.

The following managed bean class, `UserNumberBean.java`, generates a random number between 0 and 10:

```
package guessNumber;

import java.util.Random;
import javax.faces.bean.ManagedBean;
import javax.faces.bean.SessionScoped;

@ManagedBean
@SessionScoped
public class UserNumberBean {
    Integer randomInt = null;
    Integer userNumber = null;
    String response = null;
    private long maximum=10;
    private long minimum=0;

    public UserNumberBean() {
        Random randomGR = new Random();
        randomInt = new Integer(randomGR.nextInt(10));
        System.out.println("Duke's number: " + randomInt);
    }
    public void setUserNumber(Integer user_number) {
        userNumber = user_number;
    }

    public Integer getUserNumber() {
        return userNumber;
    }
}
```

```

    }

    public String getResponse() {
        if ((userNumber != null) && (userNumber.compareTo(randomInt) == 0))
        {
            return "Yay! You got it!";
        }
        else
        {
            return "Sorry, " + userNumber + " is incorrect.";
        }
    }

    public long getMaximum() {
        return (this.maximum);
    }

    public void setMaximum(long maximum) {
        this.maximum = maximum;
    }

    public long getMinimum() {
        return (this.minimum);
    }

    public void setMinimum(long minimum) {
        this.minimum = minimum;
    }
}

```

Note the use of the `@ManagedBean` annotation which registers the backing bean as a resource with JavaServer Faces implementation. The `@SessionScoped` annotation registers the bean scope as session.

## Creating Facelets Views

Creating a page or view is the responsibility of a page author. This task involves adding components on the pages, wiring the components to backing bean values and properties, and registering converters, validators, or listeners onto the components.

For the example application, XHTML web pages serve as the front end. The first page of the example application is a page called `greeting.xhtml`. A closer look at various sections of this web page provides more information.

The first section of the web page declares the content type for the page, which is XHTML:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE html
  PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
```

The next section declares the XML namespace for the tag libraries that are used in the web page:

```
<html xmlns="http://www.w3.org/1999/xhtml"
      xmlns:f="http://java.sun.com/jsf/core"
      xmlns:h="http://java.sun.com/jsf/html">
```

The next section uses various tags to insert components into the web page:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
      xmlns:h="http://java.sun.com/jsf/html"
      xmlns:f="http://java.sun.com/jsf/core">

<h:head>
  <title>Facelets Guess Number Application</title>
</h:head>

<h:body>
  <h:form>
    <h:graphicImage value="#{resource['images:wave.med.gif']}" />
  <h2>
    Hi,
    <p>My name is Duke. I am thinking of a number between <b>
      #{userNumberBean.minimum} and #{userNumberBean.maximum}.
    </b> Can you guess it ?</p>

    <h:inputText
      id="userNo"
      value="#{userNumberBean.userNumber}">
      <f:validateLongRange
        minimum="#{userNumberBean.minimum}"
        maximum="#{userNumberBean.maximum}" />
    </h:inputText>

    <h:commandButton id="submit" value="Submit" action="response.xhtml" />
    <h:message showSummary="true" showDetail="false"
      style="color: red;
      font-family: 'New Century Schoolbook', serif;
      font-style: oblique;
      text-decoration: overline"
      id="errors1"
      for="userNo" />
```

```

        </h2>
    </h:form>
</h:body>
</html>

```

Note the use of the Facelets HTML tags to add components, and the Facelets core tag to validate the user input. An `inputText` component accepts user input and sets the value of the backing bean property `userNumber` through the EL expression `#{userNumberBean.userNumber}`. The input value is validated for value range by the JavaServer Faces standard validator `f:validateLongRange`.

The image file `wave.med.gif`, is added to the page as a resource. For more details about the resources facility, see [“Resources” on page 112](#).

The `submit` command button starts validation of the input data. Using implicit navigation, it redirects the client to another page `response.xhtml`, which shows the response to your input.

You can now create the second page, `response.xhtml`, with the following content:

```

<!DOCTYPE html
    PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">

<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:h="http://java.sun.com/jsf/html">

<h:head>
    <title>Guess Number Facelets Application</title>
</h:head>
<h:body>
    <h:form>
        <h:graphicImage value="#{resource['images:wave.med.gif']}" />
    <h2>
        <h:outputText id="result" value="#{userNumberBean.response}" />

    </h2>

        <h:commandButton id="back" value="Back" action="greeting.xhtml" />

    </h:form>
</h:body>
</html>

```

## Configuring the Application

Configuring a JavaServer Faces application involves various configuration tasks which include adding managed-bean declarations, navigation rules and resources bundle declarations in the

application configuration resource files such as `faces-config.xml`, and mapping the Faces Servlet in the web deployment descriptor file such as a `web.xml` file. Application configuration is an advanced topic covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

If you are using an IDE such as NetBeans IDE, a web deployment descriptor is automatically created for you. In such IDE created `web.xml` files, change the default greeting page which is `index.xhtml`, to `greeting.xhtml`. Here is an example `web.xml` file:

```
<?xml version="1.0" encoding="UTF-8"?>
<web-app version="3.0" xmlns="http://java.sun.com/xml/ns/javaee"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
http://java.sun.com/xml/ns/javaee/web-app_3_0.xsd">
  <context-param>
    <param-name>javax.faces.PROJECT_STAGE</param-name>
    <param-value>Development</param-value>
  </context-param>
  <servlet>
    <servlet-name>Faces Servlet</servlet-name>
    <servlet-class>javax.faces.webapp.FacesServlet</servlet-class>
    <load-on-startup>1</load-on-startup>
  </servlet>
  <servlet-mapping>
    <servlet-name>Faces Servlet</servlet-name>
    <url-pattern>/faces/*</url-pattern>
  </servlet-mapping>
  <session-config>
    <session-timeout>
      30
    </session-timeout>
  </session-config>
  <welcome-file-list>
    <welcome-file>faces/greeting.xhtml</welcome-file>
  </welcome-file-list>
</web-app>
```

Note the use of parameter `PROJECT_STAGE`. `ProjectStage` is a context parameter identifying the status of a JavaServer Faces application in the software lifecycle.

The stage of an application can affect the behavior of the application. For example, if the project stage is defined as `Development`, debugging information is automatically generated for the user. If not defined by the user, the default project stage is considered as `Production`. Project Stage is covered in more detail in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

# Building, Packaging, Deploying and Running the Application

The example Facelets application described in this chapter can be built, packaged, and deployed using the Java EE 6 SDK with NetBeans IDE. For details on how to obtain this software and configure your environment to run the examples, see [Chapter 2, “Using the Tutorial Examples.”](#) The source code for this example is also available in the *tut-install/examples/web/guessnumber* directory.

## ▼ To Create the Example Facelets Application with NetBeans IDE

To create the example Facelets project, use the following procedure.

- 1 **In NetBeans IDE, from the File menu, choose New Project.**

The New Project wizard opens.

- 2 **In the wizard, select Java Web as the category and Web Application as the project type and click Next.**

The New Web Application wizard opens.

- 3 **In the Project Name field, type `guessNumber`, and click Next.**

- 4 **In the Server and Settings page, select Server as GlassFish v3 from the Server menu, select Java EE version as Java EE 6 Web from the Java EE version menu, and then click Next.**

- 5 **In the Frameworks page, select the JavaServer Faces checkbox and click Finish.**

A new Project is created and is available in the Projects window. A default file, `index.xhtml`, is created and opened in the Editor.

## ▼ To Create the Application

- 1 **Right-click the Project node, and select New→Java package.**

- 2 **In the Package Name field, type `guessNumber` and click Finish.**

A new package is created and placed under Source Packages node of the Project.

- 3 **Right-click the Source Packages node and select New→Java Class.**

- 4 **Type the name of the class file as `UserNumberBean`, select the name of package as `guessNumber` and click Finish.**

A new Java class file is created and opened in the IDE.

- 5 **Replace the content of the Java class file with the example code from the `UserNumberBean.java` file listed in “[Developing a Backing Bean](#)” on page 100, and save the file.**
- 6 **Create two new XHTML pages and name them `greeting.xhtml` and `response.xhtml` respectively:**
  - a. **Right-click the project node and choose `New`→`Other`.**  
The New File wizard opens.
  - b. **Choose Category as `Web` and then File Type as `XHTML` and click `Next`.**
  - c. **Enter `greeting.xhtml` in the XHTML File name field and click `Finish`.**  
A new XHTML web page is created and placed under `Web Pages` node.
  - d. **Repeat the above steps but enter the name of file as `response.xhtml` to create a second web page.**
- 7 **Edit the XHTML files and add Facelets content to them:**
  - a. **Replace the content of `greeting.xhtml` with the example `greeting.xhtml` code listed in “[Creating Facelets Views](#)” on page 101 and save the file.**
  - b. **Similarly replace the content of `response.xhtml` with the example `response.xhtml` code and save the file.**
- 8 **Add Duke's image as part of the application by copying the `wave.med.gif` image file from the tutorial example and saving it as a resource.**
  - a. **Create a folder named `resources` under `Web Pages`.**
  - b. **Create a subfolder, `images` under `resources` folder.**
  - c. **Save the `wave.med.gif` image in `resources/images` folder.**
- 9 **Edit the `web.xml` file to modify the welcome page to `greeting.html`.**
- 10 **Right-click the Project Node and select `Build` from the menu, to compile and build the application.**
- 11 **Right-click the Project Node and select `Deploy`, to deploy the application to Sun GlassFish™ Enterprise Server v3.**
- 12 **Access the application by typing the following URL in the browser:**  
`http://localhost:8080/guessNumber`

# Templating

JavaServer Faces 2.0 provides the tools to implement user interfaces that are easy to extend and reuse. Templating is a useful feature available with Facelets that allows you to create a page that will act as the base or template for the other pages in an application. By using templates, you can reuse code and avoid recreating similarly constructed pages. Templating also helps in maintaining a standard look and feel in an application with a large number of pages.

The following table lists Facelets tags that are used for templating and their respective functionality:

TABLE 5-2 Facelets Templating Tags

Tag	Function
<code>ui:component</code>	Defines a component that is created and added to the component tree.
<code>ui:composition</code>	Defines a page composition that optionally uses a template. Content outside of this tag is ignored.
<code>ui:debug</code>	Defines a debug component that is created and added to the component tree.
<code>ui:define</code>	Defines content that is inserted into a page by a template
<code>ui:decorate</code>	Similar to composition tag but does not disregard content outside this tag.
<code>ui:fragment</code>	Similar to component tag but does not disregard content outside this tag.
<code>ui:include</code>	Encapsulate and reuse content for multiple pages.
<code>ui:insert</code>	Inserts content into a template.
<code>ui:param</code>	Used to pass parameters to an included file.
<code>ui:repeat</code>	Used as an alternative for loop tags such as <code>c:forEach</code> or <code>h:dataTable</code> .
<code>ui:remove</code>	Removes content from a page.

For more information on Facelets templating tags, see the PDL at <http://java.sun.com/javaee/javaserverfaces/2.0/docs/pdl/docs/facelets/index.html>.

The Facelets tag library includes the main templating tag `<ui:insert>`. A template page is created with this tag, it allows defining a default structure for a page. A template page can be reused as a template for other pages, usually referred to as a `client` pages.

Here is an example of a template saved as `template.xhtml`:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
      xmlns:ui="http://java.sun.com/jsf/facelets"
      xmlns:h="http://java.sun.com/jsf/html">

  <h:head>
    <meta http-equiv="Content-Type" content="text/html; charset=UTF-8" />
    <link href="./resources/css/default.css" rel="stylesheet" type="text/css" />
    <link href="./resources/css/cssLayout.css" rel="stylesheet" type="text/css" />
    <title>Facelets Template</title>
  </h:head>

  <h:body>
    <div id="top" class="top">
      <ui:insert name="top">Top Section</ui:insert>
    </div>
    <div>
      <div id="left">
        <ui:insert name="left">Left Section</ui:insert>
      </div>
      <div id="content" class="left_content">
        <ui:insert name="content">Main Content</ui:insert>
      </div>
    </div>
  </h:body>

</html>
```

The example page defines a HTML page that is divided into 3 sections, a top section, a left section and a main section. The sections have stylesheets associated with them. The same structure can be reused for the other pages of the application.

The client page invokes the template by using `<ui:composition>` tag. In the following example, a client page named `templateclient.xhtml`, invokes the template page from the preceding example named `template.xhtml`. A client page allows content to be inserted with the help of the `<ui:define>` tag.

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
      xmlns:ui="http://java.sun.com/jsf/facelets"
      xmlns:h="http://java.sun.com/jsf/html">

  <h:body>

    <ui:composition template="./template.xhtml">
      <ui:define name="top">
```

```
        Welcome to Template Client Page
    </ui:define>

    <ui:define name="left">
        <h:outputLabel value="You are in the Left Section"/>
    </ui:define>

    <ui:define name="content">
        <h:graphicImage value="#{resource['images:wave.med.gif']}/>
        <h:outputText value="You are in the Main Content Section"/>
    </ui:define>

</ui:composition>

</h:body>
</html>
```

You can use the NetBeans IDE to create Facelets template and client pages. For more information on creating these pages, see <http://netbeans.org/kb/docs/web/jsf20-intro.html>.

## Composite Components

The JavaServer Faces offers the concept of composite components with Facelets. A composite component can be considered a special type of template that acts as a component.

Any component essentially is a piece of reusable code that is capable of a certain functionality. For example, an `inputText` component is capable of accepting user input. A component also has validators, converters, and listeners attached to it to perform certain defined actions.

A composite component is a component that consists of a collection of markups and other existing components. It is a reusable, user-created component that is capable of a customized, defined functionality and can have validators, converters and listeners attached to it like any other JavaServer Faces component.

With Facelets, any XHTML page that is inserted with markups and other components, can be converted into a composite component. Using the resources facility, the composite component can be stored in a library that is available to the application from the defined resources location.

The following table lists the most commonly used composite tags and their functions:

TABLE 5-3 Composite Component Tags

Tag	Function
<code>composite:interface</code>	Declares the usage contract for a composite component. The composite component can be used as a single component whose feature set is the union of the features declared in the usage contract.
<code>composite:implementation</code>	Defines the implementation of the composite component. If a <code>&lt;composite:interface&gt;</code> element appears, there must be a corresponding <code>&lt;composite:implementation&gt;</code> .
<code>composite:attribute</code>	Declares an attribute that may be given to an instance of the composite component, in which this tag is declared.
<code>composite:insertChildren</code>	Any child components or template text within the composite component tag in the using page will be re-parented into the composite component at the point indicated by this tag's placement within the <code>composite:implementation</code> section.
<code>composite:valueHolder</code>	Declares that the composite component whose contract is declared by the <code>composite:interface</code> in which this element is nested exposes an implementation of <code>ValueHolder</code> suitable for use as the target of attached objects in the using page.
<code>composite:editableValueHolder</code>	Declares that the composite component whose contract is declared by the <code>composite:interface</code> in which this element is nested exposes an implementation of <code>EditableValueHolder</code> suitable for use as the target of attached objects in the using page.
<code>composite:actionSource</code>	Declares that the composite component whose contract is declared by the <code>composite:interface</code> in which this element is nested exposes an implementation of <code>ActionSource2</code> suitable for use as the target of attached objects in the using page.

For more information and a complete list of Facelets composite tags, see the PDL

at <http://java.sun.com/>

[javaee/javaserverfaces/2.0/docs/pdl/docs/facelets/index.html](http://javaee/javaserverfaces/2.0/docs/pdl/docs/facelets/index.html).

The following example shows a composite component that accepts an email address as input:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
xmlns:composite="http://java.sun.com/jsf/composite"
xmlns:h="http://java.sun.com/jsf/html">

<h:head>
<title>This content will not be displayed
</title>
```

```

</h:head>
<h:body>

<composite:interface>
<composite:attribute name="value" required="false"/>
</composite:interface>

<composite:implementation>
  <h:outputLabel value="Email id: ">
  </h:outputLabel>
  <h:inputText value="#{cc.attrs.value}">
  </h:inputText>
</composite:implementation>

</h:body>
</html>

```

Note the use of `cc.attrs.value` when defining the value of the `inputText` component. The word `cc` in JavaServer Faces is a reserved word for composite components. The `#{cc.attrs.ATTRIBUTE_NAME}` expression is used to access the attributes defined for the composite component's interface which in this case happens to be `value`.

The preceding example content is stored as a file named `email.xhtml`, in a folder named `resources/emcomp` under the application web root directory. This directory is considered a library by the JavaServer Faces, and a UI component can be accessed from such library. For more information on resources, see [“Resources” on page 112](#).

The web page that uses this composite component is generally called a using page. The using page includes a reference to the composite component, in the `xml` namespace declarations:

```

<!DOCTYPE html
PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
xmlns:h="http://java.sun.com/jsf/html"
xmlns:em="http://java.sun.com/jsf/composite/emcomp">

<h:head>
<title>Using a sample composite component</title>
</h:head>

<body>
<h:form>
<em:email value="Enter your email id" />

</h:form>
</body>
</html>

```

The local composite component library is defined in the `xml` namespace with the declaration `xmlns:em="http://java.sun.com/jsf/composite/emcomp/".` the component it self is accessed through the use of the tag `em:email`. The preceding example content can be stored as a web page named `emuserpage.xhtml` under web root directory. When compiled and deployed on a server it can be accessed with the following URL:

```
http://localhost:8080/<application_name>/faces/emuserpage.xhtml
```

## Resources

Resources refers to any software artifacts that the application requires for proper rendering. They include images, script files and any user-created component libraries. As of JavaServer Faces 2.0, resources must be collected in a standard location, which can be one of the following:

- A resource packaged in the web application root must be in a subdirectory of a resources directory at the web application root: `resources/<resource-identifier>`.
- A resource packaged in the web application's `classpath` must be in a subdirectory of the `META-INF/resources` directory within a web application: `META-INF/resources/<resource-identifier>`.

The JavaServer Faces runtime will look for the resources in the above listed locations, in that order.

Resource identifiers are unique strings that conform to the following format:

```
[localePrefix/][libraryName/][libraryVersion/]resource name[/resourceVersion]
```

Elements of the resource identifier in brackets ( `[]` ) are optional. This indicates that only a resource name is a required element, which is usually a file name.

Resources can be considered as a library location. Any artifacts, like a composite component or template that is stored under resources directory, becomes accessible to the other components of the application which can use them to create a resource instance.

# Unified Expression Language

---

This chapter introduces the Unified Expression Language (also referred to as EL) which provides an important mechanism for enabling the presentation layer (web pages) to communicate with the application logic (backing beans). EL is used by JavaServer Pages™ (JSP™) as well as JavaServer™ Faces technologies.

Introduced as a primary feature of JSP 2.1, the EL represents a union of the expression language offered by JSP 2.0 and the expression language created for JavaServer Faces technology.

## Overview of EL

The unified expression language allows page authors to use simple expressions to dynamically access data from JavaBeans™ components. For example, the `test` attribute of the following conditional tag is supplied with an EL expression that compares the number of items in the session-scoped bean named `cart` with 0.

```
<c:if test="${sessionScope.cart.numberOfItems > 0}">
  ...
</c:if>
```

JavaServer Faces technology uses EL for the following functions:

- Deferred and immediate evaluation of expressions
- The ability to set as well as get data
- The ability to invoke methods

See “[Using the EL to Reference Backing Beans](#)” on page 173 for more information on how to use the EL in JavaServer Faces applications.

To summarize, the unified expression language provides a way to use simple expressions to perform the following tasks:

- Dynamically read application data stored in JavaBeans components, various data structures, and implicit objects
- Dynamically write data, such as user input into forms, to JavaBeans components
- Invoke arbitrary static and public methods
- Dynamically perform arithmetic operations

The EL is also used to specify the following kinds of expressions that a custom tag attribute will accept:

- *Immediate evaluation expressions* or *deferred evaluation expressions*. An immediate evaluation expression is evaluated immediately by the underlying technology such as JavaServer Faces. A deferred evaluation expression can be evaluated later by the underlying technology using the expression language.
- *Value expression* or *method expression*. A value expression references data, whereas a method expression invokes a method.
- *Rvalue expression* or *lvalue expression*. An rvalue expression can only read a value, whereas an lvalue expression can both read and write that value to an external object.

Finally, the EL provides a pluggable API for resolving expressions so custom resolvers that can handle expressions not already supported by the EL can be implemented.

This section gives an overview of the unified expression language features by explaining the following topics:

- [“Immediate and Deferred Evaluation Syntax” on page 114](#)
- [“Value and Method Expressions” on page 116](#)
- [“Defining a Tag Attribute Type” on page 122](#)
- [“Literal Expressions” on page 123](#)
- [“Operators” on page 124](#)
- [“Reserved Words” on page 125](#)
- [“Examples of EL Expressions” on page 125](#)

In addition to the above, JSP technology-related topics such as implicit objects and functions are also relevant to EL users but they are not covered in this tutorial.

## Immediate and Deferred Evaluation Syntax

The EL supports both immediate and deferred evaluation of expressions. *Immediate evaluation* means that the expression is evaluated and the result is returned immediately when the page is first rendered. *Deferred evaluation* means that the technology using the expression language can employ its own machinery to evaluate the expression sometime later during the page’s lifecycle, whenever it is appropriate to do so.

Those expressions that are evaluated immediately use the `{}` syntax. Expressions whose evaluation is deferred use the `#{}`  syntax.

Because of its multiphase lifecycle, JavaServer Faces technology mostly uses deferred evaluation expressions. During the lifecycle, component events are handled, data is validated, and other tasks are performed in a particular order. Therefore, a JavaServer Faces implementation must defer evaluation of expressions until the appropriate point in the lifecycle.

Other technologies using the EL might have different reasons for using deferred expressions.

## Immediate Evaluation

All expressions using the `{}` syntax are evaluated immediately. These expressions can only be used within template text or as the value of a tag attribute that can accept runtime expressions.

The following example shows a tag whose value attribute references an immediate evaluation expression that gets the total price from the session-scoped bean named `cart`:

```
<fmt:formatNumber value="{sessionScope.cart.total}"/>
```

The JavaServer Faces implementation evaluates the expression, `{sessionScope.cart.total}`, converts it, and passes the returned value to the tag handler.

Immediate evaluation expressions are always read-only value expressions. The example expression shown above can only get the total price from the `cart` bean; it cannot set the total price.

## Deferred Evaluation

Deferred evaluation expressions take the form `#{expr}` and can be evaluated at other phases of a page lifecycle as defined by whatever technology is using the expression. In the case of JavaServer Faces technology, its controller can evaluate the expression at different phases of the lifecycle depending on how the expression is being used in the page.

The following example shows a JavaServer Faces `inputText` tag, which represents a text field component into which a user enters a value. The `inputText` tag's `value` attribute references a deferred evaluation expression that points to the `name` property of the customer bean.

```
<h:inputText id="name" value="#{customer.name}" />
```

For an initial request of the page containing this tag, the JavaServer Faces implementation evaluates the `#{customer.name}` expression during the render response phase of the lifecycle. During this phase, the expression merely accesses the value of `name` from the customer bean, as is done in immediate evaluation.

For a postback request, the JavaServer Faces implementation evaluates the expression at different phases of the lifecycle, during which the value is retrieved from the request, validated, and propagated to the customer bean.

As shown in this example, deferred evaluation expressions can be value expressions that can be used to both read and write data. They can also be method expressions. Value expressions (both immediate and deferred) and method expressions are explained in the next section.

## Value and Method Expressions

The EL defines two kinds of expressions: value expressions and method expressions. Value expressions can either yield a value or set a value. Method expressions reference methods that can be invoked and can return a value.

### Value Expressions

Value expressions can be further categorized into rvalue and lvalue expressions. *Rvalue expressions* are those that can read data, but cannot write it. *Lvalue expressions* can both read and write data.

All expressions that are evaluated immediately use the `{ }` delimiters and are always rvalue expressions. Expressions whose evaluation can be deferred use the `# { }` delimiters and can act as both rvalue and lvalue expressions. Consider the following two value expressions:

```
${customer.name}
```

```
#{customer.name}
```

The former uses immediate evaluation syntax, whereas the latter uses deferred evaluation syntax. The first expression accesses the `name` property, gets its value, adds the value to the response, and gets rendered on the page. The same can happen with the second expression. However, the tag handler can defer the evaluation of this expression to a later time in the page lifecycle, if the technology using this tag allows.

In the case of JavaServer Faces technology, the latter tag's expression is evaluated immediately during an initial request for the page. In this case, this expression acts as an rvalue expression. During a postback request, this expression can be used to set the value of the `name` property with user input. In this case, the expression acts as an lvalue expression.

### Referencing Objects Using Value Expressions

Both rvalue and lvalue expressions can refer to the following objects and their properties or attributes:

- JavaBeans components
- Collections
- Java™ SE enumerated types
- Implicit objects

To refer to these objects, you write an expression using a variable which is the name of the object. The following expression references a backing bean (a JavaBeans component) called `customer`.

```
${customer}
```

The web container evaluates the variable that appears in an expression by looking up its value according to the behavior of `PageContext.findAttribute(String)`, where the `String` argument is the name of the variable. For example, when evaluating the expression `${customer}`, the container will look for `customer` in the page, request, session, and application scopes and will return its value. If `customer` is not found, `null` is returned.

You can alter the way variables are resolved with a custom EL resolver. For instance, you can provide an EL resolver that intercepts objects with the name `customer`, so that `${customer}` returns a value in the EL resolver instead. Creation of custom EL resolvers is an advanced topic covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

To reference an enum constant with an expression, use a `String` literal. For example, consider this Enum class:

```
public enum Suit {hearts, spades, diamonds, clubs}
```

To refer to the `Suit` constant, `Suit.hearts` with an expression, you use the `String` literal, "hearts". Depending on the context, the `String` literal is converted to the enum constant automatically. For example, in the following expression in which `mySuit` is an instance of `Suit`, "hearts" is first converted to `Suit.hearts` before it is compared to the instance.

```
${mySuit == "hearts"}
```

## Referring to Object Properties Using Value Expressions

To refer to properties of a bean or an Enum instance, items of a collection, or attributes of an implicit object, you use the `.` or `[]` notation, which is similar to the notation used by ECMAScript language. For more information on ECMAScript, see <http://www.ecmascript.org>.

If you want to reference the name property of the `customer` bean, use either the expression `${customer.name}` or the expression `${customer["name"]}`. The part inside the square brackets is a `String` literal that is the name of the property to reference.

You can use double or single quotes for the `String` literal. You can also combine the `[]` and `.` notations, as shown here:

```
${customer.address["street"]}
```

Properties of an enum can also be referenced in this way. However, as with JavaBeans component properties, the properties of an Enum class must follow JavaBeans component

conventions. This means that a property must at least have an accessor method called `get<Property>` where `<Property>` is the name of the property which can be referenced by an expression.

For example, consider an Enum class that encapsulates the names of the planets of our galaxy and includes a method to get the mass of a planet. You can use the following expression to reference the method `getMass` of the `Planet` Enum class:

```
${myPlanet.mass}
```

If you are accessing an item in an array or list, you must use either a literal value that can be converted to `int` or the `[]` notation with an `int` and without quotes. The following examples could all resolve to the same item in a list or array, assuming that `socks` can be converted to `int`:

- `${customer.orders[1]}`
- `${customer.orders.socks}`

In contrast, an item in a `Map` can be accessed using a string literal key; no coercion is required:

```
${customer.orders["socks"]}
```

An `rvalue` expression also refers directly to values that are not objects, such as the result of arithmetic operations and literal values, as shown by these examples:

- `${"literal"}`
- `${customer.age + 20}`
- `${true}`
- `${57}`

The unified expression language defines the following literals:

- Boolean: `true` and `false`
- Integer: as in Java
- Floating point: as in Java
- String: with single and double quotes; `"` is escaped as `\`, `'` is escaped as `\'`, and `\` is escaped as `\\`
- Null: `null`

You can also write expressions that perform operations on an enum constant. For example, consider the following Enum class:

```
public enum Suit {club, diamond, heart, spade }
```

After declaring an enum constant called `mySuit`, you can write the following expression to test if `mySuit` is `spade`:

```
${mySuit == "spade"}
```

When the EL resolving mechanism resolves this expression, it will invoke the `valueOf` method of the `Enum` class with the `Suit` class and the `spade` type, as shown here:

```
mySuit.valueOf(Suit.class, "spade")
```

## Where Value Expressions Can Be Used

Value expressions using the `{}` delimiters can be used in the following places:

- In static text
- In any standard or custom tag attribute that can accept an expression

The value of an expression in static text is computed and inserted into the current output. Here is an example of an expression embedded in static text:

```
<some:tag>
  some text ${expr} some text
</some:tag>
```

If the static text appears in a tag body, note that an expression *will not* be evaluated if the body is declared to be tag-dependent.

Lvalue expressions can only be used in tag attributes that can accept lvalue expressions.

There are three ways to set a tag attribute value using either an rvalue or lvalue expression:

- With a single expression construct:

```
<some:tag value="${expr}"/>
<another:tag value="#{expr}"/>
```

These expressions are evaluated and the result is converted to the attribute's expected type.

- With one or more expressions separated or surrounded by text:

```
<some:tag value="some${expr}${expr}text${expr}"/>
<another:tag value="some#{expr}#{expr}text#{expr}"/>
```

These kinds of expression are called *composite expressions*. They are evaluated from left to right. Each expression embedded in the composite expression is converted to a `String` and then concatenated with any intervening text. The resulting `String` is then converted to the attribute's expected type.

- With text only:

```
<some:tag value="sometext"/>
```

This expression is called a *literal expression*. In this case, the attribute's `String` value is converted to the attribute's expected type. Literal value expressions have special syntax rules. See [“Literal Expressions” on page 123](#) for more information. When a tag attribute has an

enum type, the expression that the attribute uses must be a literal expression. For example, the tag attribute can use the expression "hearts" to mean `Suit.hearts`. The literal is converted to `Suit` and the attribute gets the value `Suit.hearts`.

All expressions used to set attribute values are evaluated in the context of an expected type. If the result of the expression evaluation does not match the expected type exactly, a type conversion will be performed. For example, the expression `#{1.2E4}` provided as the value of an attribute of type `float` will result in the following conversion:

```
Float.valueOf("1.2E4").floatValue()
```

See Section 1.18 of the *JavaServer Pages 2.1 Expression Language Specification* (available from <http://jcp.org/aboutJava/communityprocess/final/jsr245/>) for the complete type conversion rules.

## Method Expressions

Another feature of the unified expression language is its support of deferred method expressions. A method expression is used to invoke an arbitrary public method of a bean, which can return a result.

In JavaServer Faces technology, a component tag represents a component on a page. The component tag uses method expressions to invoke methods that perform some processing for the component. These methods are necessary for handling events that the components generate and validating component data, as shown in this example:

```
<h:form>
  <h:inputText
    id="name"
    value="#{customer.name}"
    validator="#{customer.validateName}"/>
  <h:commandButton
    id="submit"
    action="#{customer.submit}" />
</h:form>
```

The `inputText` tag displays as a text field. The `validator` attribute of this `inputText` tag references a method, called `validateName`, in the bean, called `customer`.

Because a method can be invoked during different phases of the lifecycle, method expressions must always use the deferred evaluation syntax.

Like lvalue expressions, method expressions can use the `.` and the `[]` operators. For example, `#{object.method}` is equivalent to `#{object["method"]}`. The literal inside the `[]` is converted to `String` and is used to find the name of the method that matches it. Once the method is found, it is invoked or information about the method is returned.

Method expressions can be used only in tag attributes and only in the following ways:

- With a single expression construct, where `bean` refers to a JavaBeans component and `method` refers to a method of the JavaBeans component:

```
<some:tag value="#{bean.method}"/>
```

The expression is evaluated to a method expression, which is passed to the tag handler. The method represented by the method expression can then be invoked later.

- With text only:

```
<some:tag value="sometext"/>
```

Method expressions support literals primarily to support `action` attributes in JavaServer Faces technology. When the method referenced by this method expression is invoked, it returns the `String` literal, which is then converted to the expected return type, as defined in the tag's TLD.

## Parameterized Method Calls

The updated EL version 2.1.2 included in Java EE 6 offers support for parameters to method calls. Method calls can now use parameters (or arguments) without having to use static EL functions.

Both the `.` and `[]` operators can be used for invoking method calls with parameters as shown in expression syntax below:

- `expr-a[expr-b](parameters)`
- `expr-a.identifier-b(parameters)`

In the first expression syntax, `expr-a` is evaluated to represent a bean object. The expression `expr-b` is evaluated and cast to a string which represents a method in the bean represented by `expr-a`. In the second expression syntax, `expr-a` is evaluated to represent a bean object and `identifier-b` is a string that represents a method in the bean object. The parameters in parentheses are the arguments for the method invocation. Parameters can be 0 or more values or expressions, separated by commas.

Parameters are supported for both value expressions and method expressions. In the following example, which is a modified tag from `guessNumber` application, a random number is provided as an argument rather than from user input to the method call:

```
<h:inputText value="#{userNumberBean.userNumber('5')}"/>
```

The above example uses a value expression.

Consider the following example of a JavaServer Faces component tag which uses a method expression:

```
<h:commandButton action="#{trader.buy}" value="buy"/>
```

where EL expression `trader.buy` is calling the trader bean's `buy` method. You can modify the tag to pass on a parameter. Here is the revised tag where a parameter is passed:

```
<h:commandButton action="#{trader.buy('JAVA')}" value="buy"/>
```

In the above example you are passing the string 'JAVA' (a stock symbol) as a parameter to the `buy` method.

For more information on the updated EL, see <https://uel.dev.java.net>.

## Defining a Tag Attribute Type

As explained in the previous section, all kinds of expressions can be used in tag attributes. Which kind of expression and how that expression is evaluated (whether immediately or deferred) is determined by the type attribute of the tag's definition in the Page Description Language (PDL) that defines the tag.

If you plan to create custom tags, for each tag in the PDL, you need to specify what kind of expression to accept. [Table 6-1](#) shows the three different kinds of tag attributes that accept EL expressions, gives examples of expressions they accept, and the type definitions of the attributes that must be added to the PDL. You cannot use `#{}` syntax for a dynamic attribute, meaning an attribute that accepts dynamically-calculated values at runtime. Similarly, you also cannot use the `${}` syntax for a deferred attribute.

TABLE 6-1 Definitions of Tag Attributes That Accept EL Expressions

Attribute Type	Example Expression	Type Attribute Definition
dynamic	"literal"	<rtexprvalue>true</rtexprvalue>
	<code>\${literal}</code>	<rtexprvalue>true</rtexprvalue>
deferred value	"literal"	<deferred-value> <type>java.lang.String</type> </deferred-value>
	<code>#{customer.age}</code>	<deferred-value> <type>int</type> </deferred-value>

TABLE 6-1 Definitions of Tag Attributes That Accept EL Expressions (Continued)

Attribute Type	Example Expression	Type Attribute Definition
deferred method	"literal"	<pre>&lt;deferred-method&gt;   &lt;method-signature&gt;     java.lang.String submit()   &lt;/method-signature&gt; &lt;/deferred-method&gt;</pre>
	<code>#{customer.calcTotal}</code>	<pre>&lt;deferred-method&gt;   &lt;method-signature&gt;     double calcTotal(int, double)   &lt;/method-signature&gt; &lt;/deferred-method&gt;</pre>

In addition to the tag attribute types shown in [Table 6-1](#), you can also define an attribute to accept both dynamic and deferred expressions. In this case, the tag attribute definition contains both an `rtexprvalue` definition set to `true` and either a `deferred-value` or `deferred-method` definition.

## Literal Expressions

A literal expression is evaluated to the text of the expression, which is of type `String`. It does not use the `${}` or `#{}` delimiters.

If you have a literal expression that includes the reserved `${}` or `#{}` syntax, you need to escape these characters as follows.

- By creating a composite expression as shown here:

```
#{'${}exprA}
#{'#{'}exprB}
```

The resulting values would then be the strings `${exprA}` and `#{exprB}`.

- By using the escape characters `\$` and `\#` to escape what would otherwise be treated as an eval-expression:

```
\\${exprA}
\\#{exprB}
```

The resulting values would again be the strings `${exprA}` and `#{exprB}`.

When a literal expression is evaluated, it can be converted to another type. [Table 6-2](#) shows examples of various literal expressions and their expected types and resulting values.

TABLE 6-2 Literal Expressions

Expression	Expected Type	Result
Hi	String	Hi
true	Boolean	Boolean.TRUE
42	int	42

Literal expressions can be evaluated immediately or deferred, and can be either value or method expressions. At what point a literal expression is evaluated depends on where it is being used. If the tag attribute that uses the literal expression is defined to accept a deferred value expression, then when the literal expression references a value, it is evaluated at a point in the lifecycle that is determined by other factors. The other factors include where the expression is being used and to what it is referring.

In the case of a method expression, the method that is referenced is invoked and returns the specified `String` literal. For example, the `commandButton` tag of the `guessNumber` application uses a literal method expression as a logical outcome to tell the JavaServer Faces navigation system which page to display next.

## Operators

In addition to the `.` and `[]` operators discussed in [“Value and Method Expressions” on page 116](#), the unified expression language provides the following operators, which can be used in rvalue expressions only:

- **Arithmetic:** `+`, `-` (binary), `*`, `/` and `div`, `%` and `mod`, `-` (unary)
- **Logical:** `and`, `&&`, `or`, `||`, `not`, `!`
- **Relational:** `==`, `eq`, `!=`, `ne`, `<`, `lt`, `>`, `gt`, `<=`, `ge`, `>=`, `le`. Comparisons can be made against other values, or against boolean, string, integer, or floating point literals.
- **Empty:** The empty operator is a prefix operation that can be used to determine whether a value is `null` or empty.
- **Conditional:** `A ? B : C`. Evaluate B or C, depending on the result of the evaluation of A.

The precedence of operators highest to lowest, left to right is as follows:

- `[]` `.`
- `()` (used to change the precedence of operators)
- `-` (unary) `not ! empty`
- `*` `/` `div` `%` `mod`
- `+` `-` (binary)
- `<` `>` `<=` `>=` `lt` `gt` `le` `ge`
- `==` `!=` `eq` `ne`

- && and
- || or
- ? :

## Reserved Words

The following words are reserved for the unified expression language and should not be used as identifiers.

and	or	not	eq
ne	lt	gt	le
ge	true	false	null
instanceof	empty	div	mod

## Examples of EL Expressions

Table 6–3 contains example EL expressions and the result of evaluating them.

TABLE 6–3 Example Expressions

EL Expression	Result
<code>#{1 &gt; (4/2)}</code>	false
<code>#{4.0 &gt;= 3}</code>	true
<code>#{100.0 == 100}</code>	true
<code>#{(10*10) ne 100}</code>	false
<code>#{'a' &lt; 'b'}</code>	true
<code>#{'hip' gt 'hit'}</code>	false
<code>#{4 &gt; 3}</code>	true
<code>#{1.2E4 + 1.4}</code>	12001.4
<code>#{3 div 4}</code>	0.75
<code>#{10 mod 4}</code>	2
<code>#{!empty param.Add}</code>	False if the request parameter named Add is null or an empty string

**TABLE 6-3** Example Expressions *(Continued)*

EL Expression	Result
<code>\${pageContext.request.contextPath}</code>	The context path.
<code>\${sessionScope.cart.numberOfItems}</code>	The value of the <code>numberOfItems</code> property of the session-scoped attribute named <code>cart</code> .
<code>\${param['mycom.productId']}</code>	The value of the request parameter named <code>mycom.productId</code> .
<code>\${header["host"]}</code>	The host.
<code>\${departments[deptName]}</code>	The value of the entry named <code>deptName</code> in the <code>departments</code> map.
<code>\${requestScope['javax.servlet.forward.servlet_path']}</code>	The value of the request-scoped attribute named <code>javax.servlet.forward.servlet_path</code> .
<code>#{customer.lName}</code>	Gets the value of the property <code>lName</code> from the customer bean during an initial request. Sets the value of <code>lName</code> during a postback.
<code>#{customer.calcTotal}</code>	The return value of the method <code>calcTotal</code> of the customer bean.

---

# Using JavaServer™ Faces Technology in Web Pages

---

Web pages represent the presentation layer for web applications. The process of creating web pages of a JavaServer Faces application includes tasks such as adding components to the page and wiring them to backing beans, validators, converters, and other server-side objects that are associated with the page.

This chapter explains how to create web pages using different types of component and core tags. In the next chapter you will learn about adding converters, validators and listeners to component tags that will provide additional functionality to components.

The following topics are addressed here:

- “Setting Up a Page” on page 127
- “Adding Components to a Page Using HTML Tags” on page 128
- “Using Core Tags” on page 157

## Setting Up a Page

A typical JavaServer Faces web page includes the following elements:

- A set of namespace declarations that declare the JavaServer Faces tag libraries
- Optionally, the new HTML head (`h:head`) and body (`h:body`) tags
- A form tag (`h:form`) that represents the user input components

To add the JavaServer Faces components to your web page, you need to provide the page access to the two standard tag libraries: The JavaServer Faces HTML tag library, and the JavaServer Faces core tag library. The JavaServer Faces standard HTML tag library defines tags that represent common HTML user interface components. It is linked to [HTML render kit](#). The JavaServer Faces core tag library defines tags that perform core actions.

Each JavaServer Faces tag must be described by the PDL (Page Declaration Language). For a complete list of JavaServer Faces Facelets tags and their attributes, refer to the *PDL*

*documentation* at <http://java.sun.com/javaee/javaserverfaces/2.0/docs/pdl/docs/facelets/index.html>.

To use any of the JavaServer Faces tags, you need to include appropriate directives at the top of each page specifying the tag libraries.

For Facelets applications, the XML namespace directives uniquely identify the tag library `uri` and the tag `prefix`.

For example, when creating a Facelets XHTML page, include namespace directives as follows:

```
<html xmlns="http://www.w3.org/1999/xhtml"
xmlns:h="http://java.sun.com/jsf/html"
xmlns:f="http://java.sun.com/jsf/core">
```

The XML namespace `uri` identifies the tag library location and the `prefix` value is used to distinguish the tags belonging to that specific tag library. You can also use other prefixes instead of the standard `h` or `f`. However, when including the tag in the page, you must use the prefix that you have chosen for the tag library. For example, in the following web page, the `form` tag must be referenced using the `h` prefix because the preceding tag library directive uses the `h` prefix to distinguish the tags defined in HTML tag library:

```
<h:form ...>
```

The following sections, “[Adding Components to a Page Using HTML Tags](#)” on page 128 and “[Using Core Tags](#)” on page 157, describe how to use the component tags from the JavaServer Faces standard HTML tag library and the core tags from the JavaServer Faces core tag library.

## Adding Components to a Page Using HTML Tags

The tags defined by the JavaServer Faces standard HTML tag library represent HTML form components and other basic HTML elements. These components display data or accept data from the user. This data is collected as part of a form and is submitted to the server, usually when the user clicks a button. This section explains how to use each of the component tags shown in [Table 7–1](#).

**TABLE 7–1** The Component Tags

Tag	Functions	Rendered As	Appearance
<code>column</code>	Represents a column of data in a Data component	A column of data in an HTML table	A column in a table

TABLE 7-1 The Component Tags (Continued)

Tag	Functions	Rendered As	Appearance
commandButton	Submits a form to the application	An HTML <code>&lt;input type=type&gt;</code> element, where the <code>type</code> value can be <code>submit</code> , <code>reset</code> , or <code>image</code>	A button
commandLink	Links to another page or location on a page	An HTML <code>&lt;a href&gt;</code> element	A hyperlink
dataTable	Represents a data wrapper	An HTML <code>&lt;table&gt;</code> element	A table that can be updated dynamically
form	Represents an input form (inner tags of the form receive the data that will be submitted with the form)	An HTML <code>&lt;form&gt;</code> element	No appearance
graphicImage	Displays an image	An HTML <code>&lt;img&gt;</code> element	An image
inputHidden	Allows a page author to include a hidden variable in a page	An HTML <code>&lt;input type=hidden&gt;</code> element	No appearance
inputSecret	Allows a user to input a string without the actual string appearing in the field	An HTML <code>&lt;input type=password&gt;</code> element	A text field, which displays a row of characters instead of the actual string entered
inputText	Allows a user to input a string	An HTML <code>&lt;input type=text&gt;</code> element	A text field
inputTextarea	Allows a user to enter a multiline string	An HTML <code>&lt;textarea&gt;</code> element	A multi-row text field
message	Displays a localized message	An HTML <code>&lt;span&gt;</code> tag if styles are used	A text string
messages	Displays localized messages	A set of HTML <code>&lt;span&gt;</code> tags if styles are used	A text string
outputFormat	Displays a localized message	Plain text	Plain text
outputLabel	Displays a nested component as a label for a specified input field	An HTML <code>&lt;label&gt;</code> element	Plain text
outputLink	Links to another page or location on a page without generating an action event	An HTML <code>&lt;a&gt;</code> element	A hyperlink

TABLE 7-1 The Component Tags (Continued)

Tag	Functions	Rendered As	Appearance
outputText	Displays a line of text	Plain text	Plain text
panelGrid	Displays a table	An HTML <table> element with <tr> and <td> elements	A table
panelGroup	Groups a set of components under one parent	A HTML <div> or <span> element	A row in a table
selectBooleanCheckbox	Allows a user to change the value of a Boolean choice	An HTML <input type=checkbox> element.	A check box
selectItem	Represents one item in a list of items in a SelectOne component	An HTML <option> element	No appearance
selectItems	Represents a list of items in a SelectOne component	A list of HTML <option> elements	No appearance
selectManyCheckbox	Displays a set of check boxes from which the user can select multiple values	A set of HTML <input> elements of type checkbox	A set of check boxes
selectManyListbox	Allows a user to select multiple items from a set of items, all displayed at once	An HTML <select> element	A list box
selectManyMenu	Allows a user to select multiple items from a set of items	An HTML <select> element	A scrollable combo box
selectOneListbox	Allows a user to select one item from a set of items, all displayed at once	An HTML <select> element	A list box
selectOneMenu	Allows a user to select one item from a set of items	An HTML <select> element	A scrollable combo box
selectOneRadio	Allows a user to select one item from a set of items	An HTML <input type=radio> element	A set of radio buttons

The next section explains the important tag attributes that are common to most component tags. For each of the components discussed in the following sections, [“Writing Bean Properties” on page 174](#) explains how to write a bean property bound to a particular component or its value.

## Common Component Tag Attributes

In general, most of the component tags support the following attributes:

- `id`: Uniquely identifies the component.
- `immediate`: If set to `true`, indicates that any events, validation, and conversion associated with the component should happen in the apply request values phase rather than a later phase.
- `rendered`: Specifies a condition under which the component should be rendered. If the condition is not satisfied, the component is not rendered.
- `style`: Specifies a Cascading Style Sheet (CSS) style for the tag.
- `styleClass`: Specifies a CSS stylesheet class that contains definitions of the styles.
- `value`: Identifies an external data source and binds the component's value to it.
- `binding`: Identifies a bean property and binds the component instance to it.

All of the tag attributes (except `id`) can accept expressions, as defined by the EL, described in [Chapter 6, “Unified Expression Language.”](#)

### The `id` Attribute

The `id` attribute is not usually required for a component tag. It is used when another component or a server-side class must refer to the component. If you don't include an `id` attribute, the JavaServer Faces implementation automatically generates a component ID. Unlike most other JavaServer Faces tag attributes, the `id` attribute only takes expressions using the evaluation syntax described in [The `immediate` Attribute](#), which uses the `{}` delimiters. For more information on expression syntax, see [“Value Expressions” on page 116](#).

### The `immediate` Attribute

Input components and command components (those that implement `ActionSource`, such as buttons and hyperlinks) can set the `immediate` attribute to `true` to force events, validations, and conversions to be processed during the apply request values phase of the life cycle (a sub phase in the request phase of the JavaServer Faces lifecycle).

You need to carefully consider how the combination of an input component's `immediate` value and a command component's `immediate` value determines what happens when the command component is activated.

Assume that you have a page with a button and a field for entering the quantity of a book in a shopping cart. If both the button's and the field's `immediate` attributes are set to `true`, the new value entered in the field will be available for any processing associated with the event that is generated, when the button is clicked. The event associated with the button and the event, validation, and conversion associated with the field are all handled during the apply request values phase.

If the button's `immediate` attribute is set to `true` but the field's `immediate` attribute is set to `false`, the event associated with the button is processed without updating the field's local value to the model layer. This is because any events, conversion, or validation associated with the field occurs during its usual phases of the life cycle, which come after the apply request values phase.

For a complete description of JavaServer Faces lifecycle phases, see the *JavaServer Faces 2.0 Specification*.

## The rendered Attribute

A component tag uses a Boolean EL expression, along with the `rendered` attribute, to determine whether or not the component will be rendered. For example, the `commandLink` component in the following section of a page is not rendered if the cart contains no items:

```
<h:commandLink id="check"
  ...
  rendered="#{cart.numberOfItems > 0}">
  <h:outputText
    value="#{bundle.CartCheck}"/>
</h:commandLink>
```

Unlike nearly every other JavaServer Faces tag attribute, the `rendered` attribute is restricted to using `rvalue` expressions. As explained in [“Value and Method Expressions” on page 116](#), these `rvalue` expressions can only read data; they cannot write the data back to the data source. Therefore, expressions used with `rendered` attributes can use the arithmetic operators and literals that `rvalue` expressions can use but `lvalue` expressions cannot use. For example, the expression in the preceding example uses the `>` operator.

## The style and styleClass Attributes

The `style` and `styleClass` attributes allow you to specify Cascading Style Sheets (CSS) styles for the rendered output of your tags. [“Displaying Error Messages With the `h:message` and `h:messages` Tags” on page 152](#) describes an example of using the `style` attribute to specify styles directly in the attribute. A component tag can instead refer to a CSS stylesheet class.

The following example shows the use of a `dataTable` tag that references the style class `list-background`:

```
<h:dataTable id="books"
  ...
  styleClass="list-background"
  value="#{bookDBAO.books}"
  var="book">
```

The stylesheet that defines this class is `stylesheet.css`, which will be included in the application. For more information on defining styles, see *Cascading Style Sheets Specification* at <http://www.w3.org/Style/CSS/>.

## The value and binding Attributes

A tag representing a Output component or a subclass of Output component class uses value and binding attributes to bind its component's value or instance respectively to an external data source.

## Adding HTML Head and Body Tags

The new HTML head (h:head) and body (h:body) tags add HTML type page structure to JavaServer Faces web pages.

- The h:head tag represents the head element of a HTML page
- The h:body tag represents the body element of a HTML page

The following is an example of a XHTML page using the usual head and body markups:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
<title>Add a title</title>
</head>
<body>
Add Content
</body>
```

The following is an example of a XHTML page using h:head and h:body tags:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
xmlns:h="http://java.sun.com/jsf/html">
<h:head>
Add a title
</h:head>
<h:body>
Add Content
</h:body>
```

Both of the above example code segments render the same HTML elements. The head and body tags are mainly useful for resource relocation. For more information on resource relocation see, [“Resource Relocation using h:output Tags” on page 154.](#)

## Adding a Form Component

A h:form tag represents an input form, which includes child components that can contain data, that is either presented to the user or submitted with the form.

Figure 7–1 shows a typical login form in which a user enters a user name and password, then submits the form by clicking the Login button.



The image shows a simple login form. It consists of two text input fields stacked vertically. The first field is labeled "User Name:" and contains the text "Duke". The second field is labeled "Password:" and contains eight asterisks "\*\*\*\*\*". Below these fields is a rectangular button with the text "Login" in a bold, sans-serif font.

FIGURE 7-1 A Typical Form

The `h:form` tag represents the Form component on the page and encloses all the components that display or collect data from the user, as shown here:

```
<h:form>
... other JavaServer Faces tags and other content...
</h:form>
```

The `h:form` tag can also include HTML markup to lay out the components on the page. Note that the `h:form` tag itself does not perform any layout; its purpose is to collect data and to declare attributes that can be used by other components in the form.

A page can include multiple `h:form` tags, but only the values from the form submitted by user will be included in the postback request.

## Using Text Components

Text components allow users to view and edit text in web applications. The basic types of text components are as follows:

- Label, which displays read-only text.
- Text field, which allows users to enter text, often to be submitted as part of a form.
- Text area, which is a type of text field that allows users to enter multiple lines of text.
- Password field, which is a type of text field that displays a set of characters, such as asterisks, instead of the password text that the user enters.

Figure 7–2 shows examples of these text components.

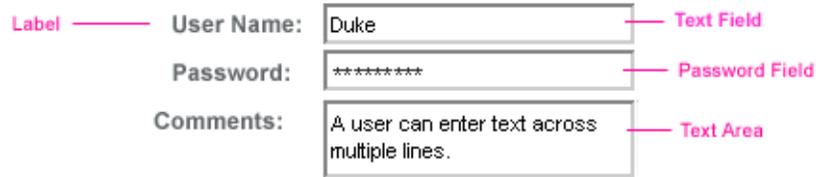


FIGURE 7-2 Example Text Components

Text components can be categorized into two types; Input and Output. A JavaServer Faces Output component is rendered as a read-only text. An example is a label. A JavaServer Faces Input component is rendered as an editable text. An example is a text field.

The Input and Output components can each be rendered in one of four ways to display more specialized text. Table 7-2 and Table 7-3 list the Input and Output components and the tags that represent the component.

---

**Note** – The name of a tag is composed of the name of the component and the name of the renderer. For example, the `h:inputText` tag refers to a Input component that is rendered with the Text renderer.

---

TABLE 7-2 Input Tags

Component	Tag	Function
Input	<code>h:inputHidden</code>	Allows a page author to include a hidden variable in a page
	<code>h:inputSecret</code>	The standard password field: Accepts one line of text with no spaces and displays it as a set of asterisks as it is typed
	<code>h:inputText</code>	The standard text field: Accepts a text string of one line
	<code>h:inputTextarea</code>	The standard text area: Accepts multiple lines of text

The Input tags support the following tag attributes in addition to those described in “Common Component Tag Attributes” on page 131. Note that this list does not include all the attributes supported by the Input tags, but just those that are used most often. For the complete list of attributes, refer to the PDL Documents at <http://java.sun.com/javasee/javaxserverfaces/2.0/docs/pdl/docs/facelets/index.html>.

- `converter`: Identifies a converter that will be used to convert the component’s local data. See “Using the Standard Converters” on page 159 for more information on how to use this attribute.
- `converterMessage`: Specifies an error message to display when the converter registered on the component fails.

- `dir`: Specifies the direction of the text displayed by this component. Acceptable values are LTR, meaning left-to-right, and RTL, meaning right-to-left.
- `label`: Specifies a name that can be used to identify this component in error messages.
- `lang`: Specifies the code for the language used in the rendered markup, such as `en_US`.
- `required`: Takes a boolean value that indicates whether or not the user must enter a value in this component.
- `requiredMessage`: Specifies an error message to display when the user does not enter a value into the component.
- `validator`: Identifies a method expression pointing to a backing bean method that performs validation on the component's data. See [“Referencing a Method That Performs Validation” on page 169](#) for an example of using the `f:validator` tag.
- `f:validatorMessage`: Specifies an error message to display when the validator registered on the component fails to validate the component's local value.
- `valueChangeListener`: Identifies a method expression that points to a backing bean method that handles the event of entering a value in this component. See [“Referencing a Method That Handles a Value-Change Event” on page 170](#) for an example of using `valueChangeListener`.

TABLE 7-3 Output Tags

Component	Tag	Function
Output	<code>h:outputLabel</code>	The standard read-only label: Displays a component as a label for a specified input field
	<code>h:outputLink</code>	Displays an <code>&lt;a href&gt;</code> tag that links to another page without generating an action event
	<code>h:outputFormat</code>	Displays a localized message
	<code>h:outputText</code>	Displays a text string of one line

The Output tags support the `converter` tag attribute in addition to those listed in [“Common Component Tag Attributes” on page 131](#).

The rest of this section explains how to use selected tags listed in [Table 7-2](#) and [Table 7-3](#). The other tags are written in a similar way.

## Rendering a Text Field With the `inputText` Tag

The `h:inputText` tag is used to display a text field. A similar tag, the `h:outputText` tag, displays a read-only, single-line string. This section shows you how to use the `h:inputText` tag. The `h:outputText` tag is written in a similar way.

Here is an example of an `h:inputText` tag:

```

<h:inputText id="name" label="Customer Name" size="50"
  value="#{cashier.name}"
  required="true"
  requiredMessage="#{customMessages.CustomerName}">
  <f:valueChangeListener
    type="com.sun.bookstore6.listeners.NameChanged" />
</h:inputText>

```

The `label` attribute specifies a user-friendly name that will be used in the substitution parameters of error messages displayed for this component.

The `value` attribute refers to the name property of a backing bean named `CashierBean`. This property holds the data for the name component. After the user submits the form, the value of the name property in `CashierBean` will be set to the text entered in the field corresponding to this tag.

The `required` attribute causes the page to reload with errors (displayed on the screen) if the user does not enter a value in the name text field. The JavaServer Faces implementation checks whether the value of the component is null or is an empty string.

If your component must have a not null value or a String value at least one character in length, you should add a `required` attribute to your tag and set its value to `true`. If your tag has a `required` attribute that is set to `true` and the value is null or a zero-length string, no other validators that are registered on the tag are called. If your tag does not have a `required` attribute set to `true`, other validators that are registered on the tag are called, but those validators must handle the possibility of a null or zero-length string.

## Rendering a Password Field With the `inputSecret` Tag

The `h:inputSecret` tag renders an `<input type="password">` HTML tag. When the user types a string into this field, a row of asterisks is displayed instead of the text typed by the user. Here is an example:

```

<h:inputSecret redisplay="false"
  value="#{LoginBean.password}" />

```

In this example, the `redisplay` attribute is set to `false`. This will prevent the password from being displayed in a query string or in the source file of the resulting HTML page.

## Rendering a Label With the `outputLabel` Tag

The `h:outputLabel` tag is used to attach a label to a specified input field for the purpose of making it accessible. The following page uses an `h:outputLabel` tag to render the label of a check box:

```

<h:selectBooleanCheckbox
  id="fanClub"
  binding="#{cashier.specialOffer}" />

```

```
<h:outputLabel for="fanClub"
    binding="#{cashier.specialOfferText}" >
    <h:outputText id="fanClubLabel"
        value="#{bundle.DukeFanClub}" />
</h:outputLabel>
...

```

The `for` attribute of the `h:outputLabel` tag maps to the `id` of the input field to which the label is attached. The `h:outputText` tag nested inside the `h:outputLabel` tag represents the actual label component. The `value` attribute on the `h:outputText` tag indicates the text that is displayed next to the input field.

Instead of using an `h:outputText` tag for the text displayed as a label, you can simply use the `h:outputLabel` tag's `value` attribute. The following code snippet shows what the previous code snippet would look like if it used the `value` attribute of the `h:outputLabel` tag to specify the text of the label. Here is an example:

```
<h:selectBooleanCheckbox
    id="fanClub"
    binding="#{cashier.specialOffer}" />
<h:outputLabel for="fanClub"
    binding="#{cashier.specialOfferText}"
    value="#{bundle.DukeFanClub}" />
</h:outputLabel>
...

```

## Rendering a Hyperlink With the `h:outputLink` Tag

The `h:outputLink` tag is used to render a hyperlink that, when clicked, loads another page but does not generate an action event. You should use this tag instead of the `h:commandLink` tag if you always want the URL (specified by the `h:outputLink` tag's `value` attribute) to open and do not want any processing to be performed when the user clicks the link. Here is an example:

```
<h:outputLink value="javadocs">
    Documentation for this demo
</h:outputLink>

```

The text in the body of the `outputLink` tag identifies the text that the user clicks to get to the next page.

## Displaying a Formatted Message With the `h:outputFormat` Tag

The `h:outputFormat` tag allows display of concatenated messages as a `MessageFormat` pattern, as described in the API documentation for `java.text.MessageFormat` (see <http://java.sun.com/javase/6/docs/api/java/text/MessageFormat.html>). Here is an example of an `outputFormat` tag:

```
<h:outputFormat value="Hello, {0} !">
<f:param value="Bill"
</h:outputFormat>
```

The `value` attribute specifies the `MessageFormat` pattern. The `param` tag specifies the substitution parameters for the message. The value of the parameter replaces the `{0}` in the sentence. The message displayed in the page is as follows:

```
Hello, Bill!
```

This is an example of hard-coding the data to be substituted in the message by using a literal value with the `value` attribute on the `param` tag.

A `h:outputFormat` tag can include more than one `param` tag for those messages that have more than one parameter that must be concatenated into the message. If you have more than one parameter for one message, make sure that you put the `param` tags in the proper order so that the data is inserted in the correct place in the message. Here is the preceding example modified with an additional parameter:

```
<h:outputFormat value="Hello, {0}! You are visitor number {1} to the page.">
<f:param value="Bill"
<f:param value="#{bean.numVisitor}">
</h:outputFormat>
```

The value of `{1}` is replaced by the second parameter. The parameter is an EL expression `bean.numVisitor`, where the property `numVisitor` of backing bean `bean`, keeps track of visitors to the page. This is an example of a value-expression-enabled tag attribute accepting an EL expression. The message displayed in the page is now as follows:

```
Hello, Bill! You are visitor number 10 to the page.
```

## Using Command Components for Performing Actions and Navigation

In JavaServer Faces applications, the `button` and `hyperlink` component tags are used to perform actions, such as submitting a form, and for navigating to another page. They are called command components as they perform an action when activated.

The `h:commandButton` tag is rendered as a button. The `h:commandLink` tag is rendered as a hyperlink.

In addition to the tag attributes listed in [“Common Component Tag Attributes” on page 131](#), the `h:commandButton` and `h:commandLink` tags can use the following attributes:

- `action`, which is either a logical outcome `String` or a method expression pointing to a bean method that returns a logical outcome `String`. In either case, the logical outcome `String` is used to determine what page to access when the Command component is activated.
- `actionListener`, which is a method expression pointing to a bean method that processes an action event fired by the Command component.

See [“Referencing a Method That Performs Navigation” on page 169](#) for more information on using the `action` attribute. See [“Referencing a Method That Handles an Action Event” on page 169](#) for details on using the `actionListener` attribute.

## Rendering a Button With the `h:commandButton` Tag

If you are using a `commandButton` component, when a user clicks the button, the data from the current page is processed, and the next page is opened. Here is a `h:commandButton` tag example:

```
<h:commandButton value="Submit"
  action="#{cashier.submit}"/>
```

Clicking the button will cause the `submit` method of `CashierBean` to be invoked because the `action` attribute references this method. The `submit` method performs some processing and returns a logical outcome.

The `value` attribute of the example `commandButton` tag references the button’s label. For information on how to use the `action` attribute, see [“Referencing a Method That Performs Navigation” on page 169](#).

## Rendering a Hyperlink With the `h:commandLink` Tag

The `h:commandLink` tag represents an HTML hyperlink and is rendered as an HTML `<a>` element. It acts like a form submit button and is used to submit an action event to the application.

A `h:commandLink` tag must include a nested `h:outputText` tag, which represents the text that the user clicks to generate the event. Here is an example:

```
<h:commandLink id="NAmerica" action="bookstore"
  actionListener="#{localeBean.chooseLocaleFromLink}">
  <h:outputText value="#{bundle.English}" />
</h:commandLink>
```

This tag will render the following HTML:

```
<a id="_id3:NAmerica" href="#"
  onclick="document.forms['_id3']['_id3:NAmerica'].
```

```
value='_id3:NAmerica';
document.forms['_id3'].submit();
return false;">English</a>
```

---

**Note** – The `h:commandLink` tag will render JavaScript. If you use this tag, make sure your browser is enabled for JavaScript.

---

## Adding Graphics and Images With the `h:graphicImage` Tag

In a JavaServer Faces application, the `Graphic` component represents an image. The `h:graphicImage` tag is used to render a `Graphic` component on a page.

```
<h:graphicImage id="mapImage" url="/template/world.jpg"/>
```

The `url` attribute specifies the path to the image. The URL of the example tag begins with a `/`, which adds the relative context path of the web application to the beginning of the path to the image.

Alternately, you can also use the [“Resources” on page 112](#) facility to point to the image location. Here is an example:

```
<h:graphicImage value="#{resource['images:wave.med.gif']}/>
```

## Laying Out Components With the `Panel` Component

In a JavaServer Faces application, you use the `Panel` component as a layout container for a set of other components. The `Panel` component is rendered as an HTML table. [Table 7–4](#) lists the tags corresponding to the `Panel` component.

**TABLE 7–4** `Panel` Component Tags

Tag	Attributes	Function
<code>h:panelGrid</code>	<code>columns</code> , <code>columnClasses</code> , <code>footerClass</code> , <code>headerClass</code> , <code>panelClass</code> , <code>rowClasses</code>	Displays a table
<code>h:panelGroup</code>	<code>layout</code>	Groups a set of components under one parent

The `h:panelGrid` tag is used to represent an entire table. The `h:panelGroup` tag is used to represent rows in a table. Other tags are used to represent individual cells in the rows.

The `columns` attribute is required if you want your table to have more than one column as the attribute defines how to group the data in the table. The `h:panelGrid` tag also has a set of attributes that specify CSS stylesheet classes: `columnClasses`, `footerClass`, `headerClass`, `panelClass`, and `rowClasses`. These stylesheet attributes are optional.

If the `headerClass` attribute value is specified, the `panelGrid` must have a header as its first child. Similarly, if a `footerClass` attribute value is specified, the `panelGrid` must have a footer as its last child.

Here is an example:

```
<h:panelGrid columns="3" headerClass="list-header"
  rowClasses="list-row-even, list-row-odd"
  styleClass="list-background"
  title="#{bundle.Checkout}">
  <f:facet name="header">
    <h:outputText value="#{bundle.Checkout}"/>
  </f:facet>
  <h:outputText value="#{bundle.Name}" />
  <h:inputText id="name" size="50"
    value="#{cashier.name}"
    required="true">
    <f:valueChangeListener
      type="listeners.NameChanged" />
  </h:inputText>
  <h:message styleClass="validationMessage" for="name"/>
  <h:outputText value="#{bundle.CCNumber}"/>
  <h:inputText id="ccno" size="19"
    converter="CreditCardConverter" required="true">
    <bookstore:formatValidator
      formatPatterns="9999999999999999|
        9999 9999 9999 9999|9999-9999-9999-9999"/>
  </h:inputText>
  <h:message styleClass="validationMessage" for="ccno"/>
  ...
</h:panelGrid>
```

The above `h:panelGrid` tag is rendered as a table that contains components in which the bookstore customer inputs personal information. This `h:panelGrid` tag uses stylesheet classes to format the table. The following code shows the `list-header` definition:

```
.list-header {
  background-color: #ffffff;
  color: #000000;
  text-align: center;
}
```

Because the `h:panelGrid` tag specifies a `headerClass`, the `panelGrid` must contain a header. The example `panelGrid` tag uses a `facet` tag for the header. Facets can have only one child, so a `h:panelGroup` tag is needed if you want to group more than one component within a facet. The example `h:panelGrid` tag has only one cell of data, therefore a `h:panelGroup` tag is not needed.

The `h:panelGroup` tag has an attribute, `layout`, in addition to those listed in “[Common Component Tag Attributes](#)” on page 131. If the `layout` attribute has the value `block`, then an HTML `div` element is rendered to enclose the row; otherwise, an HTML `span` element is rendered to enclose the row. If you are specifying styles for the `h:panelGroup` tag, you should set the `layout` attribute to `block` in order for the styles to be applied to the components within the `h:panelGroup` tag. You should do this because styles such as those that set width and height are not applied to inline elements, which is how content enclosed by the `span` element is defined.

A `h:panelGroup` tag can also be used to encapsulate a nested tree of components so that the tree of components appears as a single component to the parent component.

Data, represented by the nested tags, is grouped into rows according to the value of the `columns` attribute of the `h:panelGrid` tag. The `columns` attribute in the example is set to 3, and therefore the table will have three columns. The column in which each component is displayed is determined by the order in which the component is listed on the page modulo 3. So, if a component is the fifth one in the list of components, that component will be in the 5 modulo 3 column, or column 2.

## Displaying Components for Selecting One Value

Another commonly used component is one that allows a user to select one value, whether it be the only value available or one of a set of choices. The most common examples of this `selectOneComponent` are as follows:

- A `h:selectBooleanCheckbox` tag, displayed as check box, which represents `boolean` state
- A `h:selectOneRadio` tag, displayed as a set of radio buttons
- A `h:selectOneMenu` tag, displayed as a drop-down menu, with a scrollable list
- A `h:selectOneListbox` tag, displayed as a list box, with an unscrollable list

Figure 7–3 shows examples of these components.

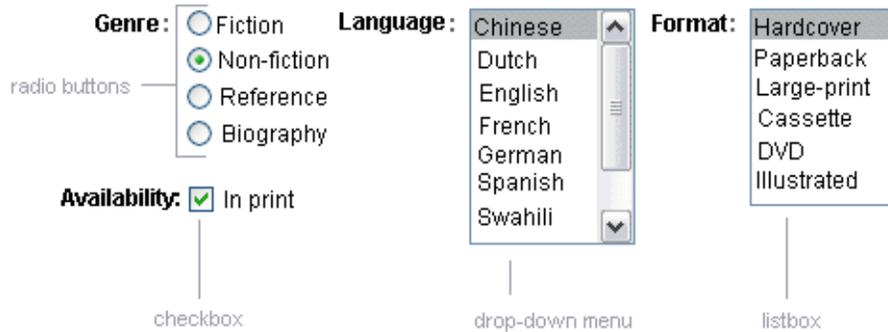


FIGURE 7-3 Example Select One Components

## Displaying a Check Box Using the `h:selectBooleanCheckbox` Tag

The `SelectBoolean` component defines tags that have a boolean value. The `h:selectBooleanCheckbox` tag is the only tag that JavaServer Faces technology provides for representing boolean state.

Here is an example that shows how to use the `selectBooleanCheckbox` tag:

```
<h:selectBooleanCheckbox
    id="fanClub"
    rendered="false"
    binding="#{cashier.specialOffer}" />
<h:outputLabel
    for="fanClub"
    rendered="false"
    binding="#{cashier.specialOfferText}">
    <h:outputText
        id="fanClubLabel"
        value="#{bundle.DukeFanClub}" />
</h:outputLabel>
```

This example tag displays a check box to allow users to indicate whether they want to join the Duke Fan Club. The label for the check box is rendered by the `outputLabel` tag. The actual text is represented by the nested `outputText` tag.

## Displaying a Menu Using the `h:selectOneMenu` Tag

A `SelectOne` component allows the user to select one value from a set of values. This component can be rendered as a list box, a set of radio buttons, or a menu. This section explains the `h:selectOneMenu` tag. The `h:selectOneRadio` and `h:selectOneListbox` tags are used in a similar way. The `h:selectOneListbox` tag is similar to the `h:selectOneMenu` tag except that `h:selectOneListbox` defines a `size` attribute that determines how many of the items are displayed at once.

The `h:selectOneMenu` tag represents a component that contains a list of items from which a user can choose one item. This menu component is also commonly known as a drop-down list or a combo box. The following code snippet shows how the `h:selectOneMenu` tag is used, to allow the user to select a shipping method:

```
<h:selectOneMenu id="shippingOption"
  required="true"
  value="#{cashier.shippingOption}">
  <f:selectItem
    itemValue="2"
    itemLabel="#{bundle.QuickShip}"/>
  <f:selectItem
    itemValue="5"
    itemLabel="#{bundle.NormalShip}"/>
  <f:selectItem
    itemValue="7"
    itemLabel="#{bundle.SaverShip}"/>
</h:selectOneMenu>
```

The `value` attribute of the `h:selectOneMenu` tag maps to the property that holds the currently selected item's value. You are not required to provide a value for the currently selected item. If you don't provide a value, the first item in the list is selected by default.

Like the `h:selectOneRadio` tag, the `selectOneMenu` tag must contain either a `f:selectItems` tag or a set of `f:selectItem` tags for representing the items in the list. [“Using The SelectItem and SelectItems Components” on page 147](#) explains these tags.

The other `selectOne` components are used in the same way.

## Rendering Components for Selecting Multiple Values

In some cases, you need to allow your users to select multiple values rather than just one value from a list of choices. The `SelectMany` components are used for this purpose. You can do this using one of the following component tags:

- A `h:selectManyCheckbox` tag, displayed as a set of check boxes
- A `h:selectManyMenu` tag, displayed as a drop-down menu
- A `h:selectManyListbox` tag, displayed as a list box

[Figure 7–4](#) shows examples of these components.



FIGURE 7-4 Example SelectMany Components

The `SelectMany` component allows the user to select zero or more values from a set of values. This section explains the `h:selectManyCheckbox` tag. The `h:selectManyListbox` tag and `h:selectManyMenu` tag are used in a similar way.

A list box differs from a menu in that it displays a subset of items in a box, whereas a menu displays only one item at a time when the user is not selecting the menu. The `size` attribute of the `h:selectManyListbox` tag determines the number of items displayed at one time. The list box includes a scroll bar for scrolling through any remaining items in the list.

The `h:selectManyCheckbox` tag renders a set of check boxes, with each check box representing one value that can be selected.

```
<h:selectManyCheckbox
  id="newsletters"
  layout="pageDirection"
  value="#{cashier.newsletters}">
  <f:selectItems
    value="#{newsletters}"/>
</h:selectManyCheckbox>
```

The `value` attribute of the `h:selectManyCheckbox` tag identifies the `newsletters` property of the `Cashier` backing bean. This property holds the values of the currently selected items from the set of check boxes. You are not required to provide a value for the currently selected items. If you don't provide a value, the first item in the list is selected by default.

The `layout` attribute indicates how the set of check boxes are arranged on the page. Because `layout` is set to `pageDirection`, the check boxes are arranged vertically. The default is `lineDirection`, which aligns the check boxes horizontally.

The `h:selectManyCheckbox` tag must also contain a tag or set of tags representing the set of check boxes. To represent a set of items, you use the `f:selectItems` tag. To represent each item individually, you use a `f:selectItem` tag. The following subsection explains these tags in more detail.

## Using The `selectItem` and `selectItems` Components

The `f:selectItem` and `f:selectItems` tags represent components that can be nested inside a `SelectOne` or a `SelectMany` component. A `f:selectItem` tag is associated with a `SelectItem` instance, which contains the value, label, and description of a single item in a `SelectOne` or `SelectMany` component. The `SelectItems` instance represents a set of `SelectItem` instances, containing the values, labels, and descriptions of the entire list of items

You can use either a set of `f:selectItem` tags or a single `f:selectItems` tag within your `SelectOne` or `SelectMany` component tag.

The advantages of using the `f:selectItems` tag are as follows:

- An item can be represented by using different data structures, including `Array`, `Map` and `Collection`. The data structure is composed of `SelectItem` instances.

The value of `f:selectItems` tag can even represent a generic collection of POJOs.

- Different lists can be concatenated together into a single `SelectMany` or `SelectOne` component and group the lists within the component.
- Values can be generated dynamically at runtime.

The advantages of using `f:selectItem` are as follows:

- Items in the list can be defined from the page.
- Less code is needed in the bean for the `selectItem` properties.

For more information on writing component properties for the `SelectItems` components, see [“Writing Bean Properties” on page 174](#). The rest of this section shows you how to use the `f:selectItems` and `f:selectItem` tags.

### Using the `f:selectItems` Tag

The following example from [“Rendering Components for Selecting Multiple Values” on page 145](#) shows how to use the `h:selectManyCheckbox` tag:

```
<h:selectManyCheckbox
  id="newsletters"
  layout="pageDirection"
  value="#{cashier.newsletters}">
  <f:selectItems
    value="#{newsletters}"/>
</h:selectManyCheckbox>
```

The `value` attribute of the `f:selectItems` tag is bound to the backing bean `newsletters`, which is configured in the application configuration resource file.

You can also create the list corresponding to a `SelectMany` or `SelectOne` component programmatically in the backing bean. See [“Writing Bean Properties” on page 174](#) for information on how to write a backing bean property corresponding to a `SelectMany` or `SelectOne` component.

The arguments to the `SelectItem` constructor are as follows:

- An `Object` representing the value of the item
- A `String` representing the label that displays in the `SelectMany` component on the page
- A `String` representing the description of the item

[“SelectItems Properties” on page 179](#) describes in more detail how to write a backing bean property for a `SelectItems` component.

## Using the `f:selectItem` Tag

The `f:selectItem` tag represents a single item in a list of items. Here is the example from [“Displaying a Menu Using the `h:selectOneMenu` Tag” on page 144](#) once again:

```
<h:selectOneMenu
    id="shippingOption" required="true"
    value="#{cashier.shippingOption}">
  <f:selectItem
    itemValue="2"
    itemLabel="#{bundle.QuickShip}"/>
  <f:selectItem
    itemValue="5"
    itemLabel="#{bundle.NormalShip}"/>
  <f:selectItem
    itemValue="7"
    itemLabel="#{bundle.SaverShip}"/>
</h:selectOneMenu>
```

The `itemValue` attribute represents the default value of the `SelectItem` instance. The `itemLabel` attribute represents the `String` that appears in the drop-down menu component on the page.

The `itemValue` and `itemLabel` attributes are value-binding-enabled, meaning that they can use value-binding expressions to refer to values in external objects. They can also define literal values, as shown in the example `h:selectOneMenu` tag.

## Using Data-Bound Table Components

Data-bound table components display relational data in a tabular format. [Figure 7–5](#) shows an example of this kind of table.

Quantity	Title	Price	
<input type="text" value="1"/>	<a href="#">Web Servers for Fun and Profit</a>	\$40.75	<input type="button" value="Remove Item"/>
<input type="text" value="3"/>	<a href="#">Web Components for Web Developers</a>	\$27.75	<input type="button" value="Remove Item"/>
<input type="text" value="1"/>	<a href="#">From Oak to Java: The Revolution of a Language</a>	\$10.75	<input type="button" value="Remove Item"/>
<input type="text" value="2"/>	<a href="#">My Early Years: Growing up on *7</a>	\$30.75	<input type="button" value="Remove Item"/>
<input type="text" value="1"/>	<a href="#">Java Intermediate Bytecodes</a>	\$30.95	<input type="button" value="Remove Item"/>
<input type="text" value="3"/>	<a href="#">Duke: A Biography of the Java Evangelist</a>	\$45.00	<input type="button" value="Remove Item"/>
<b>Subtotal:\$362.20</b>			

[Update Quantities](#)

FIGURE 7-5 Table on a Web Page

In a JavaServer Faces application, the `h:dataTable` component tag supports binding to a collection of data objects. It displays the data as an HTML table. The `h:column` tag represents a column of data within the table. It iterates over each record in the data source which is displayed as a row. Here is an example:

```
<h:dataTable id="items"
  captionClass="list-caption"
  columnClasses="list-column-center, list-column-left,
    list-column-right, list-column-center"
  footerClass="list-footer"
  headerClass="list-header"
  rowClasses="list-row-even, list-row-odd"
  styleClass="list-background"
  <h:column headerClass="list-header-left">
    <f:facet name="header">
      <h:outputText value="Quantity" />
    </f:facet>
    <h:inputText id="quantity" size="4"
      value="#{item.quantity}" />
    ...
  </h:column>
  <h:column>
    <f:facet name="header">
      <h:outputText value="Title" />
    </f:facet>
    <h:outputText value="#{item.title}" />
  </h:column>
  ...
  <f:facet name="footer">
    <h:panelGroup>
```

```
        <h:outputText value="Total"/>
        <h:outputText value="#{cart.total}" />
        <f:convertNumber type="currency" />
    </h:outputText>
</h:panelGroup>
</f:facet>
<
</h:dataTable>
```

Figure 7–5 shows a data grid that this `h:dataTable` tag can display.

The example `h:dataTable` tag displays the books in the shopping cart as well as the quantity of each book in the shopping cart, the prices, and a set of buttons, which the user can click to remove books from the shopping cart.

The `column` tags represent columns of data in a `Data` component. While the `Data` component is iterating over the rows of data, it processes the `Column` component associated with each `h:column` tag for each row in the table.

The `Data` component shown in the preceding code example iterates through the list of books (`cart.items`) in the shopping cart and displays their titles, authors, and prices. Each time `Data` iterates through the list of books, it renders one cell in each column.

The `h:dataTable` and `h:column` tags use facets to represent parts of the table that are not repeated or updated. These include headers, footers, and captions.

In the preceding example, `h:column` tags include `f:facet` tags for representing column headers or footers. The `h:column` tag allows you to control the styles of these headers and footers by supporting the `headerClass` and `footerClass` attributes. These attributes accept space-separated lists of CSS style classes, which will be applied to the header and footer cells of the corresponding column in the rendered table.

Facets can have only one child, and so a `h:panelGroup` tag is needed if you want to group more than one component within a `f:facet`. Because the facet tag representing the footer includes more than one tag, the `panelGroup` is needed to group those tags. Finally, this `h:dataTable` tag includes a `f:facet` tag with its `name` attribute set to `caption`, causing a table caption to be rendered below the table.

This table is a classic use case for a `Data` component because the number of books might not be known to the application developer or the page author at the time that application is developed. The `Data` component can dynamically adjust the number of rows of the table to accommodate the underlying data.

The `value` attribute of a `h:dataTable` tag references the data to be included in the table. This data can take the form of any of the following:

- A list of beans
- An array of beans

- A single bean
- A `javax.faces.model.DataModel`
- A `java.sql.ResultSet`
- A `javax.servlet.jsp.jstl.sql.ResultSet`
- A `javax.sql.RowSet`

All data sources for Data components have a `DataModel` wrapper. Unless you explicitly construct a `DataModel` wrapper, the JavaServer Faces implementation will create one around data of any of the other acceptable types. See [“Writing Bean Properties” on page 174](#) for more information on how to write properties for use with a Data component.

The `var` attribute specifies a name that is used by the components within the `h:dataTable` tag as an alias to the data referenced in the `value` attribute of `dataTable`.

In the example `h:dataTable` tag, the `value` attribute points to a list of books. The `var` attribute points to a single book in that list. As the Data component iterates through the list, each reference to `item` points to the current book in the list.

The Data component also has the ability to display only a subset of the underlying data. This is not shown in the preceding example. To display a subset of the data, you use the optional `first` and `rows` attributes.

The `first` attribute specifies the first row to be displayed. The `rows` attribute specifies the number of rows, starting with the first row, to be displayed. For example, if you wanted to display records 2 through 10 of the underlying data, you would set `first` to 2 and `rows` to 9. When you display a subset of the data in your pages, you might want to consider including a link or button that causes subsequent rows to display when clicked. By default, both `first` and `rows` are set to zero, and this causes all the rows of the underlying data to display.

The `h:dataTable` tag also has a set of optional attributes for adding styles to the table:

- `captionClass`: Defines styles for the table caption
- `columnClasses`: Defines styles for all the columns
- `footerClass`: Defines styles for the footer
- `headerClass`: Defines styles for the header
- `rowClasses`: Defines styles for the rows
- `styleClass`: Defines styles for the entire table

Each of these attributes can specify more than one style. If `columnClasses` or `rowClasses` specifies more than one style, the styles are applied to the columns or rows in the order that the styles are listed in the attribute. For example, if `columnClasses` specifies styles `list-column-center` and `list-column-right` and if there are two columns in the table, the first column will have style `list-column-center`, and the second column will have style `list-column-right`.

If the `style` attribute specifies more styles than there are columns or rows, the remaining styles will be assigned to columns or rows starting from the first column or row. Similarly, if the `style`

attribute specifies fewer styles than there are columns or rows, the remaining columns or rows will be assigned styles starting from the first style.

## Displaying Error Messages With the `h:message` and `h:messages` Tags

The `h:message` and `h:messages` tags are used to display error messages when conversion or validation fails. The `h:message` tag displays error messages related to a specific input component, whereas the `h:messages` tag displays the error messages for the entire page.

Here is an example `h:message` tag from the `guessNumber` application:

```
<h:inputText id="userNo" value="#{UserNumberBean.userNumber}">
  <f:validateLongRange minimum="0" maximum="10" />
  <h:commandButton id="submit"
    action="success" value="Submit" /><>
<h:message
  style="color: red;
  font-family: 'New Century Schoolbook', serif;
  font-style: oblique;
  text-decoration: overline" id="errors1" for="userNo"/>
```

The `for` attribute refers to the ID of the component that generated the error message. The error message is displayed at the same location that the `h:message` tag appears in the page. In this case, the error message will appear after the Submit button.

The `style` attribute allows you to specify the style of the text of the message. In the example in this section, the text will be red, New Century Schoolbook, serif font family, and oblique style, and a line will appear over the text. The message and messages tags support many other attributes for defining styles. For more information on these attributes, refer to the *PDL documentation* at <http://java.sun.com/javase/6/docs/api/javax.faces.component.htmlmessages.MessagesImpl.html>.

Another attribute supported by `h:messages` tag is the `layout` attribute. Its default value is `list`, which indicates that the messages are displayed in a bullet list using the HTML `ul` and `li` elements. If you set the attribute value to `table`, the messages will be rendered in a table using the HTML `table` element.

The preceding example shows a standard validator that is registered on the input component. The message tag displays the error message that is associated with this validator when the validator cannot validate the input component's value. In general, when you register a converter or validator on a component, you are queuing the error messages associated with the converter or validator on the component. The `h:message` and `h:messages` tags display the appropriate error messages that are queued on the component when the validators or converters registered on that component fail to convert or validate the component's value.

Standard error messages are provided with standard converters and standard validators. An application architect can override these standard messages and supply error messages for custom converters and validators by registering custom error messages with the application through the `message-bundle`. Creating and using custom error messages is an advanced topic covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

## Creating Bookmarkable URLs with `h:button` and `h:link` Tags

Bookmarkability or the ability to create bookmarkable URLs refers to the ability to generate hyperlinks based on specified navigation outcome and component parameters. Bookmarkable URLs are supported in JavaServer Faces 2.0.

In HTTP protocol, by default most browsers send GET requests for URL retrieval and POST requests for data processing. The GET requests can have query parameters and can be cached while it is not advised for POSTs which send data to the external servers. The other JavaServer Faces tags capable of generating hyperlinks use either simple GET requests as in the case of `h:outputLink`, or POST requests as in the case of `h:commandLink` or `h:commandButton` tags. Get requests with query parameters provide finer granularity to URL strings. These URLs are created with a one or more `name=value` parameters appended to the simple URL after a `?` character and separated by either `&`; or `&amp;`; strings.

Bookmarkable URLs can be created with the help of the `OutcomeTarget` component, which is rendered as one of the following two HTML tags:

- `h:button`
- `h:link`

Both of these tags are capable of generating a hyperlink based on the `outcome` attribute of the component. For example:

```
<h:link outcome="response" value="Message">
<f:param name="Result" value="#{sampleBean.result}"/>
</h:link>
```

The `h:link` tag will generate a URL link that points to the `response.xhtml` file on the same server, appended with the single query parameter created by the `f:param` tag. When processed, the parameter `Result` is assigned the value of backing bean's `result` method `#{sampleBean.result}`. A sample HTML generated from the above set of tags is as follows assuming the value of the parameter is `success`:

```
<a href="http://localhost:8080/guessNumber/guess/response.xhtml?Result=success">Response</a>
```

This is a simple GET request. To create more complex GET requests and utilize the `h:link` tag's functionality, you may use `ViewParameters`.

## Using View Parameters

The core tags `metadata` and `f:viewParam`, are used as a source of parameters for configuring the URLs. View parameters are declared as part of `metadata` for a page as shown in the following example:

```
<h:body>
<f:metadata>
<f:viewParam id="name" name="Name" value="#{sampleBean.username}"/>
<f:viewParam id="ID" name="uid" value="#{sampleBean.useridentity}"/>

</f:metadata>
<h:link outcome="response" value="Message" includeViewParams="true">
  </h:link>
</h:body>
```

View parameters are declared with `f:viewParam` tag and are placed inside the `f:metadata` tag. If `includeViewParams` attribute is set on the component, the view parameters are added to the hyperlink.

The resulting URL will look like this:

```
http://localhost:8080/guessNumber/guess/response.xhtml?Name=Duke&uid=2001
```

As the URL can be the result of various parameter values, the order of the URL creation has been predefined. The order in which the various parameter values are read is as under:

1. Component
2. Navigation-case parameters
3. View parameters

When there is a GET request for the page, the `Restore View` and `Render Response` phases (sub phases of JavaServer Applications request lifecycle) are executed immediately. In case the page is using view parameters for creating a bookmarkable URL, the post-back request lifecycle is executed with all phases being processed.

## Resource Relocation using `h:output` Tags

Resource relocation refers to the ability of a JavaServer Faces application to specify the location where a resource can be rendered. Resource relocation can be defined with the following two new HTML tags introduced in JavaServer Faces 2.0.

- `h:outputScript`
- `h:outputStylesheet`

These tags have a couple of attributes, `name` and `target` which can be used for defining the render location. For a complete list of attributes for these tags, see the *PDL Documentation* at

<http://java.sun.com/javaee/javaserverfaces/2.0/docs/pdldocs/facelets/index.html>.

For `h:outputScript` tag, the name and target attributes define where the output of a resource may appear. Here is an example:

```
!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:h="http://java.sun.com/jsf/html">
  <h:head id="head">
    <title>Resource Relocation</title>
  </h:head>
  <h:body id="body">
    <h:form id="form">
      <h:outputScript name="hello.js"/>
      <h:outputStylesheet name="hello.css"/>
    </h:form>
  </h:body>
</html>
```

Since target attribute is not defined in the tag, the Stylesheet `hello.css` is rendered in head and the `hello.js` script will be rendered in the body of the page as defined by `h:head` tag.

Here is the HTML generated by the above page:

```
<html xmlns="http://www.w3.org/1999/xhtml">
  <head>
    <title>Resource Relocation</title>
    <link type="text/css" rel="stylesheet"
      href="/ctx/faces/javafx.faces.resource/hello.css"/>
  </head>
  <body>
    <form id="form" name="form" method="post" action="..." enctype="...">
      <script type="text/javascript"
        src="/ctx/faces/javafx.faces.resource/hello.js">
      </script>
    </form>
  </body>
</html>
```

The original page can be recreated setting the target attribute for the `h:outputScript` tag which allows the incoming GET request to provide the location parameter. Here is an example:

```
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:h="http://java.sun.com/jsf/html">
  <h:head id="head">
    <title>Resource Relocation</title>
```

```
</h:head>
<h:body id="body">
  <h:form id="form">
    <h:outputScript name="hello.js" target="#{param.location}"/>
    <h:outputStylesheet name="hello.css"/>
  </h:form>
</h:body>
</html>
```

In this case, if the incoming request does not provide a location parameter, the default locations will still apply: the stylesheet is rendered in the head and the script inline. However if the incoming request provides the location parameter as head, both the stylesheet and the script will be rendered in the head element.

The HTML generated by the above page is as follows:

```
<html xmlns="http://www.w3.org/1999/xhtml">
  <head>
    <title>Resource Relocation</title>
    <link type="text/css" rel="stylesheet"
      href="/ctx/faces/javax.faces.resource/hello.css"/>
    <script type="text/javascript"
      src="/ctx/faces/javax.faces.resource/hello.js">
    </script>
  </head>
  <body>
    <form id="form" name="form" method="post" action="..." enctype="...">
    </form>
  </body>
</html>
```

Similarly if the incoming request provides the location parameter as body, the script will be rendered in the body element.

The above section describes simple uses for the resource relocation. Resource relocation feature can add even more functionality for the components and pages. A page author does not have to know the location of a resource or its placement.

Component authors, by using `@ResourceDependency` annotation for the components, can define the resources for the component such as a stylesheet and script. This allows the page authors freedom from defining resource locations.

## Using Core Tags

The tags included in the JavaServer Faces core tag library are used to perform core actions that are not performed by HTML tags. Commonly used core tags are listed in [Table 7-5](#) along with the functions they perform.

TABLE 7-5 The Core Tags

Tag Categories	Tags	Functions
Event-handling tags	<code>f:actionListener</code>	Adds an action listener to a parent component
	<code>f:phaseListener</code>	Adds a <code>PhaseListener</code> to a page
	<code>f:setPropertyActionListener</code>	Registers a special action listener whose sole purpose is to push a value into a backing bean when a form is submitted
	<code>f:valueChangeListener</code>	Adds a value-change listener to a parent component
Attribute configuration tag	<code>f:attribute</code>	Adds configurable attributes to a parent component
Data conversion tags	<code>f:converter</code>	Adds an arbitrary converter to the parent component
	<code>f:convertDateTime</code>	Adds a <code>DateTime</code> converter instance to the parent component
	<code>f:convertNumber</code>	Adds a <code>Number</code> converter instance to the parent component
Facet tag	<code>f:facet</code>	Adds a nested component that has a special relationship to its enclosing tag
	<code>f:metadata</code>	Registers a facet on a parent component
Localization tag	<code>f:loadBundle</code>	Specifies a <code>ResourceBundle</code> that is exposed as a <code>Map</code>
Parameter substitution tag	<code>f:param</code>	Substitutes parameters into a <code>MessageFormat</code> instance and adds query string name-value pairs to a URL
Tags for representing items in a list	<code>f:selectItem</code>	Represents one item in a list of items in a <code>SelectOne</code> or <code>SelectMany</code> component
	<code>f:selectItems</code>	Represents a set of items in a <code>SelectOne</code> or <code>SelectMany</code> component

TABLE 7-5 The Core Tags (Continued)

Tag Categories	Tags	Functions
Validator tags	<code>f:validateDoubleRange</code>	Adds a <code>DoubleRangeValidator</code> to a component
	<code>f:validateLength</code>	Adds a <code>LengthValidator</code> to a component
	<code>f:validateLongRange</code>	Adds a <code>LongRangeValidator</code> to a component
	<code>f:validator</code>	Adds a custom validator to a component
	<code>f:validateRegEx</code>	Adds a <code>RegExValidator</code> instance to a component
	<code>f:validateBean</code>	Delegates the validation of a local value to a <code>BeanValidator</code> instance
	<code>f:validateRequired</code>	Enforces the presence of a value in a component
Ajax tag	<code>f:ajax</code>	Associates Ajax action to a single or group of components based on placement
Event tag	<code>f:event</code>	Allows installing <code>ComponentSystemEventListener</code> on a component

These tags, which are used in conjunction with component tags, are explained in other sections of this tutorial. [Table 7-6](#) lists the sections that explain how to use specific core tags.

TABLE 7-6 Where the Core Tags Are Explained

Tags	Where Explained
Event-handling tags	<a href="#">“Registering Listeners on Components” on page 164</a>
Data conversion tags	<a href="#">“Using the Standard Converters” on page 159</a>
facet	<a href="#">“Using Data-Bound Table Components” on page 148</a> and <a href="#">“Laying Out Components With the Panel Component” on page 141</a>
loadBundle	<a href="#">“Rendering Components for Selecting Multiple Values” on page 145</a>
param	<a href="#">“Displaying a Formatted Message With the h:outputFormat Tag” on page 138</a>
selectItem and selectItems	<a href="#">“Using The SelectItem and SelectItems Components” on page 147</a>
Validator tags	<a href="#">“Using the Standard Validators” on page 166</a>

## Using Converters, Listeners and Validators

---

The previous chapter described different types of components and explained how to add them to a web page. This chapter provides information on adding more functionality to the components through converters, listeners and validators.

- Converters are used to convert data that is received from the input components
- Listeners are used to listen to the events happening in the page and perform actions as defined
- Validators are used to validate the data that is received from the input components

The following topics are addressed in this chapter:

- [“Using the Standard Converters” on page 159](#)
- [“Registering Listeners on Components” on page 164](#)
- [“Using the Standard Validators” on page 166](#)
- [“Referencing a Backing Bean Method” on page 168](#)

### Using the Standard Converters

The JavaServer™ Faces implementation provides a set of Converter implementations that you can use to convert component data.

The standard Converter implementations, located in the `javax.faces.convert` package, are as follows:

- `BigDecimalConverter`
- `BigIntegerConverter`
- `BooleanConverter`
- `ByteConverter`
- `CharacterConverter`
- `DateTimeConverter`
- `DoubleConverter`

- EnumConverter
- FloatConverter
- IntegerConverter
- LongConverter
- NumberConverter
- ShortConverter

A standard error message associated with each of these converters. If you have registered one of these converters onto a component on your page, and the converter is not able to convert the component's value, the converter's error message will display on the page. For example, the error message that displays if `BigIntegerConverter` fails to convert a value is:

```
{0} must be a number consisting of one or more digits
```

In this case, the `{0}` substitution parameter will be replaced with the name of the input component on which the converter is registered.

Two of the standard converters (`DateTimeConverter` and `NumberConverter`) have their own tags, which allow you to configure the format of the component data using the tag attributes. For more information about using `DateTimeConverter`, see [“Using DateTimeConverter” on page 161](#). For more information about using `NumberConverter`, see [“Using NumberConverter” on page 163](#). The following section explains how to convert a component's value, including how to register other standard converters with a component.

## Converting a Component's Value

To use a particular converter to convert a component's value, you need to register the converter onto the component. You can register any of the standard converters in one of the following four ways:

- Nest one of the standard converter tags inside the component's tag. These tags are `convertDateTime` and `convertNumber`, which are described in [“Using DateTimeConverter” on page 161](#) and [“Using NumberConverter” on page 163](#), respectively.
- Bind the value of the component to a backing bean property of the same type as the converter.
- Refer to the converter from the component tag's `converter` attribute.
- Nest a converter tag inside of the component tag and use either the converter tag's `converterId` attribute or its `binding` attribute to refer to the converter.

As an example of the second way, if you want a component's data to be converted to an `Integer`, you can simply bind the component's value to a backing bean property. Here is an example:

```
Integer age = 0;  
public Integer getAge(){ return age;}  
public void setAge(Integer age) {this.age = age;}
```

If the component is not bound to a bean property, you can employ the third method by using the `converter` attribute directly on the component tag:

```
<h:inputText
    converter="javax.faces.convert.IntegerConverter" />
```

This example shows the `converter` attribute referring to the fully-qualified class name of the converter. The `converter` attribute can also take the ID of the component.

The data from `inputText` tag in this example will be converted to a `java.lang.Integer`. The `Integer` type is already a supported type of the `NumberConverter`. If you don't need to specify any formatting instructions using the `convertNumber` tag attributes, and if one of the standard converters will suffice, you can simply reference that converter by using the component tag's `converter` attribute.

Finally, you can nest a `converter` tag within the component tag and use either the `converter` tag's `converterId` attribute or its `binding` attribute to reference the converter.

The `converterId` attribute must reference the converter's ID. Here is an example:

```
<h:inputText value="#{LoginBean.Age}" />
    <f:converter converterId="Integer" />
</h:inputText>
```

Instead of using the `converterId` attribute, the `converter` tag can use the `binding` attribute. The `binding` attribute must resolve to a bean property that accepts and returns an appropriate `Converter` instance.

Custom converters and using `binding` attribute are advanced topics covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

## Using DateTimeConverter

You can convert a component's data to a `java.util.Date` by nesting the `convertDateTime` tag inside the component tag. The `convertDateTime` tag has several attributes that allow you to specify the format and type of the data. [Table 8-1](#) lists the attributes.

Here is a simple example of a `convertDateTime` tag:

```
<h:outputText id="shipDate" value="#{cashier.shipDate}">
    <f:convertDateTime dateStyle="full" />
</h:outputText>
```

When binding the `DateTime` converter to a component, ensure that the backing bean property to which the component is bound is of type `java.util.Date`. In the preceding example, `cashier.shipDate` must be of type `java.util.Date`.

The example tag can display the following output:

Saturday, September 26, 2009

You can also display the same date and time by using the following tag where date format is specified:

```
<h:outputText value="#{cashier.shipDate}">
  <f:convertDateTime
    pattern="EEEEEEEE, MMM dd, yyyy" />
</h:outputText>
```

If you want to display the example date in Spanish, you can use the `locale` attribute:

```
<h:inputText value="#{cashier.shipDate}">
  <f:convertDateTime dateStyle="full"
    locale="Locale.SPAIN"
    timeStyle="long" type="both" />
</h:inputText>
```

This tag would display the following output:

sabado 26 de septiembre de 2009

Refer to the *Customizing Formats* lesson of the Java Tutorial at <http://java.sun.com/docs/books/tutorial/i18n/format/simpleDateFormat.html> for more information on how to format the output using the `pattern` attribute of the `convertDateTime` tag.

TABLE 8-1 `convertDateTime` Tag Attributes

Attribute	Type	Description
<code>binding</code>	<code>DateTimeConverter</code>	Used to bind a converter to a backing bean property.
<code>dateStyle</code>	<code>String</code>	Defines the format, as specified by <code>java.text.DateFormat</code> , of a date or the date part of a date string. Applied only if <code>type</code> is <code>date</code> (or both) and <code>pattern</code> is not defined. Valid values: <code>default</code> , <code>short</code> , <code>medium</code> , <code>long</code> , and <code>full</code> . If no value is specified, <code>default</code> is used.
<code>locale</code>	<code>String</code> or <code>Locale</code>	<code>Locale</code> whose predefined styles for dates and times are used during formatting or parsing. If not specified, the <code>Locale</code> returned by <code>FacesContext.getLocale</code> will be used.
<code>pattern</code>	<code>String</code>	Custom formatting pattern that determines how the date/time string should be formatted and parsed. If this attribute is specified, <code>dateStyle</code> , <code>timeStyle</code> , and <code>type</code> attributes are ignored.

TABLE 8-1 convertDateTime Tag Attributes (Continued)

Attribute	Type	Description
timeStyle	String	Defines the format, as specified by <code>java.text.DateFormat</code> , of a time or the time part of a date string. Applied only if <code>type</code> is <code>time</code> and <code>pattern</code> is not defined. Valid values: <code>default</code> , <code>short</code> , <code>medium</code> , <code>long</code> , and <code>full</code> . If no value is specified, <code>default</code> is used.
timeZone	String or <code>TimeZone</code>	Time zone in which to interpret any time information in the date string.
type	String	Specifies whether the string value will contain a date, a time, or both. Valid values are <code>date</code> , <code>time</code> , or <code>both</code> . If no value is specified, <code>date</code> is used.
for	String	Used with composite components. Refers to one of the objects within the composite component inside which this tag is nested.

## Using NumberConverter

You can convert a component's data to a `java.lang.Number` by nesting the `convertNumber` tag inside the component tag. The `convertNumber` tag has several attributes that allow you to specify the format and type of the data. Table 8-2 lists the attributes.

The following example uses a `convertNumber` tag to display the total prices of the books in the shopping cart:

```
<h:outputText value="#{cart.total}" >
  <f:convertNumber type="currency"/>
</h:outputText>
```

When binding the `Number` converter to a component, ensure that the backing bean property to which the component is bound is of primitive type or has a type of `java.lang.Number`. In the case of the preceding example, `cart.total` is of type `java.lang.Number`.

Here is an example of a number that this tag can display:

\$934

This result can also be displayed using the following tag where currency pattern is specified:

```
<h:outputText id="cartTotal"
  value="#{cart.Total}" >
  <f:convertNumber pattern="
#####"
  />
</h:outputText>
```

See the *Customizing Formats* lesson of the Java Tutorial at <http://java.sun.com/docs/books/tutorial/i18n/format/decimalFormat.html> for more information on how to format the output using the pattern attribute of the `convertNumber` tag.

TABLE 8-2 convertNumber Attributes

Attribute	Type	Description
binding	NumberConverter	Used to bind a converter to a backing bean property.
currencyCode	String	ISO 4217 currency code, used only when formatting currencies.
currencySymbol	String	Currency symbol, applied only when formatting currencies.
groupingUsed	boolean	Specifies whether formatted output contains grouping separators.
integerOnly	boolean	Specifies whether only the integer part of the value will be parsed.
locale	String or Locale	Locale whose number styles are used to format or parse data.
maxFractionDigits	int	Maximum number of digits formatted in the fractional part of the output.
maxIntegerDigits	int	Maximum number of digits formatted in the integer part of the output.
minFractionDigits	int	Minimum number of digits formatted in the fractional part of the output.
minIntegerDigits	int	Minimum number of digits formatted in the integer part of the output.
pattern	String	Custom formatting pattern that determines how the number string is formatted and parsed.
type	String	Specifies whether the string value is parsed and formatted as a number, currency, or percentage. If not specified, number is used.
for	String	Used with composite components. Refers to one of the objects within the composite component inside which this tag is nested.

## Registering Listeners on Components

An application developer can implement listeners as classes or as backing bean methods. If a listener is a backing bean method, the page author references the method from either the component's `valueChangeListener` attribute or its `actionListener` attribute. If the listener is a class, the page author can reference the listener from either a `valueChangeListener` tag or an `actionListener` tag, and nest the tag inside the component tag, to register the listener on the component.

“Referencing a Method That Handles an Action Event” on page 169 and “Referencing a Method That Handles a Value-Change Event” on page 170 describe how a page author uses the `valueChangeListener` and `actionListener` attributes to reference backing bean methods that handle events.

This section explains how to register the `NameChanged` value-change listener and a hypothetical `LocaleChange` action listener implementation on components. Implementing value-change listeners, and implementing action listeners are advanced topics that are covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

## Registering a Value-Change Listener on a Component

A `ValueChangeListener` implementation can be registered on a component that implements `EditableValueHolder` by nesting a `valueChangeListener` tag within the component’s tag on the page. The `valueChangeListener` tag supports two attributes:

- `type`: References the fully qualified class name of a `ValueChangeListener` implementation
- `binding`: References an object that implements `ValueChangeListener`

One of these attributes must be used to reference the value-change listener. The `type` attribute can accept a literal or a value expression. The `binding` attribute can only accept a value expression, which must point to a backing bean property that accepts and returns a `ValueChangeListener` implementation.

The following example shows a value-change listener registered on a component:

```
<h:inputText id="name" size="50" value="#{cashier.name}"
  required="true">
  <f:valueChangeListener type="listeners.NameChanged" />
</h:inputText>
```

In the example, the core tag `type` attribute specifies the custom `NameChanged` listener as the `ValueChangeListener` implementation which is registered on the name component.

After this component tag is processed and local values have been validated, its corresponding component instance will queue the `ValueChangeEvent` that is associated with the specified `ValueChangeListener` to the component.

The `binding` attribute is used to bind a `ValueChangeListener` implementation to a backing bean property. It works in a similar way to the `binding` attribute supported by the standard converter tags.

## Registering an Action Listener on a Component

A page author can register an `ActionListener` implementation on a `Command` component by nesting an `actionListener` tag within the component's tag on the page. Similarly to the `valueChangeListener` tag, the `actionListener` tag supports both the `type` and `binding` attributes. One of these attributes must be used to reference the action listener.

Here is an example of `commandLink` tag, that references an `ActionListener` implementation rather than a backing bean method:

```
<h:commandLink id="NAmerica" action="bookstore">
    <f:actionListener type="listeners.LocaleChange" />
</h:commandLink>
```

The `type` attribute of the `actionListener` tag specifies the fully qualified class name of the `ActionListener` implementation. Similarly to the `valueChangeListener` tag, the `actionListener` tag also supports the `binding` attribute.

## Using the Standard Validators

JavaServer Faces technology provides a set of standard classes and associated tags that page authors and application developers can use to validate a component's data. [Table 8–3](#) lists all the standard validator classes and the tags that allow you to use the validators from the page.

TABLE 8–3 The Validator Classes

Validator Class	Tag	Function
<code>DoubleRangeValidator</code>	<code>validateDoubleRange</code>	Checks whether the local value of a component is within a certain range. The value must be floating-point or convertible to floating-point.
<code>LengthValidator</code>	<code>validateLength</code>	Checks whether the length of a component's local value is within a certain range. The value must be a <code>java.lang.String</code> .
<code>LongRangeValidator</code>	<code>validateLongRange</code>	Checks whether the local value of a component is within a certain range. The value must be any numeric type or <code>String</code> that can be converted to a <code>long</code> .
<code>RegexValidator</code>	<code>validateRegEx</code>	Checks whether the local value of a component is a match against a regular expression from <code>java.util.regex</code> package.
<code>BeanValidator</code>	<code>validateBean</code>	Registers a bean validator for the component

TABLE 8-3 The Validator Classes (Continued)

Validator Class	Tag	Function
RequiredValidator	validateRequired	Ensures that the local value is not empty on a <code>EditableValueHolder</code> component

Similar to the standard converters, each of these validators has one or more standard error messages associated with it. If you have registered one of these validators onto a component on your page, and the validator is unable to validate the component's value, the validator's error message will display on the page. For example, the error message that displays when the component's value exceeds the maximum value allowed by `LongRangeValidator` is as follows:

```
{1}: Validation Error: Value is greater than allowable maximum of "{0}"
```

In this case, the {1} substitution parameter is replaced by the component's label or id, and the {0} substitution parameter is replaced with the maximum value allowed by the validator.

## Validating a Component's Value

To validate a component's value using a particular validator, you need to register that validator on the component. You can do this in one of the following ways:

- Nest the validator's corresponding tag (shown in [Table 8-3](#)) inside the component's tag. [“Using the LongRangeValidator” on page 167](#) describes how to use the `validateLongRange` tag. You can use the other standard tags in the same way.
- Refer to a method that performs the validation from the component tag's `validator` attribute.
- Nest a validator tag inside the component tag and use either the validator tag's `validatorId` attribute or its `binding` attribute to refer to the validator.

See [“Referencing a Method That Performs Validation” on page 169](#) for more information on using the `validator` attribute.

The `validatorId` attribute works similarly to the `converterId` attribute of the `converter` tag, as described in [“Converting a Component's Value” on page 160](#).

Keep in mind that validation can be performed only on components that implement `EditableValueHolder` because these components accept values that can be validated.

## Using the LongRangeValidator

The following example shows how to use the `validateLongRange` validator on an input component tag `quantity`:

```

<h:inputText id="quantity" size="4"
  value=
  "#{item.quantity}"
  >
  <f:validateLongRange minimum="1"/>
</h:inputText>
<h:message for="quantity"/>

```

This tag requires that the user enter a number that is at least 1. The `size` attribute specifies that the number can have no more than four digits. The `validateLongRange` tag also has a `maximum` attribute, with which you can set a maximum value of the input.

The attributes of all the standard validator tags accept EL value expressions. This means that the attributes can reference backing bean properties rather than specify literal values. For example, the `validateLongRange` tag in the preceding example can reference a backing bean property called `minimum` to get the minimum value acceptable to the validator implementation as shown here:

```
<f:validateLongRange minimum="#{ShowCartBean.minimum}" />
```

## Referencing a Backing Bean Method

A component tag has a set of attributes for referencing backing bean methods that can perform certain functions for the component associated with the tag. These attributes are summarized in [Table 8–4](#).

**TABLE 8–4** Component Tag Attributes That Reference Backing Bean Methods

Attribute	Function
<code>action</code>	Refers to a backing bean method that performs navigation processing for the component and returns a logical outcome <code>String</code>
<code>actionListener</code>	Refers to a backing bean method that handles action events
<code>validator</code>	Refers to a backing bean method that performs validation on the component's value
<code>valueChangeListener</code>	Refers to a backing bean method that handles value-change events

Only components that implement `ActionSource` can use the `action` and `actionListener` attributes. Only components that implement `EditableValueHolder` can use the `validator` or `valueChangeListener` attributes.

The component tag refers to a backing bean method using a method expression as a value of one of the attributes. The method referenced by an attribute must follow a particular signature, which is defined by the tag attribute's definition in the TLD. For example, the definition of the `validator` attribute of the `inputText` tag in `html_basic.tld` is the following:

```
void validate(javax.faces.context.FacesContext,
             javax.faces.component.UIComponent, java.lang.Object)
```

The following sections give examples of how to use the four different attributes.

- [“Referencing a Method That Performs Navigation” on page 169](#)
- [“Referencing a Method That Handles an Action Event” on page 169](#)
- [“Referencing a Method That Performs Validation” on page 169](#)
- [“Referencing a Method That Handles a Value-Change Event” on page 170](#)

## Referencing a Method That Performs Navigation

If your page includes a component (such as a button or hyperlink) that causes the application to navigate to another page when the component is activated, the tag corresponding to this component must include an `action` attribute. This attribute does one of the following:

- Specifies a logical outcome `String` that tells the application which page to access next
- References a backing bean method that performs some processing and returns a logical outcome `String`

The following examples shows how to reference a navigation method:

```
<h:commandButton
    value="#{bundle.Submit}"
    action="#{cashier.submit}" />
```

## Referencing a Method That Handles an Action Event

If a component on your page generates an action event, and if that event is handled by a backing bean method, you refer to the method by using the component’s `actionListener` attribute.

The following example shows how the method is referenced:

```
<h:commandLink id="NAmerica" action="bookstore"
    actionListener="#{localeBean.chooseLocaleFromLink}">
```

The `actionListener` attribute of this component tag references the `chooseLocaleFromLink` method using a method expression. The `chooseLocaleFromLink` method handles the event when user clicks the hyperlink rendered by this component.

## Referencing a Method That Performs Validation

If the input of one of the components on your page is validated by a backing bean method, refer to the method from the component’s tag using the `validator` attribute.

The following example shows how to reference a method that performs validation on email, an input component:

```
<h:inputText id="email" value="#{checkoutFormBean.email}"
  size="25" maxLength="125"
  validator="#{checkoutFormBean.validateEmail}"/>
```

## Referencing a Method That Handles a Value-Change Event

If you want a component on your page to generate a value-change event and you want that event to be handled by a backing bean method, you refer to the method using the component's `valueChangeListener` attribute.

The following example shows how a component references a `ValueChangeListener` implementation that handles the event when a user enters a name in the name input field:

```
<h:inputText
  id="name"
  size="50"
  value="#{cashier.name}"
  required="true">
  <f:valueChangeListener type="listeners.NameChanged" />
</h:inputText>
```

To refer to this backing bean method, the tag uses the `valueChangeListener` attribute:

```
<h:inputText
  id="name"
  size="50"
  value="#{cashier.name}"
  required="true"
  valueChangeListener="#{cashier.processValueChange}" />
</h:inputText>
```

The `valueChangeListener` attribute of this component tag references the `processValueChange` method of `CashierBean` using a method expression. The `processValueChange` method handles the event of a user entering a name in the input field rendered by this component.

# Developing With JavaServer™ Faces Technology

---

The previous chapters, [Chapter 7, “Using JavaServer Faces Technology in Web Pages,”](#) and [Chapter 8, “Using Converters, Listeners and Validators,”](#) show how to add components to a page, connect them to server-side objects using the component tags, and how to provide additional functionality to the components through converters, listeners, and validators. Developing a JavaServer Faces application also involves the task of programming the server-side objects. These objects include backing beans, converters, event handlers, and validators.

This chapter provides an overview of the backing beans, and explains how to write methods and properties of backing beans that are used by a JavaServer Faces application. It also introduces the new bean validation (JSR 303) feature.

## Backing Beans

A typical JavaServer Faces application includes one or more backing beans, each of which is a type of JavaServer Faces managed bean that can be associated with the components used in a particular page. This section introduces the basic concepts on creating, configuring, and using backing beans in an application.

### Creating a Backing Bean

A backing bean is created with a constructor with no arguments (like all JavaBeans™ components), and also a set of properties and a set of methods that perform functions for a component.

Each of the component properties can be bound to one of the following:

- A component’s value
- A component instance
- A converter instance

- A listener instance
- A validator instance

The most common functions that backing bean methods perform include the following:

- Validating a component's data
- Handling an event fired by a component
- Performing processing to determine the next page to which the application must navigate

As with all JavaBeans components, a property consists of a private data field and a set of accessor methods, as shown by this code:

```
Integer userNumber = null;
...
public void setUserNumber(Integer user_number) {
    userNumber = user_number;
}
public Integer getUserNumber() {
    return userNumber;
}
public String getResponse() {
    ...
}
```

When a bean property is bound to a component's value, it can be any of the basic primitive and numeric types, or any Java object type for which the application has access to an appropriate converter. For example, a property can be of type `Date` if the application has access to a converter that can convert the `Date` type to a `String` and back again. See [“Writing Bean Properties” on page 174](#) for information on which types are accepted by which component tags.

When a bean property is bound to a component instance, the property's type must be the same as the component object. For example, if a `SelectBoolean` component is bound to the property, the property must accept and return a `SelectBoolean` object.

Likewise, if the property is bound to a converter, validator, or listener instance, then the property must be of the appropriate converter, validator, or listener type.

For more information on writing beans and their properties, see [“Writing Bean Properties” on page 174](#).

## Using the EL to Reference Backing Beans

To bind component values and objects to backing bean properties or to reference backing bean methods from component tags, page authors use the unified expression language (EL) syntax. As explained in “[Overview of EL](#)” on page 113, the following are some of the features that EL offers:

- Deferred evaluation of expressions
- The ability to use a value expression to both read and write data
- Method expressions

Deferred evaluation of expressions is important because the JavaServer Faces lifecycle is split into several phases where component event handling, data conversion and validation, and data propagation to external objects are all performed in an orderly fashion. The implementation must be able to delay the evaluation of expressions until the proper phase of the life cycle has been reached. Therefore, its tag attributes always use deferred evaluation syntax, which is distinguished by the `{ }` delimiter.

In order to store data in external objects, almost all JavaServer Faces tag attributes use lvalue value expressions, which are expressions that allow both getting and setting data on external objects.

Finally, some component tag attributes accept method expressions that reference methods that handle component events, or validate or convert component data.

To illustrate a JavaServer Faces tag using EL, let’s suppose that a tag of an application referenced a method to perform the validation of user input:

```
<h:inputText id="userNo"
    value="#{UserNumberBean.userNumber}"
    validator="#{UserNumberBean.validate}" />
```

This tag binds the `userNo` component’s value to the `UserNumberBean.userNumber` backing bean property using an lvalue expression. It uses a method expression to refer to the `UserNumberBean.validate` method, which performs validation of the component’s local value. The local value is whatever the user enters into the field corresponding to this tag. This method is invoked when the expression is evaluated, which is during the process validation phase of the life cycle.

Nearly all JavaServer Faces tag attributes accept value expressions. In addition to referencing bean properties, value expressions can also reference lists, maps, arrays, implicit objects, and resource bundles.

Another use of value expressions is binding a component instance to a backing bean property. A page author does this by referencing the property from the binding attribute:

```
<inputText binding="#{UserNumberBean.userNoComponent}" />
```

In addition to using expressions with the standard component tags, you can also configure your custom component properties to accept expressions by creating `ValueExpression` or `MethodExpression` instances for them.

For information on EL, see [Chapter 6, “Unified Expression Language.”](#)

For information on referencing backing bean methods from component tags, see [“Referencing a Backing Bean Method” on page 168.](#)

## Writing Bean Properties

As explained in [“Backing Beans” on page 171](#), a backing bean property can be bound to one of the following items:

- A component value
- A component instance
- A converter implementation
- A listener implementation
- A validator implementation

These properties follow the conventions of JavaBeans components (also called beans). For more information on JavaBeans components, see [JavaBeans Tutorial](#).

The component’s tag binds the component’s value to a backing bean property using its `value` attribute and binds the component’s instance to a backing bean property using its `binding` attribute.

Likewise, all the converter, listener, and validator tags use their `binding` attributes to bind their associated implementations to backing bean properties. Binding is an advanced topic covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

To bind a component’s value to a backing bean property, the type of the property must match the type of the component’s value to which it is bound. For example, if a backing bean property is bound to a `SelectBoolean` component’s value, the property should accept and return a `boolean` value or a `Boolean` wrapper `Object` instance.

To bind a component instance to a backing bean property, the property must match the type of component. For example, if a backing bean property is bound to a `SelectBoolean` instance, the property should accept and return `SelectBoolean` value.

Similarly, to bind a converter, listener, or validator implementation to a backing bean property, the property must accept and return the same type of converter, listener, or validator object. For example, if you are using the `convertDateTime` tag to bind a `DateTime` converter to a property, that property must accept and return a `DateTime` instance.

The rest of this section explains how to write properties that can be bound to component values, to component instances for the component objects described in “[Adding Components to a Page Using HTML Tags](#)” on page 128, and to converter, listener, and validator implementations.

## Writing Properties Bound to Component Values

To write a backing bean property that is bound to a component’s value, you must match the property type to the component’s value .

Table 9–1 lists the component classes described in “[Adding Components to a Page Using HTML Tags](#)” on page 128 and the acceptable types of their values.

TABLE 9–1 Acceptable Types of Component Values

Component	Acceptable Types of Component Values
Input, Output, SelectItem, SelectOne	Any of the basic primitive and numeric types or any Java programming language object type for which an appropriate Converter implementation is available
Data	array of beans, List of beans, single bean, java.sql.ResultSet, javax.servlet.jsp.jstl.sql.Result, javax.sql.RowSet
SelectBoolean	boolean or Boolean
SelectItems	java.lang.String, Collection, Array, Map
SelectMany	array or List though elements of the array or List can be any of the standard types

When page authors bind components to properties using the value attributes of the component tags, they need to ensure that the corresponding properties match the types of the components’ values.

### Input and Output Properties

In the following example, an h:inputText tag binds the value of component to the name property of a backing bean called CashierBean.

```
<h:inputText id="name" size="50"
  value="#{cashier.name}"

</h:inputText>
```

The following code segment from the backing bean CashierBean, shows the bean property type bound by the preceding component tag:

```
protected String name = null;
public void setName(String name) {
    this.name = name;
}
public String getName() {
    return this.name;
}
```

As described in [“Using the Standard Converters” on page 159](#), to convert the value of a Input or Output component, you can either apply a converter or create the bean property bound to the component with the matching type.

Here is the example tag from [“Using DateTimeConverter” on page 161](#) that displays the date that books will be shipped.

```
<h:outputText value="#{cashier.shipDate}">
    <f:convertDateTime dateStyle="full" />
</h:outputText>
```

The bean property represented by this tag must be of a type of `java.util.Date`. The following code segment shows the `shipDate` property from the backing bean `CashierBean`, that is bound by the tag's value in the preceding example:

```
protected Date shipDate;
public Date getShipDate() {
    return this.shipDate;
}
public void setShipDate(Date shipDate) {
    this.shipDate = shipDate;
}
```

## Data Properties

Data components must be bound to one of the backing bean property types listed in [Table 9–1](#). The Data component is discussed in [“Using Data-Bound Table Components” on page 148](#). Here is part of the start tag of `dataTable` from that section:

```
<h:dataTable id="items"
    ...
    value="#{cart.items}"
    var="item" >
```

The value expression points to the `items` property of a shopping cart bean named `cart`. The cart bean maintains a map of `ShoppingCartItem` beans.

The `getItems` method from cart bean populates a `List` with `ShoppingCartItem` instances that are saved in the `items` map from when the customer adds books to the cart, as shown in the following code segment:

```
public synchronized List.getItems() {
    List results = new ArrayList();
    results.addAll(this.items.values());
    return results;
}
```

All the components contained in the Data component are bound to the properties of the cart bean that is bound to the entire Data component. For example, here is the `h:outputText` tag that displays the book title in the table:

```
<h:commandLink action="#{showcart.details}">
    <h:outputText value="#{item.item.title}"/>
</h:commandLink>
```

## SelectBoolean Properties

Backing bean properties that hold the `SelectBoolean` component's data must be of `boolean` or `Boolean` type. The example `selectBooleanCheckbox` tag from the section [“Displaying Components for Selecting One Value” on page 143](#) binds a component to a property. The following example shows a tag that binds a component value to a `boolean` property:

```
<h:selectBooleanCheckbox title="#{bundle.receiveEmails}"
    value="#{custFormBean.receiveEmails}" >
</h:selectBooleanCheckbox>
<h:outputText value="#{bundle.receiveEmails}">
```

Here is an example property that can be bound to the component represented by the example tag:

```
protected boolean receiveEmails = false;
...
public void setReceiveEmails(boolean receiveEmails) {
    this.receiveEmails = receiveEmails;
}
public boolean getReceiveEmails() {
    return receiveEmails;
}
```

## SelectMany Properties

Because a `SelectMany` component allows a user to select one or more items from a list of items, this component must map to a bean property of type `List` or `array`. This bean property represents the set of currently selected items from the list of available items.

The following example of `selectManyCheckbox` tag comes from [“Rendering Components for Selecting Multiple Values” on page 145](#):

```
<h:selectManyCheckbox
  id="newsletters"
  layout="pageDirection"
  value="#{cashier.newsletters}">
  <f:selectItems value="#{newsletters}"/>
</h:selectManyCheckbox>
```

Here is the bean property that maps to the value of the `selectManyCheckbox` tag from the preceding example:

```
protected String newsletters[] = new String[0];

public void setNewsletters(String newsletters[]) {
    this.newsletters = newsletters;
}
public String[] getNewsletters() {
    return this.newsletters;
}
```

The `SelectItem` and `SelectItems` components are used to represent all the values in a `SelectMany` component. See [“SelectItem Properties” on page 179](#) and [“SelectItems Properties” on page 179](#) for information on writing the bean properties for the `SelectItem` and `SelectItems` components.

## SelectOne Properties

`SelectOne` properties accept the same types as `Input` and `Output` properties, because a `SelectOne` component represents the single selected item from a set of items. This item can be any of the primitive types and anything else for which you can apply a converter.

Here is an example of the `selectOneMenu` tag from [“Displaying a Menu Using the h:selectOneMenu Tag” on page 144](#):

```
<h:selectOneMenu id="shippingOption"
  required="true"
  value="#{cashier.shippingOption}">
  <f:selectItem
    itemValue="2"
    itemLabel="#{bundle.QuickShip}"/>
  <f:selectItem
    itemValue="5"
    itemLabel="#{bundle.NormalShip}"/>
  <f:selectItem
    itemValue="7"
    itemLabel="#{bundle.SaverShip}"/>
</h:selectOneMenu>
```

Here is the bean property corresponding to this tag:

```
protected String shippingOption = "2";

public void setShippingOption(String shippingOption) {
    this.shippingOption = shippingOption;
}
public String getShippingOption() {
    return this.shippingOption;
}
```

Note that `shippingOption` represents the currently selected item from the list of items in the `SelectOne` component.

The `SelectItem` and `SelectItems` components are used to represent all the values in a `SelectOne` component. This is explained in the section [“Displaying a Menu Using the `h:selectOneMenu` Tag” on page 144](#).

For information on how to write the backing bean properties for the `SelectItem` and `SelectItems` components, see [“SelectItem Properties” on page 179](#) and [“SelectItems Properties” on page 179](#).

## SelectItem Properties

A `SelectItem` component represents a single value in a set of values in a `SelectMany` or `SelectOne` component. A `SelectItem` component can be bound to a backing bean property of type `SelectItem`. A `SelectItem` object is composed of an `Object` representing the value, along with two `Strings` representing the label and description of the `SelectItem` object.

The example `selectOneMenu` tag from [“Displaying a Menu Using the `h:selectOneMenu` Tag” on page 144](#) contains `selectItem` tags that set the values of the list of items in the page. Here is an example of a bean property that can set the values for this list in the bean:

```
SelectItem itemOne = null;

SelectItem getItemOne(){
    return itemOne;
}

void setItemOne>SelectItem item) {
    itemOne = item;
}
```

## SelectItems Properties

The `SelectItems` components are children of `SelectMany` and `SelectOne` components. Each `SelectItems` component is composed of a set of either `SelectItem` instances or any collection of objects such as an array, or a list or even POJOs..

The following section describes how to write the properties for `selectItems` tags containing `SelectItem` instances.

## Properties for `SelectItems` Composed of `SelectItem` Instances

You can populate the `SelectItems` with `SelectItem` instances programmatically in the backing bean.

1. In your backing bean, create a list that is bound to the `SelectItem` component.
2. Then define a set of `SelectItem` objects, set their values, and populate the list with the `SelectItem` objects.

Here is an example code snippet from a backing bean that shows how to create a `SelectItems` property:

```
import javax.faces.component.SelectItem;
...
protected ArrayList options = null;
protected SelectItem newsletter0 =
    new SelectItem("200", "Duke's Quarterly", "");
...
//in constructor, populate the list
options.add(newsletter0);
options.add(newsletter1);
options.add(newsletter2);
...
public SelectItem getNewsletter0(){
    return newsletter0;
}

void setNewsletter0(SelectItem firstNL) {
    newsletter0 = firstNL;
}
// Other SelectItem properties

public Collection[] getOptions(){
    return options;
}
public void setOptions(Collection[] options){
    this.options = new ArrayList(options);
}
```

The code first initializes `options` as a list. Each newsletter property is defined with values. Then, each newsletter `SelectItem` is added to the list. Finally, the code includes the obligatory `setOptions` and `getOptions` accessor methods.

## Writing Properties Bound to Component Instances

A property bound to a component instance returns and accepts a component instance rather than a component value. The following components bind a component instance to a backing bean property:

```
<h:selectBooleanCheckbox
    id="fanClub"
    rendered="false"
    binding="#{cashier.specialOffer}" />
<h:outputLabel for="fanClub"
    rendered="false"
    binding="#{cashier.specialOfferText}" >
    <h:outputText id="fanClubLabel"
        value="#{bundle.DukeFanClub}" />
</h:outputLabel>
```

The `selectBooleanCheckbox` tag renders a check box and binds the `fanClub` `SelectBoolean` component to the `specialOffer` property of `CashierBean`. The `outputLabel` tag binds the `fanClubLabel` component (which represents the check box's label) to the `specialOfferText` property of `CashierBean`. If the user orders more than \$100 worth of books and clicks the Submit button, the `submit` method of `CashierBean` sets both components' `rendered` properties to `true`, causing the check box and label to display when the page is re-rendered.

Because the components corresponding to the example tags are bound to the backing bean properties, these properties must match the components' types. This means that the `specialOfferText` property must be of `Output` type, and the `specialOffer` property must be of `SelectBoolean` type:

```
UIOutput specialOfferText = null;

public UIOutput getSpecialOfferText() {
    return this.specialOfferText;
}
public void setSpecialOfferText(UIOutput specialOfferText) {
    this.specialOfferText = specialOfferText;
}

UISelectBoolean specialOffer = null;

public UISelectBoolean getSpecialOffer() {
    return this.specialOffer;
}
public void setSpecialOffer(UISelectBoolean specialOffer) {
    this.specialOffer = specialOffer;
}
```

For more general information on component binding, see [“Backing Beans” on page 171](#).

For information on how to reference a backing bean method that performs navigation when a button is clicked, see [“Referencing a Method That Performs Navigation” on page 169](#).

For more information on writing backing bean methods that handle navigation, see [“Writing a Method to Handle Navigation” on page 183](#).

## Writing Properties Bound to Converters, Listeners, or Validators

All of the standard converter, listener, and validator tags that are included with JavaServer Faces technology support binding attributes that allow binding converter, listener, or validator implementations to backing bean properties.

The following example shows a standard `convertDateTime` tag using a value expression with its binding attribute to bind the `DateTimeConverter` instance to the `convertDate` property of `LoginBean`:

```
<h:inputText value="#{LoginBean.birthDate}">
  <f:convertDateTime binding="#{LoginBean.convertDate}" />
</h:inputText>
```

The `convertDate` property must therefore accept and return a `DateTimeConverter` object, as shown here:

```
private DateTimeConverter convertDate;
public DateTimeConverter getConvertDate() {
    ...
    return convertDate;
}
{
public void setConvertDate(DateTimeConverter convertDate) {
    convertDate.setPattern("EEEEEEEE, MMM dd, yyyy");
    this.convertDate = convertDate;
}
```

Because the converter is bound to a backing bean property, the backing bean property can modify the attributes of the converter or add new functionality to it. In the case of the preceding example, the property sets the date pattern that the converter uses to parse the user’s input into a `Date` object.

The backing bean properties that are bound to validator or listener implementations are written in the same way and have the same general purpose.

## Writing Backing Bean Methods

Methods of a backing bean can perform several application-specific functions for components on the page. These functions include:

- Performing validation on the component's value
- Handling action events
- Handling value-change events
- Performing processing associated with navigation

By using a backing bean to perform these functions, you eliminate the need to implement the `Validator` interface to handle the validation or the `Listener` interface to handle events. Also, by using a backing bean instead of a `Validator` implementation to perform validation, you eliminate the need to create a custom tag for the `Validator` implementation. Creating custom validators is an advanced topic covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

In general, it's good practice to include these methods in the same backing bean that defines the properties for the components referencing these methods. The reason for doing so is that the methods might need to access the component's data to determine how to handle the event or to perform the validation associated with the component.

This section describes the requirements for writing the backing bean methods. The following topics explain writing different types of backing bean methods:

- [“Writing a Method to Handle Navigation” on page 183](#)
- [“Writing a Method to Handle an Action Event” on page 185](#)
- [“Writing a Method to Perform Validation” on page 185](#)
- [“Writing a Method to Handle a Value-Change Event” on page 186](#)

### Writing a Method to Handle Navigation

A backing bean method that handles navigation processing, called an action method, must be a public method that takes no parameters and returns an `Object`, which is the logical outcome that the navigation system uses to determine the page to display next. This method is referenced using the component tag's `action` attribute.

The following action method is from a backing bean named `CashierBean`, which is invoked when a user clicks the Submit button on the page. If the user has ordered more than \$100 worth of books, this method sets the rendered properties of the `fanClub` and `specialOffer` components to `true`, causing them to be displayed on the page the next time that page is rendered.

After setting the components' rendered properties to `true`, this method returns the logical outcome `null`. This causes the JavaServer Faces implementation to re-render the page without creating a new view of the page, retaining the customer's input. If this method were to return `purchase` which is the logical outcome to use to advance to a payment page, the page would re-render without retaining the customer's input.

If the user does not purchase more than \$100 worth of books or the `thankYou` component has already been rendered, the method returns `receipt`. The JavaServer Faces implementation loads the page after this method returns.

```
public String submit() {
    ...
    if(cart().getTotal() > 100.00 &&
        !specialOffer.isRendered())
    {
        specialOfferText.setRendered(true);
        specialOffer.setRendered(true);
        return null;
    } else if (specialOffer.isRendered() &&
        !thankYou.isRendered()){
        thankYou.setRendered(true);
        return null;
    } else {
        clear();
        return "receipt";
    }
}
```

Typically, an action method will return a `String` outcome, as shown in the previous example. Alternatively, you can define an `Enum` class that encapsulates all possible outcome strings, and then make an action method return an enum constant, which represents a particular `String` outcome defined by the `Enum` class. In this case, the value returned by a call to the `Enum` class's `toString` method must match that specified by the `from-outcome` element in the appropriate navigation rule configuration defined in the application configuration file.

The following example uses an `Enum` class to encapsulate all logical outcomes:

```
public enum Navigation {
    main, accountHist, accountList, atm, atmAck, transferFunds,
    transferAck, error
}
```

When an action method returns an outcome, it uses the dot notation to reference the outcome from the `Enum` class:

```
public Object submit(){
    ...
    return Navigation.accountHist;
}
```

The section [“Referencing a Method That Performs Navigation” on page 169](#) explains how a component tag references this method. The section [“Writing Properties Bound to Component Instances” on page 181](#) discusses how to write the bean properties to which the components are bound.

## Writing a Method to Handle an Action Event

A backing bean method that handles an action event must be a public method that accepts an action event and returns `void`. This method is referenced using the component tag's `actionListener` attribute. Only components that implement `ActionSource` can refer to this method.

In the following example, a method from a backing bean named `LocaleBean` processes the event of a user clicking one of the hyperlinks on the page:

```
public void chooseLocaleFromLink(ActionEvent event) {
    String current = event.getComponent().getId();
    FacesContext context = FacesContext.getCurrentInstance();
    context.getViewRoot().setLocale((Locale)
        locales.get(current));
}
```

This method gets the component that generated the event from the event object, and then it gets the component's ID, which indicates a region of the world. The method matches the ID against a `HashMap` object that contains the locales available for the application. Finally, it sets the locale using the selected value from the `HashMap` object.

[“Referencing a Method That Handles an Action Event” on page 169](#) explains how a component tag references this method.

## Writing a Method to Perform Validation

Instead of implementing the `Validator` interface to perform validation for a component, you can include a method in a backing bean to take care of validating input for the component.

A backing bean method that performs validation must accept a `FacesContext`, the component whose data must be validated, and the data to be validated, just as the `validate` method of the `Validator` interface does. A component refers to the backing bean method by using its `validator` attribute. Only values of `Input` components or values of components that extend `Input` can be validated.

Here is an example of a backing bean method that validates user input:

```
public void validateEmail(FacesContext context,
    UIComponent toValidate, Object value) {

    String message = "";
    String email = (String) value;
    if (email.contains('@')) {
        ((UIInput)toValidate).setValid(false);
        message = CoffeeBreakBean.loadErrorMessage(context,
```

```

        CoffeeBreakBean.CB_RESOURCE_BUNDLE_NAME,
        "EmailError");
    context.addMessage(toValidate.getClientId(context),
        new FacesMessage(message));
    }
}

```

Let's take a closer look at the above code segment:

1. The `validateEmail` method first gets the local value of the component.
2. It then checks whether the `@` character is contained in the value.
3. If not, the method sets the component's `valid` property to `false`.
4. The method then loads the error message and queues it onto the `FacesContext` instance, associating the message with the component ID.

See [“Referencing a Method That Performs Validation” on page 169](#) for information on how a component tag references this method.

## Writing a Method to Handle a Value-Change Event

A backing bean that handles a value-change event must use a public method that accepts a value-change event and returns `void`. This method is referenced using the component's `valueChangeListener` attribute.

This section explains how to write a backing bean method to replace the `ValueChangeListener` implementation.

The following example tag comes from [“Registering a Value-Change Listener on a Component” on page 165](#), where the `h:inputText` tag with the `id` of `name`, has a `ValueChangeListener` instance registered on it. This `ValueChangeListener` instance handles the event of entering a value in the field corresponding to the component. When the user enters a value, a value-change event is generated, and the `processValueChange(ValueChangeEvent)` method of the `ValueChangeListener` class is invoked.

```

<h:inputText id="name" size="50" value="#{cashier.name}"
    required="true">
    <f:valueChangeListener type="listeners.NameChanged" />
</h:inputText>

```

Instead of implementing `ValueChangeListener`, you can write a backing bean method to handle this event. To do this, you move the `processValueChange(ValueChangeEvent)` method from the `ValueChangeListener` class, called `NameChanged`, to your backing bean.

Here is the backing bean method that processes the event of entering a value in the `name` field on the page:

```

public void processValueChange(ValueChangeEvent event)
    throws AbortProcessingException {
    if (null != event.getNewValue()) {
        FacesContext.getCurrentInstance().
            getExternalContext().getSessionMap().
                put("name", event.getNewValue());
    }
}

```

To make this method handle the `ValueChangeEvent` that is generated by an Input component, reference this method from the component tag's `valueChangeListener` attribute. See [“Referencing a Method That Handles a Value-Change Event” on page 170](#) for more information.

## Bean Validation

Bean validation (JSR 303) is a new feature that is available in Java EE 6. A JavaServer Faces 2.0 implementation must support bean validation if the server runtime (such as Java EE 6) requires it.

Validation can take place at different layers in even the simplest of applications, as shown in the `guessNumber` example application from the earlier chapter. The `guessNumber` example application validates the user input (in the `<h:inputText>` tag) for numerical data at the presentation layer and for a valid range of numbers at the business layer.

The bean validation model is supported by constraints in the form of annotations placed on a field, method, or class of a JavaBeans component such as a backing bean.

Constraints can be built-in or user-defined. Several built-in annotations are available in the `javax.validation.constraints` package. Some of the commonly used built-in annotations are listed below:

- `@Min`: The annotated element must be a number whose value must be higher or equal to the specified minimum.
- `@Max`: The annotated element must be a number whose value must be lower or equal to the specified maximum.
- `@Size`: The annotated element must be between specified minimum and maximum boundaries.
- `@NotNull`: The annotated element must not be null.
- `@Null`: The annotated element must be null.
- `@Pattern`: The annotated element must match the specified Java regular expression.

For a complete list of built-in constraint annotations, see API documentation for `javax.validation.constraints` class at <http://java.sun.com/javaee/6/docs/api/>.

In the following example, a constraint is placed on a field using the built-in `@NotNull` constraint:

```
public class Name {
    @NotNull
    private String firstname;
    @NotNull
    private String lastname;
}
```

You can also place more than one constraint on a single JavaBeans component object. For example, you can place an additional constraint for size of field on the first name and the last name fields:

```
public class Name {
    @NotNull
    @Size(min=1, max=16)
    private String firstname;
    @NotNull
    @Size(min=1, max=16)
    private String lastname;
}
```

The following example shows a user-defined constraint placed on a method which checks for a predefined email address pattern such as a corporate email account:

```
@validEmail
public String getEmailAddress()
{
    return emailAddress;
}
```

A user-defined constraint also needs a validation implementation. For a built-in constraint, a default implementation is already available. Any validation failures are gracefully handled and can be displayed by `h:messages` tag.

# Configuring JavaServer Faces Applications

---

The responsibilities of the application architect include the following:

- Registering back-end objects with the application so that all parts of the application have access to them
- Configuring backing beans and model beans so that they are instantiated with the proper values when a page makes reference to them
- Defining navigation rules for each of the pages in the application so that the application has a smooth page flow
- Packaging the application to include all the pages, resources, and other files so that the application can be deployed on any compliant container

This chapter explains how to perform the responsibilities of the application architect.

## Application Configuration Resource File

JavaServer Faces technology provides a portable configuration format (as an XML document) for configuring resources. An application architect creates one or more files, called application configuration resource files, that use this format to register and configure objects and resources, and to define navigation rules. An application configuration resource file is usually called `faces-config.xml`.

The application configuration resource file must be valid against the schema located at [http://java.sun.com/xml/ns/javaee/web-facesconfig\\_2\\_0.xsd](http://java.sun.com/xml/ns/javaee/web-facesconfig_2_0.xsd).

In addition, each file must include the following, in this order:

- The XML version number:  
`<?xml version="1.0"?>`
- A `faces-config` tag enclosing all the other declarations:

```
<faces-config xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
    http://java.sun.com/xml/ns/javaee/web-facesconfig_2_0.xsd"
  version="2.0">
  ...
</faces-config>
```

You can have more than one application configuration resource file. The JavaServer Faces implementation finds the file or files by looking for the following:

- A resource named `/META-INF/faces-config.xml` in any of the JAR files in the web application's `/WEB-INF/lib/` directory and in parent class loaders. If a resource with this name exists, it is loaded as a configuration resource. This method is practical for a packaged library containing some components and renderers.
- A context initialization parameter, `javax.faces.application.CONFIG_FILES`, that specifies one or more (comma-delimited) paths to multiple configuration files for your web application. This method will most likely be used for enterprise-scale applications that delegate to separate groups the responsibility for maintaining the file for each portion of a big application.
- A resource named `faces-config.xml` in the `/WEB-INF/` directory of your application. This is the way that most simple applications will make their configuration files available.

To access resources registered with the application, an application developer uses an instance of the `Application` class, which is automatically created for each application. The `Application` instance acts as a centralized factory for resources that are defined in the XML file.

When an application starts up, the JavaServer Faces implementation creates a single instance of the `Application` class and configures it with the information that you provided in the application configuration resource file.

## Ordering of Application Configuration Resources

As the JavaServer Faces 2.0 specification allows use of multiple application configuration resource files, the order in which they will be loaded needs to be determined. This can be done through an `<ordering>` entry in the application configuration resource file itself. The ordering of application configuration resource files can be absolute or relative.

When using absolute ordering, the order in which application configuration resource files will be loaded is specified by user. The following example shows an entry for absolute ordering:

```
<faces-config>
<name>myJSF</name>
<absolute-ordering>
<name>A</name>
```

```
<name>B</name>  
<name>C</name>  
</absolute-ordering>  
</faces-config>
```

In this example, A, B, and C represent different application configuration resource files and are considered in that order.

When using relative ordering, the order in which application configuration resource files will be loaded is calculated by considering entries from different files. The following example shows some of these considerations:

File A:

```
<faces-config>  
<name>myJSF</name>  
<ordering>  
<before>  
<name>B</name>  
</before>  
</ordering>  
</faces-config>
```

File C:

```
<faces-config>  
<name>myJSF</name>  
<ordering>  
<after>  
<name>B</name>  
</after>  
</ordering>  
</faces-config>
```

In this example, A, B, and C represent different application configuration resource files that will be loaded in this order: A before B, and C after B.

If there is no ordering declared in an application configuration file, then the file will receive the lowest order.

## Configuring Project Stage

ProjectStage is a context parameter identifying the status of a JavaServer Faces application in the software lifecycle. The stage of an application can affect the behavior of the application.

The possible stage values are as follows:

- Production
- Development
- UnitTest
- SystemTest
- Extension

ProjectStage is configured through an init context parameter in the web deployment descriptor file. Here is an example:

```
<context-param>
<param-name>avax.faces.PROJECT_STAGE </param-name>
<param-value>Development </param-value>
</context-param>
```

If no ProjectStage is defined, the default stage is considered as Development. You can also add custom stages according to your requirement.

The ProjectStage value can also be configured through JNDI. JNDI configuration of ProjectStage is an advanced topic covered in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

## Configuring Beans

To instantiate backing beans and other managed beans used in a JavaServer Faces application and store them in scope, you use the managed bean creation facility. This facility is configured in the application configuration resource file using managed-bean XML elements to define each bean. This file is processed at application startup time. When a page references a bean, the JavaServer Faces implementation initializes it according to its configuration in the application configuration resource file or by @ManagedBean annotation in the bean class. For information on using this annotation, see [“Using Annotations” on page 201](#).

With the managed bean creation facility, you can:

- Create beans in one centralized file that is available to the entire application, rather than conditionally instantiate beans throughout the application
- Customize the bean’s properties without any additional code
- Customize the bean’s property values directly from within the configuration file so that it is initialized with these values when it is created
- Using value elements, set the property of one managed bean to be the result of evaluating another value expression

This section shows you how to initialize beans using the managed bean creation facility. See “Writing Bean Properties” on page 174 and “Writing Backing Bean Methods” on page 183 for information on programming backing beans.

## Using the managed - bean Element

You create a bean using a managed - bean element, which represents an instance of a bean class that must exist in the application. At runtime, the JavaServer Faces implementation processes the managed - bean element. If a page references the bean, and if no bean instance exists, the JavaServer Faces implementation instantiates the bean as specified by the element configuration.

Here is an example managed bean configuration:

```
<managed-bean>
  <managed-bean-name> UserNumberBean</managed-bean-name>
    <managed-bean-class>
      guessnNumber.UserNumberBean
    </managed-bean-class>
    <managed-bean-scope> session </managed-bean-scope>
    <managed-property>
      <property-name> maximum </property-name>
      <property-class> long </property-class>
      <value> 10 </value>
    </managed-property>
    . . .
  </managed-bean-name>
</managed-bean>
```

Using NetBeans IDE, you can add a managed bean declaration by doing the following:

1. After opening your project in NetBeans IDE, expand the project node in the Projects pane.
2. Expand the Web Pages and WEB-INF nodes of the project node.
3. Double-click `faces-config.xml`.
4. After `faces-config.xml` opens in the editor pane, right-click in the editor pane.
5. Select JavaServer Faces → Add Managed Bean.
6. In the Add Managed Bean dialog box:
  - a. Enter the display name of the bean in the Bean Name field.
  - b. Click Browse to locate the bean’s class.
7. In the Browse Class dialog box:
  - a. Start typing the name of the class that you are looking for in the Class Name field. While you are typing, the dialog will show the matching classes.
  - b. Select the class from the Matching Classes box.

- c. Click OK.
8. In the Add Managed Bean dialog box:
    - a. Select the bean's scope from the Scope menu.
    - b. Click Add.

The preceding steps will add the managed-bean element and three elements inside of that element: a managed-bean-name element, a managed-bean-class element, and a managed-bean-scope element. You will need to edit the XML of the configuration file directly to further configure this managed bean.

The managed-bean-name element defines the key under which the bean will be stored in a scope. For a component's value to map to this bean, the component tag's value attribute must match the managed-bean-name up to the first period.

The managed-bean-class element defines the fully qualified name of the JavaBeans component class used to instantiate the bean.

The managed-bean element can contain zero or more managed-property elements, each corresponding to a property defined in the bean class. These elements are used to initialize the values of the bean properties. If you don't want a particular property initialized with a value when the bean is instantiated, do not include a managed-property definition for it in your application configuration resource file.

If a managed-bean element does not contain other managed-bean elements, it can contain one map-entries element or list-entries element. The map-entries element configures a set of beans that are instances of Map. The list-entries element configures a set of beans that are instances of List.

To map to a property defined by a managed-property element, you must ensure that the part of a component tag's value expression after the period matches the managed-property element's property-name element. In the earlier example, the shape property is initialized with the value poly. The next section explains in more detail how to use the managed-property element.

## Using managed-bean Scopes

The managed-bean-scope element defines the scope in which the bean will be stored. The acceptable scopes are none, request, session, or application. In addition you have new two new scopes, view and custom, available in JavaServer Faces 2.0.

If you define the bean with a none scope, the bean is instantiated anew each time that it is referenced, and so it does not get saved in any scope. One reason to use a scope of none is that a managed bean references another managed bean. The second bean should be in none scope if it is supposed to be created only when it is referenced. See [“Initializing Managed Bean Properties” on page 199](#) for an example of initializing a managed bean property.

If you are configuring a backing bean that is referenced by a component tag's `binding` attribute, you should define the bean with a request scope. If you placed the bean in session or application scope instead, the bean would need to take precautions to ensure thread safety because `UIComponent` instances depend on running inside of a single thread.

If you are configuring a bean that allows attributes to be associated with the view, you can use the new `view` scope. The attributes persist until the user has navigated to the next view.

You can also create a custom scope which is a user-defined, non-standard scope. Its value must be configured as a map in the application configuration file or as an annotation entered in the `managed-bean` class.

## Eager application-scoped Beans

Managed beans are lazy instantiated, in the sense that they are instantiated when a request is made from the application.

If the `managed-bean` `eager` attribute is declared as `true`, then they are instantiated when the application is started and placed in the application scope.

Managed-beans can be declared as `eager` either through annotations:

```
@ManagedBean(eager=true)
```

The attribute can also be set from the application resource configuration file:

```
<managed-bean eager="true">
```

## Initializing Properties Using the `managed-property` Element

A `managed-property` element must contain a `property-name` element, which must match the name of the corresponding property in the bean. A `managed-property` element must also contain one of a set of elements (listed in [Table 10-1](#)) that defines the value of the property. This value must be of the same type as that defined for the property in the corresponding bean. Which element you use to define the value depends on the type of the property defined in the bean. [Table 10-1](#) lists all the elements that are used to initialize a value.

**TABLE 10-1** Sub-elements of `managed-property` Elements That Define Property Values

Element	Value That It Defines
<code>list-entries</code>	Defines the values in a list

**TABLE 10-1** Sub-elements of managed-property Elements That Define Property Values *(Continued)*

Element	Value That It Defines
map-entries	Defines the values of a map
null-value	Explicitly sets the property to null
value	Defines a single value, such as a String, int, or JavaServer Faces EL expression

“Using the managed-bean Element” on page 193 includes an example of initializing String properties using the value sub-element. You also use the value sub-element to initialize primitive and other reference types. The rest of this section describes how to use the value sub-element and other sub-elements to initialize properties of Java Enum types, java.util.Map, array, and Collection, as well as initialization parameters.

## Referencing a Java Enum Type

As of version 1.2 of JavaServer Faces technology, a managed bean property can also be a Java Enum type (see <http://java.sun.com/javase/6/docs/api/java/lang/Enum.html>). In this case, the value element of the managed-property element must be a String that matches one of the String constants of the Enum. In other words, the String must be one of the valid values that can be returned if you were to call `valueOf(Class, String)` on enum, where `Class` is the Enum class and `String` is the contents of the value subelement. For example, suppose the managed bean property is the following:

```
public enum Suit { Hearts, Spades, Diamonds, Clubs}
...
public Suit getSuit() { ... return Suit.Hearts; }
```

Assuming that you want to configure this property in the application configuration file, the corresponding managed-property element would look like this:

```
<managed-property>
  <property-name>Suit</property-name>
  <value>Hearts</value>
</managed-property>
```

When the system encounters this property, it iterates over each of the members of the enum and calls `toString()` on each member until it finds one that is exactly equal to the value from the value element.

## Referencing an Initialization Parameter

Another powerful feature of the managed bean creation facility is the ability to reference implicit objects from a managed bean property.

Suppose that you have a page that accepts data from a customer, including the customer's address. Suppose also that most of your customers live in a particular area code. You can make the area code component render this area code by saving it in an implicit object and referencing it when the page is rendered.

You can save the area code as an initial default value in the context `initParam` implicit object by adding a context parameter to your web application and setting its value in the deployment descriptor. For example, to set a context parameter called `defaultAreaCode` to `650`, add a `context-param` element to the deployment descriptor, and give the parameter the name `defaultAreaCode` and the value `650`.

Next, you write a `managed-bean` declaration that configures a property that references the parameter:

```
<managed-bean>
  <managed-bean-name>customer</managed-bean-name>
  <managed-bean-class>CustomerBean</managed-bean-class>
  <managed-bean-scope>request</managed-bean-scope>
  <managed-property>
    <property-name>areaCode</property-name>
    <value>#{initParam.defaultAreaCode}</value>
  </managed-property>
  ...
</managed-bean>
```

To access the area code at the time that the page is rendered, refer to the property from the area component tag's `value` attribute:

```
<h:inputText id=area value="#{customer.areaCode}"
```

Retrieving values from other implicit objects is done in a similar way.

## Initializing Map Properties

The `map-entries` element is used to initialize the values of a bean property with a type of `java.util.Map` if the `map-entries` element is used within a `managed-property` element. A `map-entries` element contains an optional `key-class` element, an optional `value-class` element, and zero or more `map-entry` elements.

Each of the `map-entry` elements must contain a `key` element and either a `null-value` or `value` element. Here is an example that uses the `map-entries` element:

```
<managed-bean>
  ...
  <managed-property>
    <property-name>prices</property-name>
    <map-entries>
```

```
<map-entry>
  <key>My Early Years: Growing Up on *7</key>
  <value>30.75</value>
</map-entry>
<map-entry>
  <key>Web Servers for Fun and Profit</key>
  <value>40.75</value>
</map-entry>
</map-entries>
</managed-property>
</managed-bean>
```

The map that is created from this `map-entries` tag contains two entries. By default, all the keys and values are converted to `java.lang.String`. If you want to specify a different type for the keys in the map, embed the `key-class` element just inside the `map-entries` element:

```
<map-entries>
  <key-class>java.math.BigDecimal</key-class>
  ...
</map-entries>
```

This declaration will convert all the keys into `java.math.BigDecimal`. Of course, you must make sure that the keys can be converted to the type that you specify. The key from the example in this section cannot be converted to a `java.math.BigDecimal` because it is a `String`.

If you also want to specify a different type for all the values in the map, include the `value-class` element after the `key-class` element:

```
<map-entries>
  <key-class>int</key-class>
  <value-class>java.math.BigDecimal</value-class>
  ...
</map-entries>
```

Note that this tag sets only the type of all the `value` subelements.

The first `map-entry` in the preceding example includes a `value` subelement. The `value` subelement defines a single value, which will be converted to the type specified in the bean.

The second `map-entry` defines a `value` element, which references a property on another bean. Referencing another bean from within a bean property is useful for building a system from fine-grained objects. For example, a request-scoped form-handling object might have a pointer to an application-scoped database mapping object. Together the two can perform a form-handling task. Note that including a reference to another bean will initialize the bean if it does not already exist.

Instead of using a `map-entries` element, it is also possible to assign the entire map using a `value` element that specifies a map-typed expression.

## Initializing Array and List Properties

The `list-entries` element is used to initialize the values of an array or `List` property. Each individual value of the array or `List` is initialized using a `value` or `null-value` element. Here is an example:

```
<managed-bean>
  ...
  <managed-property>
    <property-name>books</property-name>
    <list-entries>
      <value-class>java.lang.String</value-class>
      <value>Web Servers for Fun and Profit</value>
      <value>#{myBooks.bookId[3]}</value>
      <null-value/>
    </list-entries>
  </managed-property>
</managed-bean>
```

This example initializes an array or a `List`. The type of the corresponding property in the bean determines which data structure is created. The `list-entries` element defines the list of values in the array or `List`. The `value` element specifies a single value in the array or `List` and can reference a property in another bean. The `null-value` element will cause the `setBooks` method to be called with an argument of `null`. A `null` property cannot be specified for a property whose data type is a Java primitive, such as `int` or `boolean`.

## Initializing Managed Bean Properties

Sometimes you might want to create a bean that also references other managed beans so that you can construct a graph or a tree of beans. For example, suppose that you want to create a bean representing a customer's information, including the mailing address and street address, each of which is also a bean. The following `managed-bean` declarations create a `CustomerBean` instance that has two `AddressBean` properties: one representing the mailing address, and the other representing the street address. This declaration results in a tree of beans with `CustomerBean` as its root and the two `AddressBean` objects as children.

```
<managed-bean>
  <managed-bean-name>customer</managed-bean-name>
  <managed-bean-class>
    com.mycompany.mybeans.CustomerBean
  </managed-bean-class>
  <managed-bean-scope> request </managed-bean-scope>
  <managed-property>
    <property-name>mailingAddress</property-name>
    <value>#{addressBean}</value>
  </managed-property>
  <managed-property>
```

```

        <property-name>streetAddress</property-name>
        <value>#{addressBean}</value>
    </managed-property>
</managed-property>
    <property-name>customerType</property-name>
    <value>New</value>
</managed-property>
</managed-bean>
<managed-bean>
    <managed-bean-name>addressBean</managed-bean-name>
    <managed-bean-class>
        com.mycompany.mybeans.AddressBean
    </managed-bean-class>
    <managed-bean-scope> none </managed-bean-scope>
    <managed-property>
        <property-name>street</property-name>
        <null-value/>
    </managed-property>
    ...
</managed-bean>

```

The first `CustomerBean` declaration (with the `managed-bean-name` of `customer`) creates a `CustomerBean` in request scope. This bean has two properties: `mailingAddress` and `streetAddress`. These properties use the `value` element to reference a bean named `addressBean`.

The second `managed-bean` declaration defines an `AddressBean` but does not create it because its `managed-bean-scope` element defines a scope of `none`. Recall that a scope of `none` means that the bean is created only when something else references it. Because both the `mailingAddress` and the `streetAddress` properties reference `addressBean` using the `value` element, two instances of `AddressBean` are created when `CustomerBean` is created.

When you create an object that points to other objects, do not try to point to an object with a shorter life span because it might be impossible to recover that scope's resources when it goes away. A session-scoped object, for example, cannot point to a request-scoped object. And objects with `none` scope have no effective life span managed by the framework, so they can point only to other `none` scoped objects. [Table 10–2](#) outlines all of the allowed connections.

**TABLE 10–2** Allowable Connections Between Scoped Objects

An Object of This Scope	May Point to an Object of This Scope
none	none
application	none, application
session	none, application, session

TABLE 10-2 Allowable Connections Between Scoped Objects (Continued)

An Object of This Scope	May Point to an Object of This Scope
request	none, application, session, request, view
view	none, application, session, view

Be sure not to allow cyclical references between objects. For example, neither of the `AddressBean` objects in the preceding example should point back to the `CustomerBean` object because `CustomerBean` already points to the `AddressBean` objects.

## Initializing Maps and Lists

In addition to configuring `Map` and `List` properties, you can also configure a `Map` and a `List` directly so that you can reference them from a tag rather than referencing a property that wraps a `Map` or a `List`.

## Using Annotations

Support for annotations is an important new feature introduced in JavaServer Faces 2.0.

The `@ManagedBean` annotation in a class automatically registers that class as a managed bean class with the server runtime. Such a registered managed bean does not need `<managed-bean>` configuration entries in the application configuration resource file.

An example of using `managed-bean` annotation on a class is as follows:

```
@ManagedBean
@SessionScoped
public class DukesBday{
    ...
}
```

The above code snippet shows a bean that is managed by the JavaServer Faces implementation and is available for the length of that session. You do not need to configure the `managed-bean` instance in the `faces-config.xml` file. In effect, it is an alternative to the application configuration file approach and reduces the task of configuring `managed-bean`.

You can also define the scope of the `managed-bean` within the class file as shown in the above example. This facility is available to beans declared in `request`, `session`, or `application` scope but not in `view` scope.

All classes will be scanned for annotations at startup unless the `<faces-config>` element in the `faces-config.xml` file, has the `<metadata-complete>` attribute set to `true`.

Annotations are also available for other artifacts such as components, converters, validators, and renderers to be used in place of application configuration resource file entries. They are discussed along with registration of custom listeners, custom validators, and custom converters in *Java EE 6 Tutorial, Volume II: Advanced Topics*.

## Registering Custom Error Messages as a Resource Bundle

If you create custom error messages (which are displayed by the `message` and `messages` tags) for your custom converters or validators, you must make them available at application startup time. You do this in one of two ways: by queuing the message onto the `FacesContext` instance programmatically, or by registering the messages with your application using the application configuration resource file.

Here is an example of the file that registers the messages for an application:

```
<application>
  <message-bundle>
    com.sun.bookstore6.resources.ApplicationMessages
  </message-bundle>
  <locale-config>
    <default-locale>en</default-locale>
    <supported-locale>es</supported-locale>
    <supported-locale>de</supported-locale>
    <supported-locale>fr</supported-locale>
  </locale-config>
</application>
```

This set of elements will cause your `Application` instance to be populated with the messages that are contained in the specified resource bundle.

The `message-bundle` element represents a set of localized messages. It must contain the fully qualified path to the resource bundle containing the localized messages (in this case, `resources.ApplicationMessages`).

The `locale-config` element lists the default locale and the other supported locales. The `locale-config` element enables the system to find the correct locale based on the browser's language settings.

The `supported-locale` and `default-locale` tags accept the lowercase, two-character codes as defined by ISO 639 (see <http://www.ics.uci.edu/pub/ietf/http/related/iso639.txt>). Make sure that your resource bundle actually contains the messages for the locales that you specify with these tags.

To access the localized message, the application developer merely references the key of the message from the resource bundle.

## Registering Custom Localized Static Text

Any custom localized static text that you create that is not loaded into the page using the `loadBundle` tag must be registered with the application using the `resource-bundle` element in the application configuration resource file for your pages to have access to the text. Likewise, any custom error messages that are referenced by the `converterMessage`, `requiredMessage`, or `validatorMessage` attributes of an input component tag must also be made available to the application by way of the `loadBundle` tag or the `resource-bundle` element of the application configuration file.

Here is the part of the file that registers some custom error messages for the `guessNumber` application:

```
<application>
  ...
  <resource-bundle>
    <base-name>
      guessNumber.ApplicationMessages
    </base-name>
    <var>customMessages</var>
  </resource-bundle>
  ...
</application>
```

Similarly to the `loadBundle` tag, the value of the `base-name` subelement specifies the fully-qualified class name of the `ResourceBundle` class, which in this case is located in the `resources` package of the application.

Also similarly to the `var` attribute of the `loadBundle` tag, the `var` subelement of the `resource-bundle` element is an alias to the `ResourceBundle` class. This alias is used by tags in the page to identify the resource bundle.

The `locale-config` element shown in the previous section also applies to the messages and static text identified by the `resource-bundle` element. As with resource bundles identified by the `message-bundle` element, make sure that the resource bundle identified by the `resource-bundle` element actually contains the messages for the locales that you specify with these `locale-config` elements.

## Default Validator

In addition to the validators you declare on the UI components, you can also specify zero or more default validators in the application configuration resource files. The default validator applies to all `UIInput` instances in a view or component tree and is appended after the local defined validators. Here is an example of a default validator registered in the application configuration resource file:

```
<faces-config>
<application>
<default-validators>
<validator-id>javax.faces.Bean</validator-id>
</default-validators>
<application/>
</faces-config>
```

## Configuring Navigation Rules

Navigation is a set of rules for choosing the next page to be displayed after a button or hyperlink component is clicked. Navigation rules are defined in the application configuration resource file.

Each navigation rule specifies how to navigate from one page to a set of other pages. The JavaServer Faces implementation chooses the proper navigation rule according to which page is currently displayed.

After the proper navigation rule is selected, the choice of which page to access next from the current page depends on two factors:

- The action method that was invoked when the component was clicked
- The logical outcome that is referenced by the component's tag or was returned from the action method

The outcome can be anything that the developer chooses, but [Table 10–3](#) lists some outcomes commonly used in web applications.

**TABLE 10–3** Common Outcome Strings

Outcome	What It Means
success	Everything worked. Go on to the next page.
failure	Something is wrong. Go on to an error page.
logon	The user needs to log on first. Go on to the logon page.
no results	The search did not find anything. Go to the search page again.

Usually, the action method performs some processing on the form data of the current page. For example, the method might check whether the user name and password entered in the form match the user name and password on file. If they match, the method returns the outcome `success`. Otherwise, it returns the outcome `failure`. As this example demonstrates, both the method used to process the action and the outcome returned are necessary to determine the proper page to access.

Here is a navigation rule that could be used with the example just described:

```
<navigation-rule>
  <from-view-id>/logon.jsp</from-view-id>
  <navigation-case>
    <from-action>#{LogonForm.logon}</from-action>
    <from-outcome>success</from-outcome>
    <to-view-id>/storefront.jsp</to-view-id>
  </navigation-case>
  <navigation-case>
    <from-action>#{LogonForm.logon}</from-action>
    <from-outcome>failure</from-outcome>
    <to-view-id>/logon.jsp</to-view-id>
  </navigation-case>
</navigation-rule>
```

This navigation rule defines the possible ways to navigate from `logon.jsp`. Each `navigation-case` element defines one possible navigation path from `logon.jsp`. The first `navigation-case` says that if `LogonForm.logon` returns an outcome of `success`, then `storefront.jsp` will be accessed. The second `navigation-case` says that `logon.jsp` will be re-rendered if `LogonForm.logon` returns `failure`.

The configuration of an application's page flow consists of a set of navigation rules. Each rule is defined by the `navigation-rule` element in the `faces-config.xml` file.

Each `navigation-rule` element corresponds to one component tree identifier defined by the optional `from-view-id` element. This means that each rule defines all the possible ways to navigate from one particular page in the application. If there is no `from-view-id` element, the navigation rules defined in the `navigation-rule` element apply to all the pages in the application. The `from-view-id` element also allows wildcard matching patterns. For example, this `from-view-id` element says that the navigation rule applies to all the pages in the `books` directory:

```
<from-view-id>/books/*</from-view-id>
```

As shown in the example navigation rule, a `navigation-rule` element can contain zero or more `navigation-case` elements. The `navigation-case` element defines a set of matching criteria. When these criteria are satisfied, the application will navigate to the page defined by the `to-view-id` element contained in the same `navigation-case` element.

The navigation criteria are defined by optional `from-outcome` and `from-action` elements. The `from-outcome` element defines a logical outcome, such as `success`. The `from-action` element uses a method expression to refer to an action method that returns a `String`, which is the logical outcome. The method performs some logic to determine the outcome and returns the outcome.

The navigation-case elements are checked against the outcome and the method expression in this order:

- Cases specifying both a from-outcome value and a from-action value. Both of these elements can be used if the action method returns different outcomes depending on the result of the processing it performs.
- Cases specifying only a from-outcome value. The from-outcome element must match either the outcome defined by the action attribute of the UICommand component or the outcome returned by the method referred to by the UICommand component.
- Cases specifying only a from-action value. This value must match the action expression specified by the component tag.

When any of these cases is matched, the component tree defined by the to-view-id element will be selected for rendering.

Using NetBeans IDE, you can configure a navigation rule by doing the following:

1. After opening your project in NetBeans IDE, expand the project node in the Projects pane.
2. Expand the Web Pages and WEB-INF nodes of the project node.
3. Double-click `faces-config.xml`.
4. After `faces-config.xml` opens in the editor pane, right-click in the editor pane.
5. Select JavaServer Faces→Add Navigation Rule.
6. In the Add Navigation Rule dialog:
  - a. Enter or browse for the page that represents the starting view for this navigation rule.
  - b. Click Add.
7. Right-click again in the editor pane.
8. Select JavaServer Faces → Add Navigation Case.
9. In the Add Navigation Case dialog box:
  - a. From the From View menu, select the page that represents the starting view for the navigation rule (from Step 6 a).
  - b. (optional) In the From Action field, enter the action method invoked when the component that triggered navigation is activated.
  - c. (optional) In the From Outcome field, enter the logical outcome string that the activated component references from its action attribute.
  - d. From the To View menu, select or browse for the page that will be opened if this navigation case is selected by the navigation system.
  - e. Click Add.

[“Referencing a Method That Performs Navigation” on page 169](#) explains how to use a component tag’s action attribute to point to an action method. [“Writing a Method to Handle Navigation” on page 183](#) explains how to write an action method.

## Implicit Navigation Rules

Starting from JavaServer Faces 2.0, implicit navigation rules are available for Facelets applications. Implicit navigation rules come into picture if no `navigation-rules` are configured in the application resource configuration files.

When you add a UI component such as a `commandButton`, and assign a page as the value for its `action` property, the default navigation handler will try to match a suitable page within the application.

```
<h:commandButton value="submit" action="response">
```

In the above example, the default navigation handler will try to locate `response.xhtml` page and navigate to it.

## Basic Requirements of a JavaServer Faces Application

In addition to configuring your application, you must satisfy other requirements of JavaServer Faces applications, including properly packaging all the necessary files and providing a deployment descriptor. This section describes how to perform these administrative tasks.

JavaServer Faces 2.0 applications must be compliant with at least version 2.5 of the Servlet specification, and at least version 2.1 of the JavaServer Pages specification. All applications compliant with these specifications are packaged in a WAR file, which must conform to specific requirements to execute across different containers. At a minimum, a WAR file for a JavaServer Faces application must contain the following:

- A web application deployment descriptor, called `web.xml`, to configure resources required by a web application
- An application configuration resource file, which configures application resources
- A specific set of JAR files containing essential classes
- A set of application classes, JavaServer Faces pages, and other required resources, such as image files

For example, a Java Server Faces web application WAR file using Facelets typically has the following directory structure:

```
$PROJECT_DIR [Web Pages] +-[xhtml documents] +-/resources +-/WEB-INF +-/classes  
+-/lib +-/web.xml +-/faces-config.xml +-/sun-web.xml
```

The `web.xml` file (or deployment descriptor), the set of JAR files, and the set of application files must be contained in the `WEB-INF` directory of the WAR file.

## Configuring an Application With a Deployment Descriptor

Web applications are configured using elements that are contained in the web application deployment descriptor. The deployment descriptor for a JavaServer Faces application must specify certain configurations, which include the following:

- The servlet used to process JavaServer Faces requests
- The servlet mapping for the processing servlet
- The path to the configuration resource file if it is not located in a default location

The deployment descriptor can also specify other, optional configurations, including:

- Specifying where component state is saved
- Encrypting state saved on the client
- Compressing state saved on the client
- Restricting access to pages containing JavaServer Faces tags
- Turning on XML validation
- Verifying custom objects

This section gives more details on these configurations. Where appropriate, it also describes how you can make these configurations using NetBeans IDE.

### Identifying the Servlet for Life Cycle Processing

A requirement of a JavaServer Faces application is that all requests to the application that reference previously saved JavaServer Faces components must go through `FacesServlet`. A `FacesServlet` instance manages the request processing life cycle for web applications and initializes the resources required by JavaServer Faces technology.

Before a JavaServer Faces application can launch the first web page, the web container must invoke the `FacesServlet` instance in order for the application life cycle process to start.

Here is an example of defining the `FacesServlet`:

```
<servlet>
    <servlet-name>FacesServlet</servlet-name>
    <servlet-class>javax.faces.webapp.FacesServlet</servlet-class>
</servlet>
```

To make sure that the `FacesServlet` instance is invoked, you provide a mapping to it. The mapping to `FacesServlet` can be a prefix mapping, such as `/guess/*`, or an extension mapping, such as `*.faces`. The mapping is used to identify a JSP page as having JavaServer Faces content. Because of this, the URL to the first JSP page of the application must include the mapping.

In the case of prefix mapping, there are two ways to accomplish this:

- The page author can include a simple HTML page, such as an `index.html` file in the application that provides the URL to the first page. This URL must include the path to `FacesServlet`, as shown by this tag, which uses the mapping defined in the `guessNumber` application:

```
<a href="guess/greeting.xhtml">
```

- Users of the application can include the path to `FacesServlet` in the URL to the first page when they enter it in their browser, as shown in the example below :

```
http://localhost:8080/guessNumber/guess/greeting.xhtml
```

The second method allows users to start the application from the first page, rather than start it from an HTML page. However, the second method requires users to identify the first page. When you use the first method, users need only enter the following:

```
http://localhost:8080/guessNumber
```

In the case of extension mapping, if a request comes to the server for a page with a `.faces` extension, the container will send the request to the `FacesServlet` instance, which will expect a corresponding page of the same name to exist containing the content.

If you are using NetBeans IDE, the time to map the `FacesServlet` instance is when you create your JavaServer Faces project with NetBeans IDE:

1. In NetBeans IDE, select File→New Project.
2. In the New Project dialog, select Web from the Categories tree.
3. Select Web Application from the Projects panel.
4. Click Next.
5. Fill out the information in the Name and Location screen of the wizard.
6. Click Next.
7. Select the JavaServer Faces check box in the Frameworks screen.
8. Enter the mapping, such as `*.faces`, to the `FacesServlet` instance in the Servlet URL Mapping field.
9. Click Finish.

After your project is open in NetBeans IDE, you can change the mapping to the `FacesServlet` instance by doing the following:

1. Expand the node of your project in the Projects pane.
2. Expand the Web Pages and WEB-INF nodes that are under the project node.
3. Double-click `web.xml`.
4. After the `web.xml` file appears in the editor pane, click Servlets at the top of the editor pane. The `FacesServlet` configuration appears in the editor pane.

If you prefer to edit the `web.xml` file directly, perform the following steps to configure a mapping to the `FacesServlet` instance:

1. Include a `servlet` element in the deployment descriptor.
2. Inside the `servlet` element, include a `display-name` element and set it to `FacesServlet`.
3. Also inside the `servlet` element, add a `servlet-name` element and set it to `FacesServlet`.
4. Add a third element, called `servlet-class`, inside the `servlet` element and set it to `javax.faces.webapp.FacesServlet`. This is the fully-qualified class name of the `FacesServlet` class.
5. After the `servlet` element, add a `servlet-mapping` element.
6. Inside the `servlet-mapping` element, add a `servlet-name` element and set it to `FacesServlet`. This must match the name identified by the `servlet-name` element described in step 3.
7. Also inside the `servlet-mapping` element, add a `url-pattern` element and set it to whatever mapping you prefer. This will be the path to `FacesServlet`. Users of the application will include this path in the URL when they access the application. For the `guessNumber` application, the path is `/guess/*`.

## Specifying a Path to an Application Configuration Resource File

As explained in “[Application Configuration Resource File](#)” on page 189, an application can have multiple application configuration resource files. If these files are not located in the directories that the implementation searches by default or the files are not named `faces-config.xml`, you need to specify paths to these files.

To specify these paths using NetBeans IDE, do the following:

1. Expand the node of your project in the Projects pane.
2. Expand the Web Pages and WEB-INF nodes that are under the project node.
3. Double-click `web.xml`.
4. After the `web.xml` file appears in the editor pane, click General at the top of the editor pane.
5. Expand the Context Parameters node.
6. Click Add.
7. In the Add Context Parameter dialog:
  - a. Enter `javax.faces.CONFIG_FILES` in the Param Name field.
  - b. Enter the path to your configuration file in the Param Value field.
  - c. Click OK.
8. Repeat steps 1 through 7 for each configuration file.

To specify paths to the files by editing the deployment descriptor directly follow these steps:

1. Add a `context-param` element to the deployment descriptor.
2. Add a `param-value` element inside the `context-param` element and call it `javax.faces.CONFIG_FILES`.
3. Add a `param-value` element inside the `context-param` element and give it the path to your configuration file. For example, the path to the `guessNumber` application's application configuration resource file is `/WEB-INF/faces-config.xml`.
4. Repeat steps 2 and 3 for each application configuration resource file that your application contains.

## Specifying Where State Is Saved

When implementing the state-holder methods, you specify in your deployment descriptor where you want the state to be saved, either client or server. You do this by setting a context parameter in your deployment descriptor.

To specify where state is saved using NetBeans IDE, do the following:

1. Expand the node of your project in the Projects pane.
2. Expand the Web Pages and WEB-INF nodes that are under the project node.
3. Double-click `web.xml`.
4. After the `web.xml` file appears in the editor pane, click General at the top of the editor pane.
5. Expand the Context Parameters node.
6. Click Add.
7. In the Add Context Parameter dialog:
  - a. Enter `javax.faces.STATE_SAVING_METHOD` in the Param Name field.
  - b. Enter `client` or `server` in the Param Value field.
  - c. Click OK.

To specify where state is saved by editing the deployment descriptor directly, follow these steps:

1. Add a `context-param` element to the deployment descriptor.
2. Add a `param-name` element inside the `context-param` element and give it the name `javax.faces.STATE_SAVING_METHOD`.
3. Add a `param-value` element to the `context-param` element and give it the value `client` or `server`, depending on whether you want state saved in the client or the server.

If state is saved on the client, the state of the entire view is rendered to a hidden field on the page. The JavaServer Faces implementation saves the state on the client by default. Duke's Bookstore saves its state in the client.

## Including the Classes, Pages, and Other Resources

When packaging web applications using the included build scripts, you'll notice that the scripts package resources in the following ways:

- All pages are placed at the top level of the WAR file.
- The `faces-config.xml` file and the `web.xml` file are packaged in the `WEB-INF` directory.
- All packages are stored in the `WEB-INF/classes/` directory.
- All application JAR files are packaged in the `WEB-INF/lib/` directory.
- All resource files are either under the root of the web application `/resources` directory, or in the web application's classpath, `META-INF/resources/<resourceIdentifier>` directory. For more information on resources, see [“Resources” on page 112](#).

When packaging your own applications, you can use NetBeans IDE or you can use the build scripts. You can modify the build scripts to fit your situation. However, you can continue to package your WAR files as described in this section because this technique complies with the commonly accepted practice for packaging web applications.

# Java Servlet Technology

---

Shortly after the web began to be used for delivering services, service providers recognized the need for dynamic content. Applets, one of the earliest attempts toward this goal, focused on using the client platform to deliver dynamic user experiences. At the same time, developers also investigated using the server platform for the same purpose. Initially, Common Gateway Interface (CGI) server-side scripts were the main technology used to generate dynamic content. Although widely used, CGI scripting technology had many shortcomings, including platform dependence and lack of scalability. To address these limitations, Java Servlet technology was created as a portable way to provide dynamic, user-oriented content.

## What Is a Servlet?

A *servlet* is a Java programming language class that is used to extend the capabilities of servers that host applications accessed by means of a request-response programming model. Although servlets can respond to any type of request, they are commonly used to extend the applications hosted by web servers. For such applications, Java Servlet technology defines HTTP-specific servlet classes.

The `javax.servlet` and `javax.servlet.http` packages provide interfaces and classes for writing servlets. All servlets must implement the `Servlet` interface, which defines life-cycle methods. When implementing a generic service, you can use or extend the `GenericServlet` class provided with the Java Servlet API. The `HttpServlet` class provides methods, such as `doGet` and `doPost`, for handling HTTP-specific services.

## Servlet Life Cycle

The life cycle of a servlet is controlled by the container in which the servlet has been deployed. When a request is mapped to a servlet, the container performs the following steps.

1. If an instance of the servlet does not exist, the web container
  - a. Loads the servlet class.
  - b. Creates an instance of the servlet class.
  - c. Initializes the servlet instance by calling the `init` method. Initialization is covered in [“Initializing a Servlet” on page 217](#).
2. Invokes the `service` method, passing request and response objects. Service methods are discussed in [“Writing Service Methods” on page 217](#).

If the container needs to remove the servlet, it finalizes the servlet by calling the servlet’s `destroy` method. Finalization is discussed in [“Finalizing a Servlet” on page 228](#).

## Handling Servlet Life-Cycle Events

You can monitor and react to events in a servlet’s life cycle by defining listener objects whose methods get invoked when life-cycle events occur. To use these listener objects you must define and specify the listener class.

### Defining the Listener Class

You define a listener class as an implementation of a listener interface. [Table 11-1](#) lists the events that can be monitored and the corresponding interface that must be implemented. When a listener method is invoked, it is passed an event that contains information appropriate to the event. For example, the methods in the `HttpSessionListener` interface are passed an `HttpSessionEvent`, which contains an `HttpSession`.

TABLE 11-1 Servlet Life-Cycle Events

Object	Event	Listener Interface and Event Class
Web context (see <a href="#">“Accessing the Web Context” on page 226</a> )	Initialization and destruction	<code>javax.servlet.ServletContextListener</code> and <code>ServletContextEvent</code>
	Attribute added, removed, or replaced	<code>javax.servlet.ServletContextAttributeListener</code> and <code>ServletContextAttributeEvent</code>

TABLE 11-1 Servlet Life-Cycle Events (Continued)

Object	Event	Listener Interface and Event Class
Session (See “Maintaining Client State” on page 226)	Creation, invalidation, activation, passivation, and timeout	<code>javax.servlet.http.HttpSessionListener</code> , <code>javax.servlet.http.HttpSessionActivationListener</code> , and <code>HttpSessionEvent</code>
	Attribute added, removed, or replaced	<code>javax.servlet.http.HttpSessionAttributeListener</code> and <code>HttpSessionBindingEvent</code>
Request	A servlet request has started being processed by web components	<code>javax.servlet.ServletRequestListener</code> and <code>ServletRequestEvent</code>
	Attribute added, removed, or replaced	<code>javax.servlet.ServletRequestAttributeListener</code> and <code>ServletRequestAttributeEvent</code>

## Specifying Event Listener Classes

You specify an event listener class using the `listener` element of the deployment descriptor.

You can specify an event listener using the deployment descriptor editor of NetBeans IDE by doing the following:

1. Expand your application’s project node.
2. Expand the project’s Web Pages and WEB-INF nodes.
3. Double-click `web.xml`.
4. Click General at the top of the `web.xml` editor.
5. Expand the Web Application Listeners node.
6. Click Add.
7. In the Add Listener dialog, click Browse to locate the listener class.
8. Click OK.

## Handling Servlet Errors

Any number of exceptions can occur when a servlet executes. When an exception occurs, the web container generates a default page containing the message

A Servlet Exception Has Occurred

But you can also specify that the container should return a specific error page for a given exception.

## Sharing Information

Web components, like most objects, usually work with other objects to accomplish their tasks. There are several ways they can do this. They can use private helper objects (for example, JavaBeans components), they can share objects that are attributes of a public scope, they can use a database, and they can invoke other web resources. The Java Servlet technology mechanisms that allow a web component to invoke other web resources are described in [“Invoking Other Web Resources” on page 224](#).

## Using Scope Objects

Collaborating web components share information by means of objects that are maintained as attributes of four scope objects. You access these attributes using the `[get|set]Attribute` methods of the class representing the scope. [Table 11–2](#) lists the scope objects.

TABLE 11–2 Scope Objects

Scope Object	Class	Accessible From
Web context	<code>javax.servlet.ServletContext</code>	Web components within a web context. See <a href="#">“Accessing the Web Context” on page 226</a> .
Session	<code>javax.servlet.http.HttpSession</code>	Web components handling a request that belongs to the session. See <a href="#">“Maintaining Client State” on page 226</a> .
Request	subtype of <code>javax.servlet.ServletException</code>	Web components handling the request.
Page	<code>javax.servlet.jsp.JspContext</code>	The JSP page that creates the object.

## Controlling Concurrent Access to Shared Resources

In a multithreaded server, it is possible for shared resources to be accessed concurrently. In addition to scope object attributes, shared resources include in-memory data (such as instance or class variables) and external objects such as files, database connections, and network connections.

Concurrent access can arise in several situations:

- Multiple web components accessing objects stored in the web context.
- Multiple web components accessing objects stored in a session.
- Multiple threads within a web component accessing instance variables. A web container will typically create a thread to handle each request. To ensure that a servlet instance handles only one request at a time, a servlet can implement the `SingleThreadModel` interface. If a servlet implements this interface, no two threads will execute concurrently in the servlet’s

service method. A web container can implement this guarantee by synchronizing access to a single instance of the servlet or by maintaining a pool of web component instances and dispatching each new request to a free instance. This interface does not prevent synchronization problems that result from web components accessing shared resources such as static class variables or external objects.

When resources can be accessed concurrently, they can be used in an inconsistent fashion. You prevent this by controlling the access using the synchronization techniques described in the [Threads](#) lesson in *The Java Tutorial, Fourth Edition*, by Sharon Zakhour et al. (Addison-Wesley, 2006).

## Initializing a Servlet

After the web container loads and instantiates the servlet class and before it delivers requests from clients, the web container initializes the servlet. To customize this process to allow the servlet to read persistent configuration data, initialize resources, and perform any other one-time activities, you override the `init` method of the `Servlet` interface. If a servlet cannot complete its initialization process, it throws an `UnavailableException`.

## Writing Service Methods

The service provided by a servlet is implemented in the `service` method of a `GenericServlet`, in the `doMethod` methods (where *Method* can take the value `Get`, `Delete`, `Options`, `Post`, `Put`, or `Trace`) of an `HttpServlet` object, or in any other protocol-specific methods defined by a class that implements the `Servlet` interface. The term *service method* is used for any method in a servlet class that provides a service to a client.

The general pattern for a service method is to extract information from the request, access external resources, and then populate the response based on that information.

For HTTP servlets, the correct procedure for populating the response is to first retrieve an output stream from the response, then fill in the response headers, and finally write any body content to the output stream. Response headers must always be set before the response has been committed. Any attempt to set or add headers after the response has been committed will be ignored by the web container. The next two sections describe how to get information from requests and generate responses.

## Getting Information from Requests

A request contains data passed between a client and the servlet. All requests implement the `ServletRequest` interface. This interface defines methods for accessing the following information:

- Parameters, which are typically used to convey information between clients and servlets
- Object-valued attributes, which are typically used to pass information between the servlet container and a servlet or between collaborating servlets
- Information about the protocol used to communicate the request and about the client and server involved in the request
- Information relevant to localization

You can also retrieve an input stream from the request and manually parse the data. To read character data, use the `BufferedReader` object returned by the request's `getReader` method. To read binary data, use the `ServletInputStream` returned by `getInputStream`.

HTTP servlets are passed an HTTP request object, `HttpServletRequest`, which contains the request URL, HTTP headers, query string, and so on.

An HTTP request URL contains the following parts:

```
http://[host]:[port][request-path]?[query-string]
```

The request path is further composed of the following elements:

- **Context path:** A concatenation of a forward slash (/) with the context root of the servlet's web application.
- **Servlet path:** The path section that corresponds to the component alias that activated this request. This path starts with a forward slash (/).
- **Path info:** The part of the request path that is not part of the context path or the servlet path.

If the context path is `/catalog` and for the aliases listed in [Table 11-3](#), [Table 11-4](#) gives some examples of how the URL will be parsed.

TABLE 11-3 Aliases

Pattern	Servlet
<code>/Lawn/*</code>	<code>LawnServlet</code>
<code>/*.jsp</code>	<code>JSPServlet</code>

TABLE 11-4 Request Path Elements

Request Path	Servlet Path	Path Info
/catalog/lawn/index.html	/lawn	/index.html
/catalog/help/feedback.jsp	/help/feedback.jsp	null

Query strings are composed of a set of parameters and values. Individual parameters are retrieved from a request by using the `getParameter` method. There are two ways to generate query strings:

- A query string can explicitly appear in a web page.
- A query string is appended to a URL when a form with a GET HTTP method is submitted.

## Constructing Responses

A response contains data passed between a server and the client. All responses implement the `ServletResponse` interface. This interface defines methods that allow you to:

- Retrieve an output stream to use to send data to the client. To send character data, use the `PrintWriter` returned by the response's `getWriter` method. To send binary data in a MIME body response, use the `ServletOutputStream` returned by `getOutputStream`. To mix binary and text data (as in a multipart response), use a `ServletOutputStream` and manage the character sections manually.
- Indicate the content type (for example, `text/html`) being returned by the response with the `setContentType(String)` method. This method must be called before the response is committed. A registry of content type names is kept by the Internet Assigned Numbers Authority (IANA) at <http://www.iana.org/assignments/media-types/>.
- Indicate whether to buffer output with the `setBufferSize(int)` method. By default, any content written to the output stream is immediately sent to the client. Buffering allows content to be written before anything is actually sent back to the client, thus providing the servlet with more time to set appropriate status codes and headers or forward to another web resource. The method must be called before any content is written or before the response is committed.
- Set localization information such as locale and character encoding.

HTTP response objects, `javax.servlet.http.HttpServletResponse`, have fields representing HTTP headers such as the following:

- Status codes, which are used to indicate the reason a request is not satisfied or that a request has been redirected.
- Cookies, which are used to store application-specific information at the client. Sometimes cookies are used to maintain an identifier for tracking a user's session (see [“Session Tracking” on page 227](#)).

## Filtering Requests and Responses

A *filter* is an object that can transform the header and content (or both) of a request or response. Filters differ from web components in that filters usually do not themselves create a response. Instead, a filter provides functionality that can be “attached” to any kind of web resource. Consequently, a filter should not have any dependencies on a web resource for which it is acting as a filter; this way it can be composed with more than one type of web resource.

The main tasks that a filter can perform are as follows:

- Query the request and act accordingly.
- Block the request-and-response pair from passing any further.
- Modify the request headers and data. You do this by providing a customized version of the request.
- Modify the response headers and data. You do this by providing a customized version of the response.
- Interact with external resources.

Applications of filters include authentication, logging, image conversion, data compression, encryption, tokenizing streams, XML transformations, and so on.

You can configure a web resource to be filtered by a chain of zero, one, or more filters in a specific order. This chain is specified when the web application containing the component is deployed and is instantiated when a web container loads the component.

In summary, the tasks involved in using filters are

- Programming the filter
- Programming customized requests and responses
- Specifying the filter chain for each web resource

## Programming Filters

The filtering API is defined by the `Filter`, `FilterChain`, and `FilterConfig` interfaces in the `javax.servlet` package. You define a filter by implementing the `Filter` interface.

The most important method in this interface is `doFilter`, which is passed request, response, and filter chain objects. This method can perform the following actions:

- Examine the request headers.
- Customize the request object if the filter wishes to modify request headers or data.
- Customize the response object if the filter wishes to modify response headers or data.
- Invoke the next entity in the filter chain. If the current filter is the last filter in the chain that ends with the target web component or static resource, the next entity is the resource at the end of the chain; otherwise, it is the next filter that was configured in the WAR. The filter invokes the next entity by calling the `doFilter` method on the chain object (passing in the request and response it was called with, or the wrapped versions it may have created). Alternatively, it can choose to block the request by not making the call to invoke the next entity. In the latter case, the filter is responsible for filling out the response.
- Examine response headers after it has invoked the next filter in the chain.
- Throw an exception to indicate an error in processing.

In addition to `doFilter`, you must implement the `init` and `destroy` methods. The `init` method is called by the container when the filter is instantiated. If you wish to pass initialization parameters to the filter, you retrieve them from the `FilterConfig` object passed to `init`.

## Programming Customized Requests and Responses

There are many ways for a filter to modify a request or response. For example, a filter can add an attribute to the request or can insert data in the response.

A filter that modifies a response must usually capture the response before it is returned to the client. To do this, you pass a stand-in stream to the servlet that generates the response. The stand-in stream prevents the servlet from closing the original response stream when it completes and allows the filter to modify the servlet's response.

To pass this stand-in stream to the servlet, the filter creates a response wrapper that overrides the `getWriter` or `getOutputStream` method to return this stand-in stream. The wrapper is passed to the `doFilter` method of the filter chain. Wrapper methods default to calling through to the wrapped request or response object.

To override request methods, you wrap the request in an object that extends `ServletRequestWrapper` or `HttpServletRequestWrapper`. To override response methods, you wrap the response in an object that extends `ServletResponseWrapper` or `HttpServletResponseWrapper`.

## Specifying Filter Mappings

A web container uses filter mappings to decide how to apply filters to web resources. A filter mapping matches a filter to a web component by name, or to web resources by URL pattern. The filters are invoked in the order in which filter mappings appear in the filter mapping list of a WAR. You specify a filter mapping list for a WAR in its deployment descriptor, either with NetBeans IDE or by coding the list by hand with XML.

To declare the filter and map it to a web resource using NetBeans IDE, do the following:

1. Expand the application's project node in the Project pane.
2. Expand the Web Pages and WEB-INF nodes under the project node.
3. Double-click `web.xml`.
4. Click Filters at the top of the editor pane.
5. Expand the Servlet Filters node in the editor pane.
6. Click Add Filter Element to map the filter to a web resource by name or by URL pattern.
7. In the Add Servlet Filter dialog, enter the name of the filter in the Filter Name field.
8. Click Browse to locate the servlet class to which the filter applies. You can include wildcard characters so that you can apply the filter to more than one servlet.
9. Click OK.

To constrain how the filter is applied to requests, do the following:

1. Expand the Filter Mappings node in the Filters tab of the editor pane.
2. Select the filter from the list of filters.
3. Click Add.
4. In the Add Filter Mapping dialog, select one of the following dispatcher types:
  - **REQUEST**: Only when the request comes directly from the client
  - **FORWARD**: Only when the request has been forwarded to a component (see [“Transferring Control to Another Web Component”](#) on page 225)
  - **INCLUDE**: Only when the request is being processed by a component that has been included (see [“Including Other Resources in the Response”](#) on page 225)
  - **ERROR**: Only when the request is being processed with the error page mechanism (see [“Handling Servlet Errors”](#) on page 215)

You can direct the filter to be applied to any combination of the preceding situations by selecting multiple dispatcher types. If no types are specified, the default option is REQUEST.

You can declare, map, and constrain the filter by editing the XML in the web application deployment descriptor directly by following these steps:

1. While in the `web.xml` editor pane in NetBeans IDE, click XML at the top of the editor pane.
2. Declare the filter by adding a `filter` element right after the `display-name` element. The `filter` element creates a name for the filter and declares the filter's implementation class and initialization parameters.
3. Map the filter to a web resource by name or by URL pattern using the `filter-mapping` element:
  - a. Include a `filter-name` element that specifies the name of the filter as defined by the `filter` element.
  - b. Include a `servlet-name` element that specifies to which servlet the filter applies. The `servlet-name` element can include wildcard characters so that you can apply the filter to more than one servlet.
4. Constrain how the filter will be applied to requests by specifying one of the enumerated dispatcher options (described in step 4 of the preceding set of steps) with the `dispatcher` element and adding the `dispatcher` element to the `filter-mapping` element.

You can direct the filter to be applied to any combination of the preceding situations by including multiple `dispatcher` elements. If no elements are specified, the default option is `REQUEST`.

If you want to log every request to a web application, you map the hit counter filter to the URL pattern `/*`.

You can map a filter to one or more web resources and you can map more than one filter to a web resource. This is illustrated in [Figure 11-1](#), where filter F1 is mapped to servlets S1, S2, and S3, filter F2 is mapped to servlet S2, and filter F3 is mapped to servlets S1 and S2.

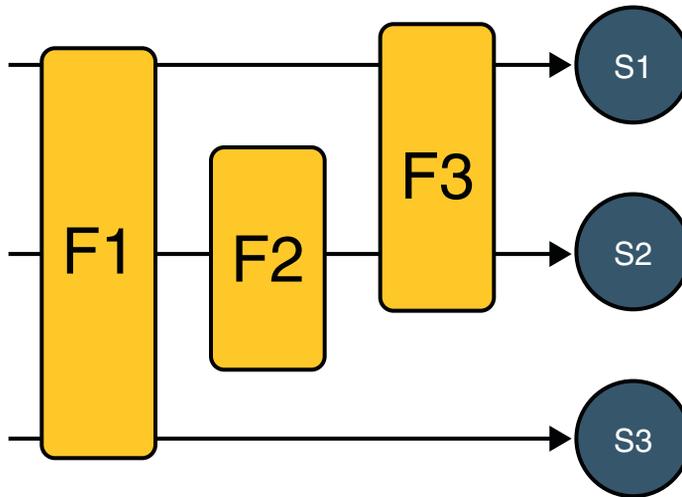


FIGURE 11-1 Filter-to-Servlet Mapping

Recall that a filter chain is one of the objects passed to the `doFilter` method of a filter. This chain is formed indirectly by means of filter mappings. The order of the filters in the chain is the same as the order in which filter mappings appear in the web application deployment descriptor.

When a filter is mapped to servlet `S1`, the web container invokes the `doFilter` method of `F1`. The `doFilter` method of each filter in `S1`'s filter chain is invoked by the preceding filter in the chain by means of the `chain.doFilter` method. Because `S1`'s filter chain contains filters `F1` and `F3`, `F1`'s call to `chain.doFilter` invokes the `doFilter` method of filter `F3`. When `F3`'s `doFilter` method completes, control returns to `F1`'s `doFilter` method.

## Invoking Other Web Resources

Web components can invoke other web resources in two ways: indirectly and directly. A web component indirectly invokes another web resource when it embeds a URL that points to another web component in content returned to a client.

A web component can also directly invoke another resource while it is executing. There are two possibilities: The web component can include the content of another resource, or it can forward a request to another resource.

To invoke a resource available on the server that is running a web component, you must first obtain a `RequestDispatcher` object using the `getRequestDispatcher("URL")` method.

You can get a `RequestDispatcher` object from either a request or the web context; however, the two methods have slightly different behavior. The method takes the path to the requested resource as an argument. A request can take a relative path (that is, one that does not begin with

a /), but the web context requires an absolute path. If the resource is not available or if the server has not implemented a `RequestDispatcher` object for that type of resource, `getRequestDispatcher` will return null. Your servlet should be prepared to deal with this condition.

## Including Other Resources in the Response

It is often useful to include another web resource (for example, banner content or copyright information) in the response returned from a web component. To include another resource, invoke the `include` method of a `RequestDispatcher` object:

```
include(request, response);
```

If the resource is static, the `include` method enables programmatic server-side includes. If the resource is a web component, the effect of the method is to send the request to the included web component, execute the web component, and then include the result of the execution in the response from the containing servlet. An included web component has access to the request object, but it is limited in what it can do with the response object:

- It can write to the body of the response and commit a response.
- It cannot set headers or call any method (for example, `setCookie`) that affects the headers of the response.

## Transferring Control to Another Web Component

In some applications, you might want to have one web component do preliminary processing of a request and have another component generate the response. For example, you might want to partially process a request and then transfer to another component depending on the nature of the request.

To transfer control to another web component, you invoke the `forward` method of a `RequestDispatcher`. When a request is forwarded, the request URL is set to the path of the forwarded page. The original URI and its constituent parts are saved as request attributes `javax.servlet.forward.[request-uri|context-path|servlet-path|path-info|query-string]`.

The `forward` method should be used to give another resource responsibility for replying to the user. If you have already accessed a `ServletOutputStream` or `PrintWriter` object within the servlet, you cannot use this method; doing so throws an `IllegalStateException`.

## Accessing the Web Context

The context in which web components execute is an object that implements the `ServletContext` interface. You retrieve the web context using the `getServletContext` method. The web context provides methods for accessing:

- Initialization parameters
- Resources associated with the web context
- Object-valued attributes
- Logging capabilities

The counter's access methods are synchronized to prevent incompatible operations by servlets that are running concurrently. A filter retrieves the counter object using the context's `getAttribute` method. The incremented value of the counter is recorded in the log.

## Maintaining Client State

Many applications require that a series of requests from a client be associated with one another. For example, a web application can save the state of a user's shopping cart across requests. Web-based applications are responsible for maintaining such state, called a *session*, because HTTP is stateless. To support applications that need to maintain state, Java Servlet technology provides an API for managing sessions and allows several mechanisms for implementing sessions.

### Accessing a Session

Sessions are represented by an `HttpSession` object. You access a session by calling the `getSession` method of a request object. This method returns the current session associated with this request, or, if the request does not have a session, it creates one.

### Associating Objects with a Session

You can associate object-valued attributes with a session by name. Such attributes are accessible by any web component that belongs to the same web context *and* is handling a request that is part of the same session.

## Notifying Objects That Are Associated with a Session

Recall that your application can notify web context and session listener objects of servlet life-cycle events (“[Handling Servlet Life-Cycle Events](#)” on page 214). You can also notify objects of certain events related to their association with a session such as the following:

- When the object is added to or removed from a session. To receive this notification, your object must implement the `javax.servlet.http.HttpSessionBindingListener` interface.
- When the session to which the object is attached will be passivated or activated. A session will be passivated or activated when it is moved between virtual machines or saved to and restored from persistent storage. To receive this notification, your object must implement the `javax.servlet.http.HttpSessionActivationListener` interface.

## Session Management

Because there is no way for an HTTP client to signal that it no longer needs a session, each session has an associated timeout so that its resources can be reclaimed. The timeout period can be accessed by using a session’s `[get|set]MaxInactiveInterval` methods.

You can also set the timeout period in the deployment descriptor using NetBeans IDE:

1. Open the `web.xml` file in the `web.xml` editor.
2. Click General at the top of the editor.
3. Enter an integer value in the Session Timeout field. The integer value represents the number of minutes of inactivity that must pass before the session times out.

To ensure that an active session is not timed out, you should periodically access the session by using service methods because this resets the session’s time-to-live counter.

When a particular client interaction is finished, you use the session’s `invalidate` method to invalidate a session on the server side and remove any session data.

## Session Tracking

A web container can use several methods to associate a session with a user, all of which involve passing an identifier between the client and the server. The identifier can be maintained on the client as a cookie, or the web component can include the identifier in every URL that is returned to the client.

If your application uses session objects, you must ensure that session tracking is enabled by having the application rewrite URLs whenever the client turns off cookies. You do this by calling the response’s `encodeURL (URL)` method on all URLs returned by a servlet. This method includes the session ID in the URL only if cookies are disabled; otherwise, it returns the URL unchanged.

## Finalizing a Servlet

When a servlet container determines that a servlet should be removed from service (for example, when a container wants to reclaim memory resources or when it is being shut down), the container calls the `destroy` method of the `Servlet` interface. In this method, you release any resources the servlet is using and save any persistent state. The `destroy` method releases the database object created in the `init` method.

All of a servlet's service methods should be complete when a servlet is removed. The server tries to ensure this by calling the `destroy` method only after all service requests have returned or after a server-specific grace period, whichever comes first. If your servlet has operations that take a long time to run (that is, operations that may run longer than the server's grace period), the operations could still be running when `destroy` is called. You must make sure that any threads still handling client requests complete; the remainder of this section describes how to do the following:

- Keep track of how many threads are currently running the service method.
- Provide a clean shutdown by having the `destroy` method notify long-running threads of the shutdown and wait for them to complete.
- Have the long-running methods poll periodically to check for shutdown and, if necessary, stop working, clean up, and return.

## Tracking Service Requests

To track service requests, include in your servlet class a field that counts the number of service methods that are running. The field should have synchronized access methods to increment, decrement, and return its value.

```
public class ShutdownExample extends HttpServlet {
    private int serviceCounter = 0;
    ...
    // Access methods for serviceCounter
    protected synchronized void enteringServiceMethod() {
        serviceCounter++;
    }
    protected synchronized void leavingServiceMethod() {
        serviceCounter--;
    }
    protected synchronized int numServices() {
        return serviceCounter;
    }
}
```

The service method should increment the service counter each time the method is entered and should decrement the counter each time the method returns. This is one of the few times that

your `HttpServlet` subclass should override the `service` method. The new method should call `super.service` to preserve the functionality of the original `service` method:

```
protected void service(HttpServletRequest req,
                       HttpServletResponse resp)
    throws ServletException, IOException {
    enteringServiceMethod();
    try {
        super.service(req, resp);
    } finally {
        leavingServiceMethod();
    }
}
```

## Notifying Methods to Shut Down

To ensure a clean shutdown, your `destroy` method should not release any shared resources until all the service requests have completed. One part of doing this is to check the service counter. Another part is to notify the long-running methods that it is time to shut down. For this notification, another field is required. The field should have the usual access methods:

```
public class ShutdownExample extends HttpServlet {
    private boolean shuttingDown;
    ...
    //Access methods for shuttingDown
    protected synchronized void setShuttingDown(boolean flag) {
        shuttingDown = flag;
    }
    protected synchronized boolean isShuttingDown() {
        return shuttingDown;
    }
}
```

Here is an example of the `destroy` method using these fields to provide a clean shutdown:

```
public void destroy() {
    /* Check to see whether there are still service methods */
    /* running, and if there are, tell them to stop. */
    if (numServices() > 0) {
        setShuttingDown(true);
    }

    /* Wait for the service methods to stop. */
    while(numServices() > 0) {
        try {
            Thread.sleep(interval);
        }
    }
}
```

```
        } catch (InterruptedException e) {  
        }  
    }  
}
```

## Creating Polite Long-Running Methods

The final step in providing a clean shutdown is to make any long-running methods behave politely. Methods that might run for a long time should check the value of the field that notifies them of shutdowns and should interrupt their work, if necessary.

```
public void doPost(...) {  
    ...  
    for(i = 0; ((i < lotsOfStuffToDo) &&  
        !isShuttingDown()); i++) {  
        try {  
            partOfLongRunningOperation(i);  
        } catch (InterruptedException e) {  
            ...  
        }  
    }  
}
```

## Further Information about Java Servlet Technology

For more information on Java Servlet technology, see:

- Java Servlet 3.0 specification:  
<http://jcp.org/en/jsr/detail?id=315>
- The Java Servlet web site:  
<http://java.sun.com/products/servlet>

( PART III

# Web Services

Part Three explores web services.



# Introduction to Web Services

---

This section of the tutorial discusses Java EE 6 web services technologies. For this book, these technologies include Java API for XML Web Services (JAX-WS) and Java API for RESTful Web Services (JAX-RS).

- “What Are Web Services?” on page 233
- “Types of Web Services” on page 233
- “Deciding Which Type of Web Service to Use” on page 236

## What Are Web Services?

*Web services* are client and server applications that communicate over the World Wide Web's (WWW) HyperText Transfer Protocol (HTTP) protocol.

As described by the World Wide Web Consortium (W3C), web services provide a standard means of interoperating between different software applications, running on a variety of platforms and frameworks. Web services are characterized by their great interoperability and extensibility, as well as their machine-processable descriptions thanks to the use of XML. They can be combined in a loosely coupled way to achieve complex operations. Programs providing simple services can interact with each other to deliver sophisticated added-value services.

## Types of Web Services

On the conceptual level, a service is a software component provided through a network-accessible endpoint. The service consumer and provider use messages to exchange invocation request and response information in the form of self-containing documents that make very few assumptions about the technological capabilities of the receiver.

On a technical level, web services can be implemented in different ways. The two types of web services discussed in this section can be distinguished as “Big” web services and “RESTful” web services.

- In Java EE 6, JAX-WS provides the functionality for “Big” web services. Big web services use Extensible Markup Language (XML) messages that follow the Simple Object Access Protocol (SOAP) standard, which is an XML language defining a message architecture and message formats. In such systems, there is often a machine-readable description of the operations offered by the service written in the Web Services Description Language (WSDL), which is an XML language for defining interfaces syntactically.

The SOAP message format and the WSDL interface definition language have gained widespread adoption and there are many development tools available, such as NetBeans IDE, that reduce the complexity of developing web service applications.

A SOAP-based design must include the following elements:

- A formal contract must be established to describe the interface that the web service offers. The Web Services Description Language (WSDL) can be used to describe the details of the contract, which may include messages, operations, bindings, and the location of the web service. You may also process SOAP messages in a JAX-WS service without publishing a WSDL.
- The architecture must address complex nonfunctional requirements. Many web service specifications address such requirements and establish a common vocabulary for them. Examples include Transactions, Security, Addressing, Trust, Coordination, and so on.
- The architecture needs to handle asynchronous processing and invocation. In such cases, the infrastructure provided by standards such as WSRM and APIs such as JAX-WS with their client-side asynchronous invocation support can be leveraged out of the box.

“Big” web services are described in [Chapter 13](#), “[Building Web Services with JAX-WS](#).”

- In Java EE 6, JAX-RS provides the functionality for REpresentational State Transfer (RESTful) Web Services. REST is well suited for basic, ad hoc integration scenarios. RESTful web services are often better integrated with HTTP than SOAP-based services are. They do not require XML messages or WSDL service-API definitions.

Project Jersey is the production-ready reference implementation for [JSR 311](#): JAX-RS: The Java API for RESTful Web Services. Jersey implements support for the annotations defined in JSR-311, making it easy for developers to build RESTful web services with Java and the Java JVM. Jersey also adds [additional features](#) not specified by the JSR.

Because RESTful web services use existing well-known W3C/IETF standards (HTTP, XML, URI, MIME), and have a lightweight infrastructure, where services can be built with minimal tooling, developing RESTful web services is inexpensive and thus has a very low barrier for adoption. You can use one of the development tools, such as NetBeans IDE, to further reduce the complexity of developing RESTful web services.

A few real-world web applications that use RESTful web services include most blog sites. These are considered RESTful in that most blog sites involve downloading XML files in RSS or Atom format which contain lists of links to other resources. Other web sites and web applications that use REST-like developer interfaces to connect to data include Twitter and Amazon Simple Storage Service (S3). With Amazon S3, buckets and objects can be created,

listed, and retrieved using either a REST-style HTTP interface or a SOAP interface. The examples that ship with Project Jersey include a storage service example with a RESTful interface. The tutorial at <http://netbeans.org/kb/docs/websvc/twitter-swing.html> uses the NetBeans IDE to create a simple, graphical, REST-based client that displays Twitter public time line messages and lets you view and update your Twitter status.

A RESTful design may be appropriate in the following situation:

- The web services are completely stateless. A good test is to consider whether the interaction can survive a restart of the server.
- A caching infrastructure can be leveraged for performance. If the data that the web service returns is not dynamically generated and can be cached, then the caching infrastructure that web servers and other intermediaries inherently provide can be leveraged to improve performance. However, the developer must take care because such caches are limited to the HTTP GET method for most servers.
- The service producer and service consumer have a mutual understanding of the context and content being passed along. Because there is no formal way to describe the web services interface, both parties must agree out of band on the schemas that describe the data being exchanged and on ways to process it meaningfully. In the real world, most commercial applications that expose services as RESTful implementations also distribute so-called value-added toolkits that describe the interfaces to developers in popular programming languages.
- Bandwidth is particularly important and needs to be limited. REST is particularly useful for limited-profile devices such as PDAs and mobile phones, for which the overhead of headers and additional layers of SOAP elements on the XML payload must be restricted.
- Web service delivery or aggregation into existing web sites can be enabled easily with a RESTful style. Developers can use technologies such as JAX-RS, Asynchronous JavaScript with XML (AJAX) and toolkits such as Direct Web Remoting (DWR) to consume the services in their web applications. Rather than starting from scratch, services can be exposed with XML and consumed by HTML pages without significantly refactoring the existing web site architecture. Existing developers will be more productive because they are adding to something they are already familiar with, rather than having to start from scratch with new technology.

RESTful web services are discussed in [Chapter 14](#), “Building RESTful Web Services with JAX-RS and Jersey.” This chapter contains information about generating the skeleton of a RESTful web service using both NetBeans IDE and the Maven project management tool.

## Deciding Which Type of Web Service to Use

Basically, you would want to use RESTful web services for integration over the Web and use Big web services in enterprise application integration scenarios that have advanced QoS requirements. This topic is discussed in more detail in the following sections.

---

**Note** – For an article that provides more in-depth analysis of this issue, see *RESTful Web Services vs. “Big” Web Services: Making the Right Architectural Decision* by Cesare Pautasso, Olaf Zimmermann, and Frank Leymann from the [WWW '08: Proceedings of the 17th International Conference on the World Wide Web \(2008\)](#), pp. 805-814.

---

### When Should I Use JAX-WS?

JAX-WS addresses advanced quality of service (QoS) requirements commonly occurring in enterprise computing. When compared to JAX-RS, JAX-WS makes it easier to support the WS-\* set of protocols (which provide standards for security and reliability, among other things) and interoperate with other WS-\* conforming clients and servers.

### When Should I Use JAX-RS?

When compared with JAX-WS, JAX-RS makes it easier to write applications for the web that apply some or all of the constraints of the REST style to induce desirable properties in the application like loose coupling (evolving the server is easier without breaking existing clients), scalability (start small and grow), and architectural simplicity (use off-the-shelf components like proxies, HTTP routers, or others). You would choose to use JAX-RS for your web application because it is easier for many types of clients to consume RESTful web services while enabling the server side to evolve and scale. Clients can choose to consume some or all aspects of the service and mash it up with other web-based services.

## Building Web Services with JAX-WS

---

JAX-WS stands for Java API for XML Web Services. JAX-WS is a technology for building web services and clients that communicate using XML. JAX-WS allows developers to write message-oriented as well as RPC-oriented web services.

In JAX-WS, a web service operation invocation is represented by an XML-based protocol such as SOAP. The SOAP specification defines the envelope structure, encoding rules, and conventions for representing web service invocations and responses. These calls and responses are transmitted as SOAP messages (XML files) over HTTP.

Although SOAP messages are complex, the JAX-WS API hides this complexity from the application developer. On the server side, the developer specifies the web service operations by defining methods in an interface written in the Java programming language. The developer also codes one or more classes that implement those methods. Client programs are also easy to code. A client creates a proxy (a local object representing the service) and then simply invokes methods on the proxy. With JAX-WS, the developer does not generate or parse SOAP messages. It is the JAX-WS runtime system that converts the API calls and responses to and from SOAP messages.

With JAX-WS, clients and web services have a big advantage: the platform independence of the Java programming language. In addition, JAX-WS is not restrictive: a JAX-WS client can access a web service that is not running on the Java platform, and vice versa. This flexibility is possible because JAX-WS uses technologies defined by the World Wide Web Consortium (W3C): HTTP, SOAP, and the Web Service Description Language (WSDL). WSDL specifies an XML format for describing a service as a set of endpoints operating on messages.

## Setting the Port

Several files in the JAX-WS examples depend on the port that you specified when you installed the Enterprise Server. The tutorial examples assume that the server runs on the default port, 8080. If you have changed the port, you must update the port number in the following file before building and running the JAX-WS examples:

```
tut-install/examples/jaxws/simpleclient/src/java/simpleclient/HelloClient.java
```

## Creating a Simple Web Service and Client with JAX-WS

This section shows how to build and deploy a simple web service and client. The source code for the service is in `tut-install/examples/jaxws/helloservice/` and the client is in `tut-install/examples/jaxws/simpleclient/`.

Figure 13–1 illustrates how JAX-WS technology manages communication between a web service and client.

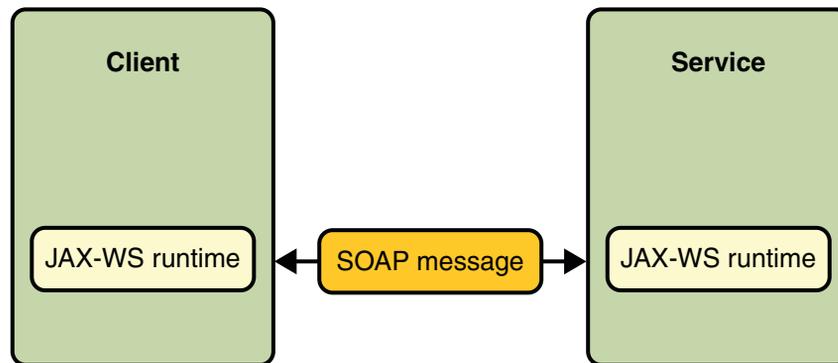


FIGURE 13–1 Communication between a JAX-WS Web Service and a Client

The starting point for developing a JAX-WS web service is a Java class annotated with the `javax.jws.WebService` annotation. The `@WebService` annotation defines the class as a web service endpoint.

A *service endpoint interface* or *service endpoint implementation* (SEI) is a Java interface or class, respectively, that declares the methods that a client can invoke on the service. An interface is not required when building a JAX-WS endpoint. The web service implementation class implicitly defines an SEI.

You may specify an explicit interface by adding the `endpointInterface` element to the `@WebService` annotation in the implementation class. You must then provide an interface that defines the public methods made available in the endpoint implementation class.

You use the endpoint implementation class and the `wsgen` tool to generate the web service artifacts that connect a web service client to the JAX-WS runtime. For reference documentation on `wsgen`, see the [Sun GlassFish Enterprise Server v3 Reference Manual](#).

Together, the `wsgen` tool and the Enterprise Server provide the Enterprise Server's implementation of JAX-WS.

These are the basic steps for creating the web service and client:

1. Code the implementation class.
2. Compile the implementation class.
3. Use `wsgen` to generate the artifacts required to deploy the service.
4. Package the files into a WAR file.
5. Deploy the WAR file. The web service artifacts (which are used to communicate with clients) are generated by the Enterprise Server during deployment.
6. Code the client class.
7. Use `wsimport` to generate and compile the web service artifacts needed to connect to the service.
8. Compile the client class.
9. Run the client.

The sections that follow cover these steps in greater detail.

## Requirements of a JAX-WS Endpoint

JAX-WS endpoints must follow these requirements:

- The implementing class must be annotated with either the `javax.jws.WebService` or `javax.jws.WebServiceProvider` annotation.
- The implementing class may explicitly reference an SEI through the `endpointInterface` element of the `@WebService` annotation, but is not required to do so. If no `endpointInterface` is specified in `@WebService`, an SEI is implicitly defined for the implementing class.
- The business methods of the implementing class must be public, and must not be declared `static` or `final`.
- Business methods that are exposed to web service clients must be annotated with `javax.jws.WebMethod`.

- Business methods that are exposed to web service clients must have JAXB-compatible parameters and return types. See [JAXB default data type bindings \(http://java.sun.com/javase/5/docs/tutorial/doc/bnazq.html#bnazs\)](http://java.sun.com/javase/5/docs/tutorial/doc/bnazq.html#bnazs).
- The implementing class must not be declared `final` and must not be `abstract`.
- The implementing class must have a default public constructor.
- The implementing class must not define the `finalize` method.
- The implementing class may use the `javax.annotation.PostConstruct` or `javax.annotation.PreDestroy` annotations on its methods for life cycle event callbacks. The `@PostConstruct` method is called by the container before the implementing class begins responding to web service clients. The `@PreDestroy` method is called by the container before the endpoint is removed from operation.

## Coding the Service Endpoint Implementation Class

In this example, the implementation class, `Hello`, is annotated as a web service endpoint using the `@WebService` annotation. `Hello` declares a single method named `sayHello`, annotated with the `@WebMethod` annotation. `@WebMethod` exposes the annotated method to web service clients. `sayHello` returns a greeting to the client, using the name passed to `sayHello` to compose the greeting. The implementation class also must define a default, public, no-argument constructor.

```
package helloservice.endpoint;

import javax.jws.WebService;

@WebService
public class Hello {
    private String message = new String("Hello, ");

    public void Hello() {}

    @WebMethod
    public String sayHello(String name) {
        return message + name + ".";
    }
}
```

## Building, Packaging, and Deploying the Service

You can build, package, and deploy the `helloservice` application using either NetBeans IDE or `ant`.

## Building, Packaging, and Deploying the Service Using NetBeans IDE

Follow these instructions to build, package, and deploy the `helloservice` example to your Application Server instance using the NetBeans IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to `tut-install/examples/jaxws/`.
3. Select the `helloservice` folder.
4. Select the Open as Main Project check box.
5. Click Open Project.
6. In the Projects tab, right-click the `helloservice` project and select Undeploy and Deploy.

This builds and packages to application into `helloservice.war`, located in `tut-install/examples/jaxws/helloservice/dist/`, and deploys this WAR file to your Application Server instance.

## Building, Packaging, and Deploying the Service Using Ant

To build and package `helloservice` using Ant, in a terminal window, go to the `tut-install/examples/jaxws/helloservice/` directory and type the following:

```
ant
```

This command calls the default target, which builds and packages the application into an WAR file, `helloservice.war`, located in the `dist` directory.

To deploy the `helloservice` example, follow these steps:

1. In a terminal window, go to `tut-install/examples/jaxws/helloservice/`.
2. Make sure the Enterprise Server is started.
3. Run `ant deploy`.

You can view the WSDL file of the deployed service by requesting the URL `http://localhost:8080/helloservice/hello?WSDL` in a web browser. Now you are ready to create a client that accesses this service.

## Undeploying the Service

At this point in the tutorial, do not undeploy the service. When you are finished with this example, you can undeploy the service by typing this command:

```
ant undeploy
```

## The all Task

As a convenience, the `all` task will build, package, and deploy the application. To do this, enter the following command:

```
ant all
```

## Testing the Service without a Client

Enterprise Server allows you to test the methods of a web service endpoint. To test the `sayHello` method of `HelloService`, do the following:

1. Open the web service test interface by entering the following URL in a web browser:

```
http://localhost:8080/helloservice/HelloService?Tester
```

2. Under Methods, enter a name as the parameter to the `sayHello` method.
3. Click the `sayHello` button.

This will take you to the `sayHello` Method invocation page.

4. Under Method returned, you'll see the response from the endpoint.

## A Simple JAX-WS Client

`HelloClient` is a standalone Java program that accesses the `sayHello` method of `HelloService`. It makes this call through a port, a local object that acts as a proxy for the remote service. The port is created at development time by the `wsimport` tool, which generates JAX-WS portable artifacts based on a WSDL file.

### Coding the Client

When invoking the remote methods on the port, the client performs these steps:

1. Uses the generated `helloservice.endpoint.HelloService` class which represents the service at the URI of the deployed service's WSDL file.

```
HelloService service = new HelloService();
```

2. Retrieves a proxy to the service, also known as a port, by invoking `getHelloPort` on the service.

```
Hello port = service.getHelloPort();
```

The port implements the SEI defined by the service.

3. Invokes the port's `sayHello` method, passing to the service a name.

```
String response = port.sayHello(name);
```

Here is the full source of `HelloClient`, which is located in the `tut-install/examples/jaxws/simpleclient/src/java/` directory.

```
package simpleclient;

import javax.xml.ws.WebServiceRef;
import helloservice.endpoint.HelloService;
import helloservice.endpoint.Hello;

public class HelloClient {

    public static void main(String[] args) {
        try {
            HelloClient client = new HelloClient();
            client.doTest(args);
        } catch (Exception e) {
            e.printStackTrace();
        }
    }

    public void doTest(String[] args) {
        try {
            System.out.println("Retrieving the port from
                the following service: " + service);
            HelloService service = new HelloService();
            Hello port = service.getHelloPort();
            System.out.println("Invoking the sayHello operation
                on the port.");

            String name;
            if (args.length > 0) {
                name = args[0];
            } else {
                name = "No Name";
            }

            String response = port.sayHello(name);
            System.out.println(response);
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

## Building and Running the Client

You can build and run the `simpleclient` application using either NetBeans IDE or `ant`. To build the client, you must first have deployed `helloservice`, as described in “[Building, Packaging, and Deploying the Service](#)” on page 240.

### Building and Running the Client in NetBeans IDE

Do the following to build and run `simpleclient`:

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to `tut-install/examples/jaxws/`.
3. Select the `simpleclient` folder.
4. Select the Open as Main Project check box.
5. Click Open Project.
6. In the Projects tab, right-click the `simpleclient` project and select Run.

You will see the output of the application client in the Output pane.

### Building and Running the Client Using Ant

In a terminal navigate to `tut-install/examples/jaxws/simpleclient/` and type the following command:

```
ant
```

This command calls the default target, which builds and packages the application into a JAR file, `simpleclient.jar`, located in the `dist` directory.

To run the client, type the following command:

```
ant run
```

## Types Supported by JAX-WS

JAX-WS delegates the mapping of Java programming language types to and from XML definitions to JAXB. Application developers don't need to know the details of these mappings, but they should be aware that not every class in the Java language can be used as a method parameter or return type in JAX-WS. For information on which types are supported by JAXB, see [JAXB default data type bindings \(http://java.sun.com/javase/5/docs/tutorial/doc/bnazq.html#bnazs\)](http://java.sun.com/javase/5/docs/tutorial/doc/bnazq.html#bnazs).

## Web Services Interoperability and JAX-WS

JAX-WS 2.0 supports the Web Services Interoperability (WS-I) Basic Profile Version 1.1. The WS-I Basic Profile is a document that clarifies the SOAP 1.1 and WSDL 1.1 specifications to promote SOAP interoperability. For links related to WS-I, see “[Further Information about JAX-WS](#)” on page 245.

To support WS-I Basic Profile Version 1.1, the JAX-WS runtime supports doc/literal and rpc/literal encodings for services, static ports, dynamic proxies, and DII.

## Further Information about JAX-WS

For more information about JAX-WS and related technologies, see:

- Java API for XML Web Services 2.0 specification  
<https://jax-ws.dev.java.net/spec-download.html>
- JAX-WS home  
<https://jax-ws.dev.java.net/>
- Simple Object Access Protocol (SOAP) 1.2 W3C Note  
<http://www.w3.org/TR/soap/>
- Web Services Description Language (WSDL) 1.1 W3C Note  
<http://www.w3.org/TR/wsd1>
- WS-I Basic Profile 1.1  
<http://www.ws-i.org>



# Building RESTful Web Services with JAX-RS and Jersey

---

PUThis chapter describes the REST architecture, RESTful web services, and Sun's reference implementation for JAX-RS (Java™ API for RESTful Web Services, [JSR-311](#)), which is referred to as *Jersey*.

## What are RESTful Web Services?

*RESTful web services* are services that are built to work best on the web. Representational State Transfer (REST) is an architectural style that specifies constraints, such as the *uniform interface*, that if applied to a web service induce desirable properties, such as performance, scalability, and modifiability, that enable services to work best on the Web. In the REST architectural style, data and functionality are considered resources, and these *resources are accessed using Uniform Resource Identifiers (URIs)*, typically links on the web. The resources are acted upon by using a set of simple, well-defined operations. The REST architectural style constrains an architecture to a client-server architecture, and is designed to use a *stateless communication protocol*, typically HTTP. In the REST architecture style, clients and servers exchange *representations of resources* using a standardized interface and protocol. These principles encourages RESTful applications to be simple, lightweight, and have high performance.

A paper that expands on the basic principles of REST technology can be found at <http://www2008.org/papers/pdf/p805-pautassoA.pdf>.

- *Resource identification through URI*. A RESTful Web service exposes a set of resources which identify the targets of the interaction with its clients. Resources are identified by URIs, which provide a global addressing space for resource and service discovery. This topic is discussed in “[The @Path Annotation and URI Path Templates](#)” on page 251.
- *Uniform interface*. Resources are manipulated using a fixed set of four create, read, update, delete operations: PUT, GET, POST, and DELETE. PUT creates a new resource, which can be then deleted using DELETE. GET retrieves the current state of a resource in some representation. POST transfers a new state onto a resource. This topic is discussed in “[Responding to HTTP Resources](#)” on page 254.

- *Self-descriptive messages.* Resources are decoupled from their representation so that their content can be accessed in a variety of formats (such as HTML, XML, plain text, PDF, JPEG, JSON, and others). Metadata about the resource is available and used, for example, to control caching, detect transmission errors, negotiate the appropriate representation format, and perform authentication or access control. This topic is discussed in [“Responding to HTTP Resources” on page 254](#) and [“Using Entity Providers to Map HTTP Response and Request Entity Bodies” on page 255](#).
- *Stateful interactions through hyperlinks.* Every interaction with a resource is stateless; that is, request messages are self-contained. Stateful interactions are based on the concept of explicit state transfer. Several techniques exist to exchange state, such as URI rewriting, cookies, and hidden form fields. State can be embedded in response messages to point to valid future states of the interaction. This topic is discussed somewhat in [“Using Entity Providers to Map HTTP Response and Request Entity Bodies” on page 255](#), is discussed somewhat in the section *Building URIs* in the [JAX-RS Overview document](#), and may be discussed in more detail in a forthcoming advanced version of this tutorial.

## Where Does Jersey Fit In?

Jersey is Sun's production quality reference implementation for [JSR 311: JAX-RS: The Java API for RESTful Web Services](#). Jersey implements support for the annotations defined in JSR-311, making it easy for developers to build RESTful web services with Java and the Java JVM. Jersey also adds [additional features](#) not specified by the JSR.

The latest version of the JAX-RS API's can be viewed at <https://jsr311.dev.java.net/nonav/javadoc/index.html>

If you are developing with Enterprise Server v3, you can install the Jersey samples and documentation using the Update Tool. Instructions for using the Update Tool can be found in the section [“Java EE 6 Tutorial Component” on page 61](#).

## Creating a RESTful Root Resource Class

*Root resource classes* are POJOs (Plain Old Java Objects) that are either annotated with `@Path` or have at least one method annotated with `@Path` or a *request method designator* such as `@GET`, `@PUT`, `@POST`, or `@DELETE`. *Resource methods* are methods of a resource class annotated with a request method designator. This section describes how to use Jersey to annotate Java objects to create RESTful web services.

## Developing RESTful Web Services with JAX-RS and Jersey

The JAX-RS API for developing RESTful web services is a Java programming language API designed to make it easy to develop applications that use the REST architecture.

The JAX-RS API uses Java programming language annotations to simplify the development of RESTful web services. Developers decorate Java programming language class files with HTTP-specific annotations to define resources and the actions that can be performed on those resources. Jersey annotations are runtime annotations, therefore, runtime reflection will generate the helper classes and artifacts for the resource, and then the collection of classes and artifacts will be built into a web application archive (WAR). The resources are exposed to clients by deploying the WAR to a Java EE or web server.

Here is a listing of some of the Java programming annotations that are defined by JAX-RS, with a brief description of how each is used. Further information on the JAX-RS API's can be viewed at <https://jsr311.dev.java.net/nonav/javadoc/index.html>.

TABLE 14-1 Summary of Jersey Annotations

Annotation	Description
@Path	The @Path annotation's value is a relative URI path indicating where the Java class will be hosted, for example, /helloWorld. You can also embed variables in the URIs to make a URI path template. For example, you could ask for the name of a user, and pass it to the application as a variable in the URI, like this, /helloWorld/{username}.
@GET	The @GET annotation is a request method designator and corresponds to the similarly named HTTP method. The Java method annotated with this request method designator will process HTTP GET requests. The behavior of a resource is determined by the HTTP method to which the resource is responding.
@POST	The @POST annotation is a request method designator and corresponds to the similarly named HTTP method. The Java method annotated with this request method designator will process HTTP POST requests. The behavior of a resource is determined by the HTTP method to which the resource is responding.
@PUT	The @PUT annotation is a request method designator and corresponds to the similarly named HTTP method. The Java method annotated with this request method designator will process HTTP PUT requests. The behavior of a resource is determined by the HTTP method to which the resource is responding.
@DELETE	The @DELETE annotation is a request method designator and corresponds to the similarly named HTTP method. The Java method annotated with this request method designator will process HTTP DELETE requests. The behavior of a resource is determined by the HTTP method to which the resource is responding.

TABLE 14-1 Summary of Jersey Annotations (Continued)

Annotation	Description
@HEAD	The @HEAD annotation is a request method designator and corresponds to the similarly named HTTP method. The Java method annotated with this request method designator will process HTTP HEAD requests. The behavior of a resource is determined by the HTTP method to which the resource is responding.
@PathParam	The @PathParam annotation is a type of parameter that you can extract for use in your resource class. URI path parameters are extracted from the request URI, and the parameter names correspond to the URI path template variable names specified in the @Path class-level annotation.
@QueryParam	The @QueryParam annotation is a type of parameter that you can extract for use in your resource class. Query parameters are extracted from the request URI query parameters.
@Consumes	The @Consumes annotation is used to specify the MIME media types of representations a resource can consume that were sent by the client.
@Produces	The @Produces annotation is used to specify the MIME media types of representations a resource can produce and send back to the client, for example, "text/plain".
@Provider	The @Provider annotation is used for anything that is of interest to the JAX-RS runtime, such as <code>MessageBodyReader</code> and <code>MessageBodyWriter</code> . For HTTP requests, the <code>MessageBodyReader</code> is used to map an HTTP request entity body to method parameters. On the response side, a return value is mapped to an HTTP response entity body using a <code>MessageBodyWriter</code> . If the application needs to supply additional metadata, such as HTTP headers or a different status code, a method can return a <code>Response</code> that wraps the entity, and which can be built using <code>Response.Builder</code> .

## Overview of a Jersey-Annotated Application

The following code sample is a very simple example of a root resource class using JAX-RS annotations. The sample shown here is from the samples that ship with Jersey, and which can be found in the following directory of that installation:

```
jersey/samples/helloworld/src/main/java/com/sun/jersey/samples/helloworld/resources/HelloWorldResource.java
```

```
package com.sun.jersey.samples.helloworld.resources;

import javax.ws.rs.GET;
import javax.ws.rs.Produces;
import javax.ws.rs.Path;

// The Java class will be hosted at the URI path "/helloworld"
@Path("/helloworld")
public class HelloWorldResource {
```

```

    // The Java method will process HTTP GET requests
    @GET
    // The Java method will produce content identified by the MIME Media
    // type "text/plain"
    @Produces("text/plain")
    public String getClichedMessage() {
        // Return some cliched textual content
        return "Hello World";
    }
}

```

The following sections describe the annotations used in this example.

- The `@Path` annotation's value is a relative URI path. In the example above, the Java class will be hosted at the URI path `/helloworld`. This is an extremely simple use of the `@Path` annotation. What makes JAX-RS so useful is that you can embed variables in the URIs. *URI path templates* are URIs with variables embedded within the URI syntax.
- The `@GET` annotation is a request method designator, along with `@POST`, `@PUT`, `@DELETE`, and `@HEAD`, that is defined by JAX-RS, and which correspond to the similarly named HTTP methods. In the example above, the annotated Java method will process HTTP GET requests. The behavior of a resource is determined by the HTTP method to which the resource is responding.
- The `@Produces` annotation is used to specify the MIME media types of representations a resource can produce and send back to the client. In this example, the Java method will produce representations identified by the MIME media type `"text/plain"`.
- The `@Consumes` annotation is used to specify the MIME media types of representations a resource can consume that were sent by the client. The above example could be modified to set the cliched message as shown in this code example.

```

@POST
@Consumes("text/plain")
public void postClichedMessage(String message) {
    // Store the message
}

```

## The `@Path` Annotation and URI Path Templates

The `@Path` annotation identifies the URI path template to which the resource responds, and is specified at the class level of a resource. The `@Path` annotation's value is a partial URI path template relative to the base URI of the server on which the resource is deployed, the context root of the WAR, and the URL pattern to which the Jersey helper servlet responds.

*URI path templates* are URIs with variables embedded within the URI syntax. These variables are substituted at runtime in order for a resource to respond to a request based on the substituted URI. Variables are denoted by curly braces. For example, look at the following `@Path` annotation:

```
@Path("/users/{username}")
```

In this type of example, a user will be prompted to enter their name, and then a Jersey web service configured to respond to requests to this URI path template will respond. For example, if the user entered their user name as Galileo, the web service will respond to the following URL:

```
http://example.com/users/Galileo
```

To obtain the value of the username variable, the `@PathParam` annotation may be used on the method parameter of a request method, as shown in the following code example.

```
@Path("/users/{username}")
public class UserResource {

    @GET
    @Produces("text/xml")
    public String getUser(@PathParam("username") String userName) {
        ...
    }
}
```

If it is required that a user name must only consist of lower and upper case numeric characters, it is possible to declare a particular regular expression that will override the default regular expression, "[^/]+?". The following example shows how this could be used with the `@Path` annotation.

```
@Path("users/{username: [a-zA-Z][a-zA-Z_0-9]}")
```

In this type of example the username variable will only match user names that begin with one upper or lower case letter and zero or more alpha numeric characters and the underscore character. If a user name does not match that template, then a 404 (Not Found) response will occur.

An `@Path` value may or may not begin with a forward slash (/), it makes no difference. Likewise, by default, an `@Path` value may or may not end in a forward slash (/), it makes no difference, and thus request URLs that end or do not end with a forward slash will both be matched. However, Jersey has a redirection mechanism, which, if enabled, automatically performs redirection to a request URL ending in a / if a request URL does not end in a / and the matching `@Path` does end in a /.

## More on URI Path Template Variables

A URI path template has one or more variables, with each variable name surrounded by curly braces, { to begin the variable name and } to end it. In the example above, username is the variable name. At runtime, a resource configured to respond to the above URI path template will attempt to process the URI data that corresponds to the location of {username} in the URI as the variable data for username.

For example, if you want to deploy a resource that responds to the URI path template `http://example.com/myContextRoot/jerseybeans/{name1}/{name2}/`, you must deploy the WAR to a Java EE server that responds to requests to the `http://example.com/myContextRoot` URI, and then decorate your resource with the following `@Path` annotation:

```
@Path("/{name1}/{name2}/")
public class SomeResource {
    ...
}
```

In this example, the URL pattern for the Jersey helper servlet, specified in `web.xml`, is the default:

```
<servlet-mapping>
    <servlet-name>My Jersey Bean Resource</servlet-name>
    <url-pattern>/jerseybeans/*</url-pattern>
</servlet-mapping>
```

A variable name can be used more than once in the URI path template.

If a character in the value of a variable would conflict with the reserved characters of a URI, the conflicting character should be substituted with percent encoding. For example, spaces in the value of a variable should be substituted with `%20`.

Be careful when defining URI path templates that the resulting URI after substitution is valid.

The following table lists some examples of URI path template variables and how the URIs are resolved after substitution. The following variable names and values are used in the examples:

- `name1:jay`
- `name2:gatsby`
- `name3:`
- `location:Main%20Street`
- `question:why`

---

**Note** – The value of the `name3` variable is an empty string.

---

TABLE 14-2 Examples of URI path templates

URI Path Template	URI After Substitution
<code>http://example.com/{name1}/{name2}/</code>	<code>http://example.com/jay/gatsby/</code>
<code>http://example.com/{question}/</code>	<code>http://example.com/why/why/why/</code>
<code>{question}/{question}/</code>	

TABLE 14-2 Examples of URI path templates (Continued)

URI Path Template	URI After Substitution
<code>http://example.com/maps/{location}</code>	<code>http://example.com/maps/Main%20Street</code>
<code>http://example.com/{name3}/home/</code>	<code>http://example.com//home/</code>

## Responding to HTTP Resources

The behavior of a resource is determined by the HTTP methods (typically, GET, POST, PUT, DELETE) to which the resource is responding.

### The Request Method Designator Annotations

A *request method designator* annotations are runtime annotations, defined by JAX-RS, and which correspond to the similarly named HTTP methods. Within a resource class file, HTTP methods are mapped to Java programming language methods using the request method designator annotations. The behavior of a resource is determined by which of the HTTP methods the resource is responding to. Jersey defines a set of request method designators for the common HTTP methods: @GET, @POST, @PUT, @DELETE, @HEAD, but you can create your own custom request method designators. Creating custom request method designators is outside the scope of this document.

The following example is an extract from the storage service sample that shows the use of the PUT method to create or update a storage container.

```
@PUT
public Response putContainer() {
    System.out.println("PUT CONTAINER " + container);

    URI uri = uriInfo.getAbsolutePath();
    Container c = new Container(container, uri.toString());

    Response r;
    if (!MemoryStore.MS.hasContainer(c)) {
        r = Response.created(uri).build();
    } else {
        r = Response.noContent().build();
    }

    MemoryStore.MS.createContainer(c);
    return r;
}
```

By default the JAX-RS runtime will automatically support the methods HEAD and OPTIONS if not explicitly implemented. For HEAD, the runtime will invoke the implemented GET method (if

present) and ignore the response entity (if set). For `OPTIONS`, the `Allow` response header will be set to the set of HTTP methods supported by the resource. In addition Jersey will return a [WADL](#) document describing the resource.

Methods decorated with request method designators must return `void`, a Java programming language type, or a `javax.ws.rs.core.Response` object. Multiple parameters may be extracted from the URI using the `PathParam` or `QueryParam` annotations as described in [“Extracting Request Parameters” on page 259](#). Conversion between Java types and an entity body is the responsibility of an entity provider, such as `MessageBodyReader` or `MessageBodyWriter`. Methods that need to provide additional metadata with a response should return an instance of `Response`. The `ResponseBuilder` class provides a convenient way to create a `Response` instance using a builder pattern. The HTTP `PUT` and `POST` methods expect an HTTP request body, so you should use a `MessageBodyReader` for methods that respond to `PUT` and `POST` requests.

As `PUT` and `POST` can post be used to create or update, here is a bit more information on when you'd use each:

- `PUT` has defined/specified semantics. `POST` can mean anything, so when using `POST` it is up to the application to define the semantics.
- When using `PUT` for creation, the client declares the URI for the newly created resource. `PUT` has very clear semantics for creating and updating. The representation the client sends must be the same representation that is received using a `GET`, given the same media type. It does not specify partial update, a mistake people often make. A common application pattern is to use `POST` to create a resource and return a `201` response with a location header whose value is the URI to the newly-created resource. Thus in this pattern, the web service declares the URI for the newly-created resource.

## Using Entity Providers to Map HTTP Response and Request Entity Bodies

*Entity providers* supply mapping services between representations and their associated Java types. There are two types of entity providers: `MessageBodyReader` and `MessageBodyWriter`. For HTTP requests, the `MessageBodyReader` is used to map an HTTP request entity body to method parameters. On the response side, a return value is mapped to an HTTP response entity body using a `MessageBodyWriter`. If the application needs to supply additional metadata, such as HTTP headers or a different status code, a method can return a `Response` that wraps the entity, and which can be built using `Response.ResponseBuilder`.

The following list contains the standard types that are supported automatically for entities. You only need to write an entity provider if you are not choosing one of the following, standard types.

- `byte[]` — All media types (`*/*`)
- `java.lang.String` — All text media types (`text/*`)
- `java.io.InputStream` — All media types (`*/*`)

- `java.io.Reader` — All media types (`/*/*`)
- `java.io.File` — All media types (`/*/*`)
- `javax.activation.DataSource` — All media types (`/*/*`)
- `javax.xml.transform.Source` — XML types (`text/xml`, `application/xml`, and `application/*+xml`)
- `javax.xml.bind.JAXBElement` and application-supplied JAXB classes XML media types (`text/xml`, `application/xml`, and `application/*+xml`)
- `MultivaluedMap<String, String>` — Form content (`application/x-www-form-urlencoded`)
- `StreamingOutput` — All media types (`/*/*`), `MessageBodyWriter` only

The following example shows how to use `MessageBodyReader` with the `@Consumes` and `@Provider` annotations:

```
@Consumes("application/x-www-form-urlencoded")
@Provider
public class FormReader implements MessageBodyReader<NameValuePair> {
```

The following example shows how to use `MessageBodyWriter` with the `@Produces` and `@Provider` annotations:

```
@Produces("text/html")
@Provider
public class FormWriter implements MessageBodyWriter<Hashtable<String, String>> {
```

The following example shows how to use `ResponseBuilder`:

```
@GET
public Response getItem() {
    System.out.println("GET ITEM " + container + " " + item);

    Item i = MemoryStore.MS.getItem(container, item);
    if (i == null)
        throw new NotFoundException("Item not found");
    Date lastModified = i.getLastModified().getTime();
    EntityTag et = new EntityTag(i.getDigest());
    ResponseBuilder rb = request.evaluatePreconditions(lastModified, et);
    if (rb != null)
        return rb.build();

    byte[] b = MemoryStore.MS.getItemData(container, item);
    return Response.ok(b, i.getMimeType())
        .lastModified(lastModified).tag(et).build();
}
```

## Using @Consumes and @Produces to Customize Requests and Responses

The information sent to a resource and then passed back to the client is specified as a MIME media type in the headers of an HTTP request or response. You can specify which MIME media types of representations a resource can respond to or produce by using the `javax.ws.rs.Consumes` and `javax.ws.rs.Produces` annotations.

By default, a resource class can respond to and produce all MIME media types of representations specified in the HTTP request and response headers.

### The @Produces Annotation

The `@Produces` annotation is used to specify the MIME media types or representations a resource can produce and send back to the client. If `@Produces` is applied at the class level, all the methods in a resource can produce the specified MIME types by default. If it is applied at the method level, it overrides any `@Produces` annotations applied at the class level.

If no methods in a resource are able to produce the MIME type in a client request, the Jersey runtime sends back an HTTP “406 Not Acceptable” error.

The value of `@Produces` is an array of `String` of MIME types. For example:

```
@Produces({"image/jpeg", "image/png"})
```

The following example shows how to apply `@Produces` at both the class and method levels:

```
@Path("/myResource")
@Produces("text/plain")
public class SomeResource {
    @GET
    public String doGetAsPlainText() {
        ...
    }

    @GET
    @Produces("text/html")
    public String doGetAsHtml() {
        ...
    }
}
```

The `doGetAsPlainText` method defaults to the MIME media type of the `@Produces` annotation at the class level. The `doGetAsHtml` method's `@Produces` annotation overrides the class-level `@Produces` setting, and specifies that the method can produce HTML rather than plain text.

If a resource class is capable of producing more than one MIME media type, the resource method chosen will correspond to the most acceptable media type as declared by the client. More specifically, the `Accept` header of the HTTP request declares what is most acceptable. For example if the `Accept` header is `Accept: text/plain`, the `doGetAsPlainText` method will be invoked. Alternatively if the `Accept` header is `Accept: text/plain;q=0.9, text/html`, which declares that the client can accept media types of `text/plain` and `text/html`, but prefers the latter, then the `doGetAsHtml` method will be invoked.

More than one media type may be declared in the same `@Produces` declaration. The following code example shows how this is done.

```
@Produces({"application/xml", "application/json"})
public String doGetAsXmlOrJson() {
    ...
}
```

The `doGetAsXmlOrJson` method will get invoked if either of the media types `application/xml` and `application/json` are acceptable. If both are equally acceptable, then the former will be chosen because it occurs first. The examples above refer explicitly to MIME media types for clarity. It is possible to refer to constant values, which may reduce typographical errors. For more information, see the constant field values of [MediaType](#).

## The `@Consumes` Annotation

The `@Consumes` annotation is used to specify which MIME media types of representations a resource can accept, or consume, from the client. If `@Consumes` is applied at the class level, all the response methods accept the specified MIME types by default. If `@Consumes` is applied at the method level, it overrides any `@Consumes` annotations applied at the class level.

If a resource is unable to consume the MIME type of a client request, the Jersey runtime sends back an HTTP “415 Unsupported Media Type” error.

The value of `@Consumes` is an array of `String` of acceptable MIME types. For example:

```
@Consumes({"text/plain, text/html"})
```

The following example shows how to apply `@Consumes` at both the class and method levels:

```
@Path("/myResource")
@Consumes("multipart/related")
public class SomeResource {
    @POST
    public String doPost(MimeMultipart mimeMultipartData) {
        ...
    }

    @POST
```

```

    @Consumes("application/x-www-form-urlencoded")
    public String doPost2(FormURLEncodedProperties formData) {
        ...
    }
}

```

The `doPost` method defaults to the MIME media type of the `@Consumes` annotation at the class level. The `doPost2` method overrides the class level `@Consumes` annotation to specify that it can accept URL-encoded form data.

If no resource methods can respond to the requested MIME type, an HTTP 415 error (Unsupported Media Type) is returned to the client.

The HelloWorld example discussed previously in this section can be modified to set the cliched message using `@Consumes`, as shown in the following code example.

```

@POST
@Consumes("text/plain")
public void postClichedMessage(String message) {
    // Store the message
}

```

In this example, the Java method will consume representations identified by the MIME media type `text/plain`. Notice that the resource method returns `void`. This means no representation is returned and response with a status code of HTTP 204 (No Content) will be returned.

## Extracting Request Parameters

Parameters of a resource method may be annotated with parameter-based annotations to extract information from a request. A previous example presented the use of the `@PathParam` parameter to extract a path parameter from the path component of the request URL that matched the path declared in `@Path`. There are six types of parameters you can extract for use in your resource class: query parameters, URI path parameters, form parameters, cookie parameters, header parameters, and matrix parameters.

*Query parameters* are extracted from the request URI query parameters, and are specified by using the `javax.ws.rs.QueryParam` annotation in the method parameter arguments. The following example (from the `sparklines` sample application) demonstrates using `@QueryParam` to extract query parameters from the Query component of the request URL.

```

@Path("smooth")
@GET
public Response smooth(
    @DefaultValue("2") @QueryParam("step") int step,
    @DefaultValue("true") @QueryParam("min-m") boolean hasMin,
    @DefaultValue("true") @QueryParam("max-m") boolean hasMax,

```

```

    @DefaultValue("true") @QueryParam("last-m") boolean hasLast,
    @DefaultValue("blue") @QueryParam("min-color") ColorParam minColor,
    @DefaultValue("green") @QueryParam("max-color") ColorParam maxColor,
    @DefaultValue("red") @QueryParam("last-color") ColorParam lastColor
  ) { ... }

```

If a query parameter "step" exists in the query component of the request URI, then the "step" value will be extracted and parsed as a 32-bit signed integer and assigned to the step method parameter. If "step" does not exist, then a default value of 2, as declared in the @DefaultValue annotation, will be assigned to the step method parameter. If the "step" value cannot be parsed as a 32-bit signed integer, then an HTTP 400 (Client Error) response is returned.

User-defined Java types such as ColorParam may be used. The following code example shows how to implement this.

```

public class ColorParam extends Color {
    public ColorParam(String s) {
        super(getRGB(s));
    }

    private static int getRGB(String s) {
        if (s.charAt(0) == '#') {
            try {
                Color c = Color.decode("0x" + s.substring(1));
                return c.getRGB();
            } catch (NumberFormatException e) {
                throw new WebApplicationException(400);
            }
        } else {
            try {
                Field f = Color.class.getField(s);
                return ((Color)f.get(null)).getRGB();
            } catch (Exception e) {
                throw new WebApplicationException(400);
            }
        }
    }
}

```

@QueryParam and @PathParam can only be used on the following Java types:

- All primitive types except char
- All wrapper classes of primitive types except Character
- Have a constructor that accepts a single String argument
- Any class with the static method named valueOf(String) that accepts a single String argument
- Any class with a constructor that takes a single String as a parameter

- `List<T>`, `Set<T>`, or `SortedSet<T>`, where *T* matches the already listed criteria. Sometimes parameters may contain more than one value for the same name. If this is the case, these types may be used to obtain all values

If `@DefaultValue` is not used in conjunction with `@QueryParam`, and the query parameter is not present in the request, then value will be an empty collection for `List`, `Set`, or `SortedSet`; null for other object types; and the Java-defined default for primitive types.

*URI path parameters* are extracted from the request URI, and the parameter names correspond to the URI path template variable names specified in the `@Path` class-level annotation. URI parameters are specified using the `javax.ws.rs.PathParam` annotation in the method parameter arguments. The following example shows how to use `@Path` variables and the `@PathParam` annotation in a method:

```
@Path("/{userName}")
public class MyResourceBean {
    ...
    @GET
    public String printUserName(@PathParam("userName") String userId) {
        ...
    }
}
```

In the above snippet, the URI path template variable name `userName` is specified as a parameter to the `printUserName` method. The `@PathParam` annotation is set to the variable name `userName`. At runtime, before `printUserName` is called, the value of `userName` is extracted from the URI and cast to a `String`. The resulting `String` is then available to the method as the `userId` variable.

If the URI path template variable cannot be cast to the specified type, the Jersey runtime returns an HTTP 400 Bad Request error to the client. If the `@PathParam` annotation cannot be cast to the specified type, the Jersey runtime returns an HTTP 404 Not Found error to the client.

The `@PathParam` parameter and the other parameter-based annotations, `@MatrixParam`, `@HeaderParam`, `@CookieParam`, and `@FormParam` obey the same rules as `@QueryParam`.

*Cookie parameters* (indicated by decorating the parameter with `javax.ws.rs.CookieParam`) extract information from the cookies declared in cookie-related HTTP headers. *Header parameters* (indicated by decorating the parameter with `javax.ws.rs.HeaderParam`) extracts information from the HTTP headers. *Matrix parameters* (indicated by decorating the parameter with `javax.ws.rs.MatrixParam`) extracts information from URL path segments. These parameters are beyond the scope of this tutorial.

*Form parameters* (indicated by decorating the parameter with `javax.ws.rs.FormParam`) extract information from a request representation that is of the MIME media type `application/x-www-form-urlencoded` and conforms to the encoding specified by HTML

forms, as described [here](#). This parameter is very useful for extracting information that is POSTed by HTML forms. The following example extracts the form parameter named "name" from the POSTed form data.

```
@POST
@Consumes("application/x-www-form-urlencoded")
public void post(@FormParam("name") String name) {
    // Store the message
}
```

If it is necessary to obtain a general map of parameter names to values, use code such as that shown in the following example, for query and path parameters.

```
@GET
public String get(@Context UriInfo ui) {
    MultivaluedMap<String, String> queryParams = ui.getQueryParameters();
    MultivaluedMap<String, String> pathParams = ui.getPathParameters();
}
```

Or code such as the following for header and cookie parameters:

```
@GET
public String get(@Context HttpHeaders hh) {
    MultivaluedMap<String, String> headerParams = ui.getRequestHeaders();
    Map<String, Cookie> pathParams = ui.getCookies();
}
```

In general `@Context` can be used to obtain contextual Java types related to the request or response.

For form parameters it is possible to do the following:

```
@POST
@Consumes("application/x-www-form-urlencoded")
public void post(MultivaluedMap<String, String> formParams) {
    // Store the message
}
```

## Overview of JAX-RS and Jersey: Further Information

The following documents contain information that you might find useful when creating applications using Jersey and JAX-RS.

- [Overview of JAX-RS 1.0 Features](#)

This document contains some of the information from this tutorial, as well as additional topics such as *Representations and Java types*, *Building Responses*, *Sub-resources*, *Building URIs*, *WebApplicationException and mapping Exceptions to Responses*, *Conditional GETs and Returning 304 (Not Modified) Responses*, *Life-cycle of root resource classes*, *Security*, *Rules of Injection*, *Use of @Context*, and *APIs defined by JAX-RS*.

- **Overview of Jersey 1.0 Features**

This document contains the following topics: Deployment, Web-Deployment Using Servlet, Embedded-Web-Deployment Using GlassFish, Embedded-Deployment Using Grizzly, Embedded-Web-Deployment Using Grizzly, Client-Side API, Client-Side Filters, Integration with Spring, JSON, JAXB, Module View Controller with JSPs, Resource Class Life-Cycle, Resource Class Instantiation, Web Application Description Language (WADL) Support, Pluggable Templates for Model View Controller, Server-Side Filters URI utilities, Web Application Reloading, Pluggable Injection, Pluggable Life-Cycle, Pluggable HTTP containers, and Pluggable IoC Integration.

## Example Applications for JAX-RS and Jersey

This section provides an introduction to creating, deploying, and running your own Jersey applications. This section demonstrates the steps that you would take to create, build, deploy, and test a very simple web application that is annotated with Jersey.

Another way that you could learn more about deploying and running Jersey applications is to review the many sample applications that ship with Jersey. These samples are installed into the *as-install/jersey/samples* directory when the Jersey Documentation and Samples are installed onto the Enterprise Server using the Update Tool. The process of installing from the Update Tool is described in “[Java EE 6 Tutorial Component](#)” on page 61. There is a `README.html` file for each sample that describes the sample and describes how to deploy and test the sample. These samples also include a Project Object Model file, `pom.xml`, that is used by Maven to build the project. The sample applications that ship with Jersey require [Maven](#) to run. The sample applications included with the tutorial will run using Ant.

## Creating a RESTful Web Service

This section discusses two ways that you can create a RESTful web service. If you choose to use NetBeans IDE to create a RESTful web service, the IDE generates a skeleton where you simply need to implement the appropriate methods. If you choose not to use an IDE, try using one of the example applications that ship with Jersey as a template to modify.

### ▼ **Creating a RESTful Web Service Using NetBeans IDE**

This section describes, using a very simple example, how to create a Jersey-annotated web application from NetBeans IDE.

**1 In NetBeans IDE, create a simple web application. This example creates a very simple “Hello, World” web application.**

- a. **Open NetBeans IDE.**
- b. **Select File→New Project.**
- c. **From Categories, select Java Web. From Projects, select Web Application. Click Next.**

---

**Note** – For this step, you could also create a RESTful web service in a Maven web project by selecting Maven as the category and Maven Web Project as the project. The remaining steps would be the same.

---

- d. **Enter a project name, HelloWorldApplication, click Next.**
- e. **Make sure the Server is Sun GlassFish v3 (or similar wording.)**
- f. **Click Finish. You may be prompted for your server Administrator User Name and Password. If so, enter this information.**

**2 The project will be created. The file `index.jsp` will display in the Source pane.**

**3 Right-click the project and select New, then select RESTful Web Services from Patterns.**

- a. **Select Simple Root Resource. Click Next.**
- b. **Enter a Resource Package name, like `helloWorld`.**
- c. **Enter `helloWorld` in the Path field. Enter `HelloWorld` in the Class Name field. For MIME Type select `text/html`.**
- d. **Click Finish.**

**e. The REST Resources Configuration page displays. Select OK.**

A new resource, `HelloWorld.java`, is added to the project and displays in the Source pane. This file provides a template for creating a RESTful web service.

**4 In `HelloWorld.java`, find the `getHtml()` method. Replace the `//TODO` comment and the exception with the following text, so that the finished product resembles the following method.**

---

**Note** – Because the MIME type that is produced is HTML, you can use HTML tags in your return statement.

---

```

/**
 * Retrieves representation of an instance of helloWorld.HelloWorld
 * @return an instance of java.lang.String
 */
@GET
@Produces("text/html")
public String getHtml() {
    return "<html><body><h1>Hello, World!!</body></h1></html>";
}

```

**5 Test the web service. To do this, right-click the project node and click Test RESTful Web Services.**

This step will deploy the application and bring up a test client in the browser.

**6 When the test client displays, select the `helloworld` resource in the left pane, and click the Test button in the right pane.**

The words `Hello, World!!` will display in the Response window below.

**7 Deploy and Run the application.**

- a. Set the Run Properties. To do this, right-click the project node, select Properties, and then select the Run category. Set the Relative URL to the location of the RESTful web service relative to the Context Path, which for this example is `resources/helloworld`.**

---

**Tip** – You can find the value for the Relative URL in the Test RESTful Web Services browser window. In the top of the right pane, after Resource, is the URL for the RESTful web service being tested. The part following the Context Path (`http://localhost:8080/HelloWorldApp`) is the Relative URL that needs to be entered here.

If you don't set this property, by default the file `index.jsp` will display when the application is run. As this file also contains `Hello World` as its default value, you might not notice that your RESTful web service isn't running, so just be aware of this default and the need to set this property, or update `index.jsp` to provide a link to the RESTful web service.

---

- b. Right-click the project and select Deploy.**

- c. Right-click the project and select Run.**

A browser window opens and displays the return value of `Hello, World!!`

**See Also** For other sample applications that demonstrate deploying and running Jersey applications using NetBeans, read “[Example: Creating a Simple Hello World Application Using JAX-RS and Jersey](#)” on page 270 and/or look at the tutorials on the NetBeans tutorial site, such as the one

titled [Getting Started with RESTful Web Services](#). This tutorial includes a section on creating a CRUD application from a database. Create, read, update and delete (CRUD) are the four basic functions of persistent storage and relational databases.

## ▼ Creating a RESTful Web Service From Examples

The easiest way to create and run an application without NetBeans IDE is to copy and edit one of the Jersey sample applications. These samples are installed into the `as-install/jersey/samples` directory when the Jersey Documentation and Samples are installed onto the Enterprise Server using the Update Tool. The process of installing from the Update Tool is described in “[Java EE 6 Tutorial Component](#)” on page 61. This task uses the simplest sample application, `helloworld-webapp`, to demonstrate one way you could go about creating your own application without NetBeans IDE.

**Before You Begin** Before you can deploy the Jersey sample applications to GlassFish from the command line, you must have downloaded and installed Maven onto your system. You can install Maven from the Maven website at <http://maven.apache.org>.

- 1 **Copy the `helloworld-webapp` application to a new directory named `helloworld2`. You will find this application in the directory `as-install/glassfish/jersey/samples/helloworld-webapp`.**
- 2 **Do a search for all *directories* named `helloworld-webapp` and rename them to `helloworld2`.**
- 3 **Search again for all *files* containing the text `helloworld-webapp` and edit these files to replace this text with `helloworld2`.**
- 4 **Using a text editor, open the file `jersey/samples/helloworld2/src/main/java/com/sun/jersey/samples/helloworld/resources/HelloWorldResource.java`.**
- 5 **Modify the text that is returned by the resource to `Hello World 2`. Save and close the file.**
- 6 **Use Maven to compile and deploy the application. For this sample application, it is deployed onto Grizzly. Enter the following command from the command line in the directory `jersey/samples/helloworld2` to compile and deploy the application: `mvn glassfish:run`.**
- 7 **Open a web browser, and enter the URL to which the application was deployed, which in this examples is `http://localhost:8080/helloworld2/helloworld`. `Hello World 2` will display in the browser.**

**See Also** You can learn more about deploying and running Jersey applications by reviewing the many sample applications that ship with Jersey. There is a `README.html` file for each sample that describes the sample and describes how to deploy and test the sample, and there is a Project Object Model file, `pom.xml`, that is used by Maven to build the project. Find a project that is similar to one you are hoping to create and use it as a template to get you started.

An example that starts from scratch can be found [here](#).

For questions regarding Jersey sample applications, visit the [Jersey Community Wiki page](#), or send an email to the users mailing list, [users@jersey.dev.java.net](mailto:users@jersey.dev.java.net).

## ▼ Creating a RESTful Web Service From Maven Archetype

Although this tutorial does not present instructions on using Maven for creating applications as a general rule, because Project Jersey is built, assembled and installed using Maven, and all of its sample applications are Maven-based, this section provides an example that creates a skeleton Jersey application from a Maven archetype.

**Before You Begin** This example requires that Maven be installed and configured to run from the command line on your system. Maven can be downloaded from <http://maven.apache.org/>.

**1 Start the Enterprise Server. For instructions on how to do this, read “Starting and Stopping the Enterprise Server” on page 63.**

**2 After Maven is installed, run the following from the command line:**

```
mvn archetype:generate -DarchetypeCatalog=http://download.java.net/maven/2
```

The archetype catalog will download. You will be prompted to select the type of archetype you want to create. As of the publication date of this tutorial, the following choices display in the command window. These options are likely to change, but are provided here to give you an idea of what they might look like.

Choose archetype:

```
1: http://download.java.net/maven/2/archetype-catalog.xml ->
jersey-quickstart-grizzly (Archetype for creating a RESTful web
application with Jersey and Grizzly)
```

```
2: http://download.java.net/maven/2/archetype-catalog.xml ->
jersey-quickstart-webapp (Archetype for creating a Jersey based RESTful
web application with WAR packaging)
```

```
3: http://download.java.net/maven/2/archetype-catalog.xml ->
jersey-quickstart-ejb (Archetype for creating a Jersey based RESTful EJB
application with WAR packaging)
```

```
4: http://download.java.net/maven/2/archetype-catalog.xml ->
jsf2-simple-example-archetype (Simple JSF project with no non-JavaEE
dependencies)
```

**3 Select the appropriate option for the type of RESTful web service you would like to create.**

With the Grizzly-based archetype (selection 1), you will get a sample Java application, which you can run directly from Java without a need to deploy it to any container. The web application archetype (selection 2) enables you to build a WAR archive, which you could deploy onto any web Servlet container.

**4 Define a value for `groupId`, such as `REStHello`.****5 Define a value for `artifactId`, such as `REStHelloApp`. This is the name of the web application as well as the directory in which the application is created.****6 Define value for `version`: `1.0-SNAPSHOT`. You can accept the default by not entering anything.****7 Define value for `package`: `groupId`, such as `restHello`. This is the directory where the main Java files will be located, which is `basedir/artifactId/src/main/java/package`. If you used the example entries, this directory will be `REStHelloApp/src/main/java/restHello`.****8 Confirm properties configuration. Enter `Y` to confirm or `N` to cancel.**

Maven generates a new project containing a simple Hello World RESTful web service.

**9 Build and run your RESTful web service. First, change into the project directory, which is the `artifactId`, or `REStHelloApp` if you used the example text.**

- For the Grizzly-based scenario (selection 1), build and run the web service on the Grizzly container using this command: `mvn clean compile exec:java`.
- If you selected the WAR-based scenario (selection 2), build your WAR file using the command `mvn clean package`. Deploy the WAR file to your favorite Servlet container. To run it using the embedded version of GlassFish V3, use this command: `mvn glassfish:run`.

**10 Test the service in your browser.**

- Enter the following URL to run the Grizzly-based application (selection 1):  
`http://localhost:9998/myresource`. This is the location where it is published by default. The browser displays the text `Got it!`
- Enter the following URL to run the WAR-based scenario (selection 2):  
`http://localhost:8080/artifactId/webresources/myresource`. If you used the example entries, enter `http://localhost:8080/REStHelloApp/webresources/myresource`. This is the location where it is published by default. The browser displays the text `Hi there!`

After starting the application using Grizzly, you should see output that looks similar to that in the example output, below:

```
[INFO] -----  
[INFO] Building REStHelloApp
```

```
[INFO] task-segment: [clean, compile, exec:java]
[INFO] -----
[INFO] [clean:clean]
[INFO] [resources:resources]
[INFO] Using default encoding to copy filtered resources.
[INFO] [compiler:compile]
[INFO] Compiling 2 source files to /export/home/japod/test/jaxrs-tutorial/
      RESTHelloApp/target/classes
[INFO] Preparing exec:java
[INFO] No goals needed for project - skipping
[INFO] [exec:java]
Starting grizzly...
Jersey app started with WADL available at http://localhost:9998/application.wadl
Hit enter to stop it...
```

After starting the application for the WAR-based scenario, you should see output that looks similar to that in the example output, below:

```
[INFO] Building RESTHelloApp Jersey Webapp
[INFO] task-segment: [glassfish:run]
[INFO] -----
[INFO] Preparing glassfish:run
[INFO] [resources:resources]
[INFO] Using default encoding to copy filtered resources.
[INFO] snapshot org.glassfish:glassfish-parent:10.0-SNAPSHOT: checking for
      updates from glassfish-maven2-repository.dev.java.net
[INFO] [compiler:compile]
[INFO] Compiling 1 source file to /export/home/japod/test/jaxrs-tutorial/
      RESTHelloApp/target/classes
[INFO] [glassfish:run]
Dec 8, 2009 1:20:34 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: HK2 initialized in 479 ms
Dec 8, 2009 1:20:34 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: com.sun.enterprise.naming.impl.ServicesHookup@1342545 Init done in 486 ms
Dec 8, 2009 1:20:34 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: com.sun.enterprise.v3.server.Globals@6de609 Init done in 488 ms
Dec 8, 2009 1:20:34 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: com.sun.enterprise.v3.server.SystemTasks@e7e8eb Init done in 493 ms
Dec 8, 2009 1:20:34 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: com.sun.enterprise.v3.services.impl.HouseKeeper@1a6518 Init done in 503 ms
Dec 8, 2009 1:20:34 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: com.sun.enterprise.v3.services.impl.CmdLineParamProcessor@8390b0
      Init done in 506 ms
JMXMP connector server URL = service:jmx:jmxmp://localhost:8888
Dec 8, 2009 1:20:35 AM com.sun.enterprise.v3.services.impl.GrizzlyProxy start
INFO: Listening on port 8080
Dec 8, 2009 1:20:35 AM com.sun.enterprise.v3.server.AppServerStartup run
INFO: com.sun.enterprise.v3.services.impl.GrizzlyService@59cbda startup
      done in 815 ms
```

```
Dec 8, 2009 1:20:35 AM com.sun.enterprise.v3.services.impl.  
    ApplicationLoaderService postConstruct  
INFO: loader service postConstruct started at 1260231635181  
Dec 8, 2009 1:20:35 AM com.sun.enterprise.v3.server.AppServerStartup run  
INFO: Application Loader startup done in 883 ms  
Dec 8, 2009 1:20:35 AM com.sun.enterprise.v3.server.AppServerStartup run  
INFO: Glassfish v3 started in 883 ms  
Dec 8, 2009 1:20:38 AM com.sun.enterprise.web.WebModuleContextConfig  
    authenticatorConfig  
SEVERE: webModuleContextConfig.missingRealm  
Dec 8, 2009 1:20:38 AM com.sun.jersey.api.core.PackagesResourceConfig init  
INFO: Scanning for root resource and provider classes in the packages:  
    restHello  
Dec 8, 2009 1:20:38 AM com.sun.jersey.api.core.PackagesResourceConfig init  
INFO: Root resource classes found:  
    class restHello.MyResource  
Dec 8, 2009 1:20:38 AM com.sun.jersey.api.core.PackagesResourceConfig init  
INFO: Provider classes found:  
Dec 8, 2009 1:20:38 AM com.sun.jersey.server.impl.application.  
    WebApplicationImpl initiate  
INFO: Initiating Jersey application, version 'Jersey: 1.1.4.1 11/24/2009 01:30 AM'  
Hit ENTER for redeploy
```

## Example: Creating a Simple Hello World Application Using JAX-RS and Jersey

This section discusses the simple RESTful web service that is included with the tutorial examples in the directory `jaxrs/JAXRSHelloWorld`. This example was created by following the steps similar to those described in [“Creating a RESTful Web Service Using NetBeans IDE” on page 263](#).

### JAXRSHelloWorld Example: Discussion

With this simple application, a simple root resource for a RESTful web service was selected. This generates a RESTful root resource class with GET and PUT methods. This design is useful for creating examples such as this simple Hello World service.

In this example, the method `getHtml()` is annotated with `@GET` and the `@Produces("text/html")` annotation. This method will process HTTP GET requests and produce content in HTML. To finish this example, you simply need to replace the current contents of this example with a statement that returns `Hello World`. This example has also replaced the name of the method with the name `sayHello`. Here is the code for the completed `sayHello()` method:

```
@GET
@Produces("text/html")
public String sayHello() {
    return "Hello World";
}
```

## ▼ Testing the JAXRSHelloWorld Example

- 1 Open the project *javaeetutorial/jaxrs/JAXRSHelloWorld* in NetBeans IDE.
- 2 Right-click the project node, JAXRSHelloWorld, and select Test RESTful Web Services.
- 3 Click the `helloWorld` service in the left pane.
- 4 The `Get(text/html)` method is selected by default. Click Test.
- 5 The response `Hello World`, displays in the lower pane, as shown in the following figure.

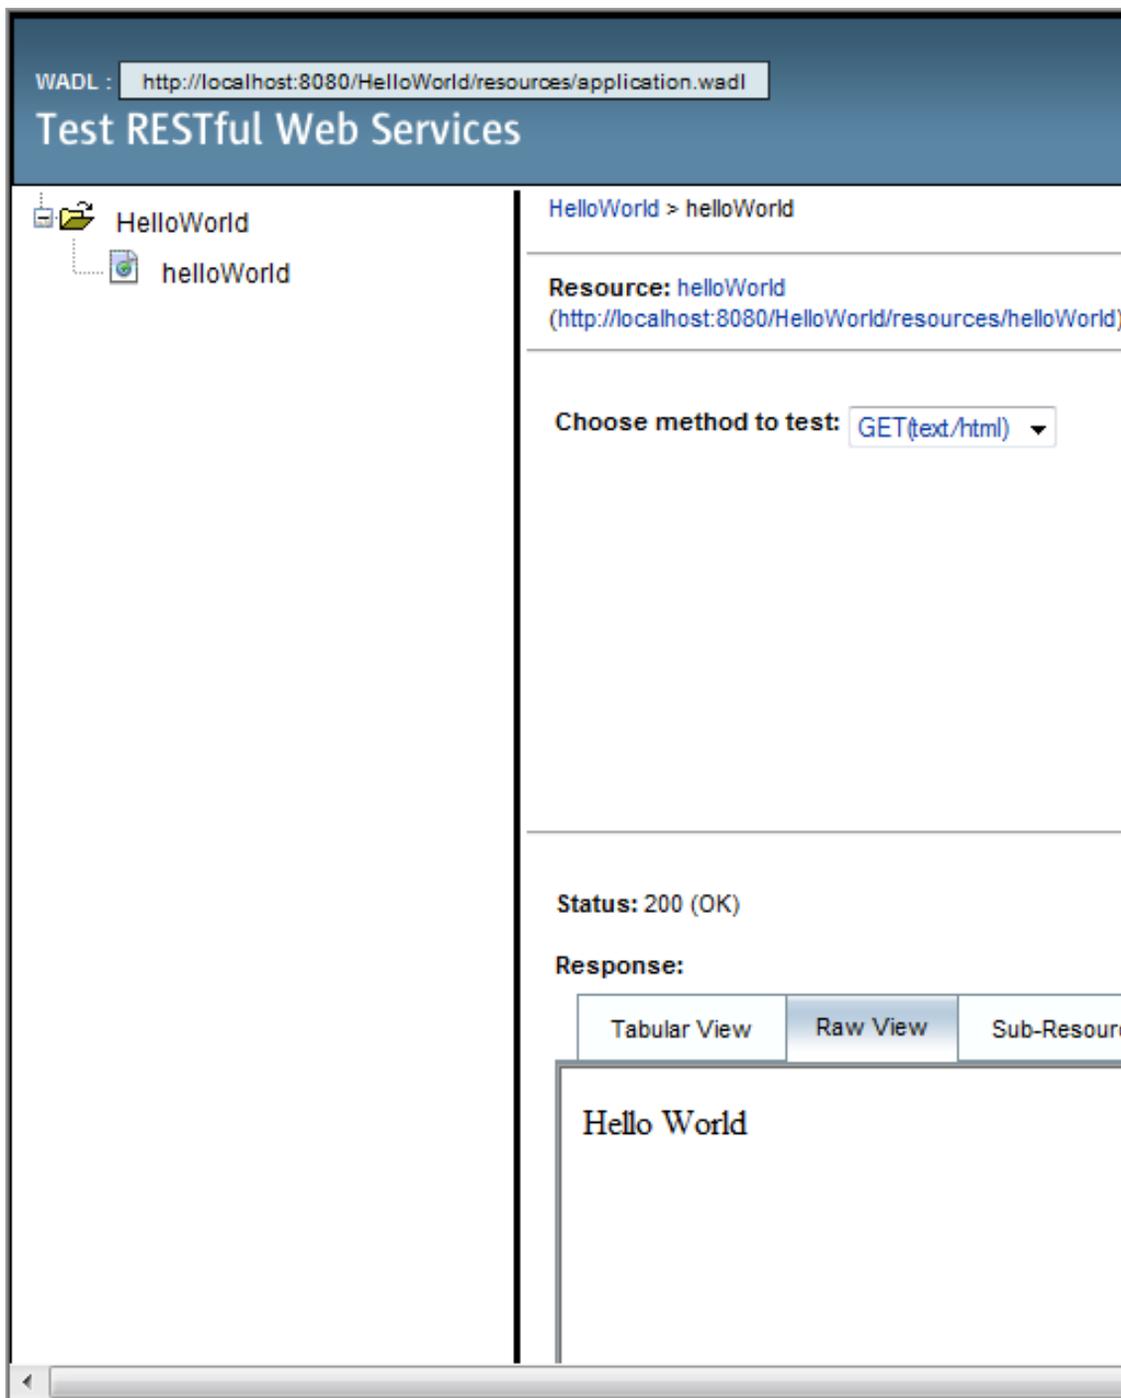


FIGURE 14-1 Testing JAXRSHelloWorld Web Service  
The Java EE 6 Tutorial, Volume II • December 2009

## ▼ Deploying and Running the JAXRSHelloWorld Example

**Before You Begin** The application's Run properties must be set to run the RESTful web service. For the provided application, this task has been completed. For future reference, right-click the project node, select Properties, then select Run, and enter the Relative URL. For this example, you would enter `/resources/helloWorld`.

- 1 **Right-click the project node, JAXRSHelloWorld, and select Deploy.**
- 2 **Right-click the project node, JAXRSHelloWorld, and select Run.**
- 3 **A browser opens and displays Hello World at the URL**  
`http://localhost:8080/HelloWorld/resources/helloWorld`.  
 The browser displays the text "Hello World".

## Example: Adding on to the Simple Hello World RESTful Web Service

This section discusses the simple RESTful web service that is included with the tutorial examples in the directory `jaxrs/HelloWorld3`. This example was created by following the steps similar to those described in "[Creating a RESTful Web Service Using NetBeans IDE](#)" on [page 263](#).

### HelloWorld3 Example: Discussion

This example takes the simple Hello World application discussed in the previous section and adds to it. In this example, there are methods for getting a user's name, and then the name is appended to the Hello World greeting. An annotation that wasn't used in the previous example, `@QueryParam`, is used in this example.

In this example, there is a simple RESTful web service that returns HTML messages. To accomplish this task, you would first create a class that uses Java Architecture for XML Binding (JAXB). This class represents the HTML message in Java (`RESTGreeting.java`), then creates a RESTful web service that returns an HTML message (`HelloGreetingService.java`.)

The JAXB class that represents the HTML message gets the message and the name. This file, `RESTGreeting.java`, is basic Java code that creates a new instance of `RESTGreeting` and the getter and setter methods for its parameters.

The RESTful web service that returns an HTML message is in the file `HelloGreetingService.java`. You may notice that method that is annotated with JAX-RS annotations is similar to the one described in the previous example, however, this example adds an `@QueryParam` annotation to extract query parameters from the Query component of the request URL. The following code example shows the JAX-RS-annotated method:

```
@GET
@Produces("text/html")
public RESTGreeting getHtml(@QueryParam("name")String name) {
    return new RESTGreeting(getGreeting(), name);
}
private String getGreeting(){
    return "Hello ";
}
```

## ▼ Testing the HelloWorld3 Example

- 1 **Open the project `javaeetutorial/jaxrs/HelloWorld3` in NetBeans IDE.**
- 2 **Right-click the project node, `HelloWorld3`, and select `Test RESTful Web Services`.**
- 3 **Click the `helloGreeting` service in the left pane.**
- 4 **Enter a name in the `name` text field.**
- 5 **The `Get (text/html)` method is selected by default. Click `Test`.**
- 6 **The response `Hello name`, displays in the Response pane, under the `Raw View` tab.**

## ▼ Deploying and Running the HelloWorld3 Example

**Before You Begin** The application's Run properties must be set to run the RESTful web service. For the provided application, this task has been completed. For future reference, right-click the project node, select `Properties`, then select `Run`, and enter the Relative URL. For this example, you would enter `/helloGreeting`.

- 1 **Right-click the project node, `HelloWorld3`, and select `Deploy`.**
- 2 **Right-click the project node, `HelloWorld3`, and select `Run`.**

The Run property does not specify a particular name, so none is shown in the browser window when it displays. The browser window simply shows the message `Hello`.
- 3 **Append a name to the URL in the web browser, so that the URL looks like this:**  
`http://localhost:8080/HelloWorld3/helloGreeting?name=your_name`.
- 4 **The message `Hello` and the name `your_name` display in the browser.**

## JAX-RS in the First Cup Example

JAX-RS is used in the *Your First Cup of Java* example, which you will find at *Your First Cup: An Introduction to the Java EE Platform*

## Real World Examples

A few real-world web applications that use RESTful web services include most blog sites. These are considered RESTful in that most blog sites involve downloading XML files in RSS or Atom format which contain lists of links to other resources. Other web sites and web applications that use REST-like developer interfaces to data include Twitter and Amazon S3 (Simple Storage Service). With Amazon S3, buckets and objects can be created, listed, and retrieved using either a REST-style HTTP interface or a SOAP interface. The examples that ship with Jersey include a storage service example with a RESTful interface. The tutorial at <http://netbeans.org/kb/docs/websvc/twitter-swing.html> uses the NetBeans IDE to create a simple, graphical, REST-based client that displays Twitter public time line messages and lets you view and update your Twitter status.

## Further Information

The information in this tutorial focuses on learning about JAX-RS and Jersey. If you are interested in learning more about RESTful Web Services in general, here are a few links to get you started.

- The Community Wiki for Project Jersey has loads of information on all things RESTful. You'll find it at <http://wikis.sun.com/display/Jersey/Main>.
- *Fielding Dissertation: Chapter 5: Representational State Transfer (REST)*, at [http://www.ics.uci.edu/~fielding/pubs/dissertation/rest\\_arch\\_style.htm](http://www.ics.uci.edu/~fielding/pubs/dissertation/rest_arch_style.htm).
- *Representational State Transfer*, from Wikipedia, [http://en.wikipedia.org/wiki/Representational\\_State\\_Transfer](http://en.wikipedia.org/wiki/Representational_State_Transfer).
- *RESTful Web Services*, by Leonard Richardson and Sam Ruby. Available from O'Reilly Media at <http://oreilly.com/catalog/9780596529260/>.

Some of the Jersey team members discuss topics out of the scope of this tutorial on their blogs. A few are listed below:

- Earthly Powers, by Paul Sandoz, at <http://blogs.sun.com/sandoz/category/REST>.
- Marc Hadley's Blog, at <http://weblogs.java.net/blog/mhadley/>
- Japod's Blog, by Jakub Podlesak, at <http://blogs.sun.com/japod/category/REST>.

You can always get the latest technology and information by visiting the Java Developer's Network. The links are listed below:

- Get the latest on JSR-311, the Java API's for RESTful Web Services (JAX-RS), at <https://jsr311.dev.java.net/>.
- Get the latest on Jersey, the open source JAX-RS reference implementation, at <https://jersey.dev.java.net/>.

◆ ◆ ◆ CHAPTER 15

# Advanced Topics in JAX-RS

---

This chapter introduces more advanced material for JAX-RS.

## **Advanced Topics in JAX-RS**

Add more information on advanced topics in JAX-RS here.



**PART IV**

**Enterprise Beans**

Part Four explores Enterprise JavaBeans.



# Enterprise Beans

---

Enterprise beans are Java EE components that implement Enterprise JavaBeans (EJB) technology. Enterprise beans run in the EJB container, a runtime environment within the Enterprise Server (see “Container Types” on page 42). Although transparent to the application developer, the EJB container provides system-level services such as transactions and security to its enterprise beans. These services enable you to quickly build and deploy enterprise beans, which form the core of transactional Java EE applications.

## What Is an Enterprise Bean?

Written in the Java programming language, an *enterprise bean* is a server-side component that encapsulates the business logic of an application. The business logic is the code that fulfills the purpose of the application. In an inventory control application, for example, the enterprise beans might implement the business logic in methods called `checkInventoryLevel` and `orderProduct`. By invoking these methods, clients can access the inventory services provided by the application.

## Benefits of Enterprise Beans

For several reasons, enterprise beans simplify the development of large, distributed applications. First, because the EJB container provides system-level services to enterprise beans, the bean developer can concentrate on solving business problems. The EJB container, rather than the bean developer, is responsible for system-level services such as transaction management and security authorization.

Second, because the beans rather than the clients contain the application’s business logic, the client developer can focus on the presentation of the client. The client developer does not have to code the routines that implement business rules or access databases. As a result, the clients are thinner, a benefit that is particularly important for clients that run on small devices.

Third, because enterprise beans are portable components, the application assembler can build new applications from existing beans. These applications can run on any compliant Java EE server provided that they use the standard APIs.

## When to Use Enterprise Beans

You should consider using enterprise beans if your application has any of the following requirements:

- The application must be scalable. To accommodate a growing number of users, you may need to distribute an application's components across multiple machines. Not only can the enterprise beans of an application run on different machines, but also their location will remain transparent to the clients.
- Transactions must ensure data integrity. Enterprise beans support transactions, the mechanisms that manage the concurrent access of shared objects.
- The application will have a variety of clients. With only a few lines of code, remote clients can easily locate enterprise beans. These clients can be thin, various, and numerous.

## Types of Enterprise Beans

Table 16-1 summarizes the two types of enterprise beans. The following sections discuss each type in more detail.

TABLE 16-1 Enterprise Bean Types

Enterprise Bean Type	Purpose
Session	Performs a task for a client; optionally may implement a web service
Message-Driven	Acts as a listener for a particular messaging type, such as the Java Message Service API

---

**Note** – Entity beans have been replaced by Java Persistence API entities. For information about entities, see [Chapter 22, “Introduction to the Java Persistence API.”](#)

---

# What Is a Session Bean?

A *session bean* encapsulates business logic that can be invoked programmatically by a client over local, remote, or web service client views. To access an application that is deployed on the server, the client invokes the session bean's methods. The session bean performs work for its client, shielding the client from complexity by executing business tasks inside the server.

A session bean is not persistent. (That is, its data is not saved to a database.)

For code samples, see [Chapter 18, “Running the Enterprise Bean Examples.”](#)

## Types of Session Beans

There are three types of session beans: stateful, stateless, and singleton.

### Stateful Session Beans

The state of an object consists of the values of its instance variables. In a *stateful* session bean, the instance variables represent the state of a unique client-bean session. Because the client interacts (“talks”) with its bean, this state is often called the *conversational state*.

As its name suggests, a session bean is similar to an interactive session. A session bean is not shared; it can have only one client, in the same way that an interactive session can have only one user. When the client terminates, its session bean appears to terminate and is no longer associated with the client.

The state is retained for the duration of the client-bean session. If the client removes the bean, the session ends and the state disappears. This transient nature of the state is not a problem, however, because when the conversation between the client and the bean ends there is no need to retain the state.

### Stateless Session Beans

A *stateless* session bean does not maintain a conversational state with the client. When a client invokes the methods of a stateless bean, the bean's instance variables may contain a state specific to that client, but only for the duration of the invocation. When the method is finished, the client-specific state should not be retained. Clients may, however, change the state of instance variables in pooled stateless beans, and this state is held over to the next invocation of the pooled stateless bean. Except during method invocation, all instances of a stateless bean are equivalent, allowing the EJB container to assign an instance to any client. That is, the state of a stateless session bean should apply across all clients.

Because stateless session beans can support multiple clients, they can offer better scalability for applications that require large numbers of clients. Typically, an application requires fewer stateless session beans than stateful session beans to support the same number of clients.

A stateless session bean can implement a web service, but a stateful session bean cannot.

## Singleton Session Beans

A *singleton* session bean is instantiated once per application, and exists for the lifecycle of the application. Singleton session beans are designed for circumstances where a single enterprise bean instance is shared across and concurrently accessed by clients.

Singleton session beans offer similar functionality to stateless session beans, but differ from stateless session beans in that there is only one singleton session bean per application, as opposed to a pool of stateless session beans, any of which may respond to a client request. Like stateless session beans, singleton session beans can implement web service endpoints.

Singleton session beans maintain their state between client invocations, but are not required to maintain their state across server crashes or shutdowns.

Applications that use a singleton session bean may specify that the singleton should be instantiated upon application startup, which allows the singleton to perform initialization tasks for the application. The singleton may perform cleanup tasks on application shutdown as well, because the singleton will operate throughout the lifecycle of the application.

## When to Use Session Beans

Stateful session beans are appropriate if any of the following conditions are true:

- The bean's state represents the interaction between the bean and a specific client.
- The bean needs to hold information about the client across method invocations.
- The bean mediates between the client and the other components of the application, presenting a simplified view to the client.
- Behind the scenes, the bean manages the work flow of several enterprise beans.

To improve performance, you might choose a stateless session bean if it has any of these traits:

- The bean's state has no data for a specific client.
- In a single method invocation, the bean performs a generic task for all clients. For example, you might use a stateless session bean to send an email that confirms an online order.
- The bean implements a web service.

Singleton session beans are appropriate in the following circumstances:

- State needs to be shared across the application.
- A single enterprise bean needs to be accessed by multiple threads concurrently.
- The application needs an enterprise bean to perform tasks upon application startup and shutdown.

- The bean implements a web service.

## What Is a Message-Driven Bean?

A *message-driven bean* is an enterprise bean that allows Java EE applications to process messages asynchronously. It normally acts as a JMS message listener, which is similar to an event listener except that it receives JMS messages instead of events. The messages can be sent by any Java EE component (an application client, another enterprise bean, or a web component) or by a JMS application or system that does not use Java EE technology. Message-driven beans can process JMS messages or other kinds of messages.

## What Makes Message-Driven Beans Different from Session Beans?

The most visible difference between message-driven beans and session beans is that clients do not access message-driven beans through interfaces. Interfaces are described in the section “[Accessing Enterprise Beans](#)” on page 286. Unlike a session bean, a message-driven bean has only a bean class.

In several respects, a message-driven bean resembles a stateless session bean.

- A message-driven bean’s instances retain no data or conversational state for a specific client.
- All instances of a message-driven bean are equivalent, allowing the EJB container to assign a message to any message-driven bean instance. The container can pool these instances to allow streams of messages to be processed concurrently.
- A single message-driven bean can process messages from multiple clients.

The instance variables of the message-driven bean instance can contain some state across the handling of client messages (for example, a JMS API connection, an open database connection, or an object reference to an enterprise bean object).

Client components do not locate message-driven beans and invoke methods directly on them. Instead, a client accesses a message-driven bean through, for example, JMS by sending messages to the message destination for which the message-driven bean class is the `MessageListener`. You assign a message-driven bean’s destination during deployment by using Enterprise Server resources.

Message-driven beans have the following characteristics:

- They execute upon receipt of a single client message.
- They are invoked asynchronously.
- They are relatively short-lived.

- They do not represent directly shared data in the database, but they can access and update this data.
- They can be transaction-aware.
- They are stateless.

When a message arrives, the container calls the message-driven bean's `onMessage` method to process the message. The `onMessage` method normally casts the message to one of the five JMS message types and handles it in accordance with the application's business logic. The `onMessage` method can call helper methods, or it can invoke a session bean to process the information in the message or to store it in a database.

A message can be delivered to a message-driven bean within a transaction context, so all operations within the `onMessage` method are part of a single transaction. If message processing is rolled back, the message will be redelivered. For more information, see [Chapter 31, "Transactions."](#)

## When to Use Message-Driven Beans

Session beans allow you to send JMS messages and to receive them synchronously, but not asynchronously. To avoid tying up server resources, do not use blocking synchronous receives in a server-side component, and in general JMS messages should not be sent or received synchronously. To receive messages asynchronously, use a message-driven bean.

## Accessing Enterprise Beans

---

**Note** – The material in this section applies only to session beans and not to message-driven beans. Because they have a different programming model, message-driven beans do not have interfaces or no-interface views that define client access.

---

Clients access enterprise beans either through a *no-interface view* or through a *business interface*. A no-interface view of an enterprise bean exposes the public methods of the enterprise bean implementation class to clients. Clients using the no-interface view of an enterprise bean may invoke any public methods in the enterprise bean implementation class, or any superclasses of the implementation class. A business interface is a standard Java programming language interface that contains the business methods of the enterprise bean.

A client can access a session bean only through the methods defined in the bean's business interface, or through the public methods of an enterprise bean that has a no-interface view. The business interface or no-interface view defines the client's view of an enterprise bean. All other aspects of the enterprise bean (method implementations and deployment settings) are hidden from the client.

Well-designed interfaces and no-interface views simplify the development and maintenance of Java EE applications. Not only do clean interfaces and no-interface views shield the clients from any complexities in the EJB tier, but they also allow the enterprise beans to change internally without affecting the clients. For example, if you change the implementation of a session bean business method, you won't have to alter the client code. But if you were to change the method definitions in the interfaces, then you might have to modify the client code as well. Therefore, it is important that you design the interfaces and no-interface views carefully to isolate your clients from possible changes in the enterprise beans.

Session beans can have more than one business interface. Session beans should, but are not required to, implement their business interface or interfaces.

## Using Enterprise Beans in Clients

The client of an enterprise bean obtains a reference to an instance of an enterprise bean either through *dependency injection*, using Java programming language annotations, or *JNDI lookup*, using the Java Naming and Directory Interface syntax to find the enterprise bean instance.

Dependency injection is the simplest way of obtaining an enterprise bean reference. Clients that run within a Java EE server-managed environment, like JSF web applications, JAX-RS web services, other enterprise beans, or Java EE application clients support dependency injection using the `javax.ejb.EJB` annotation.

Applications that run outside a Java EE server-managed environment, such as Java SE applications, must perform an explicit lookup. JNDI supports a global syntax for identifying Java EE components to simplify this explicit lookup.

### Portable JNDI Syntax

There are three JNDI namespaces used for portable JNDI lookups: `java:global`, `java:module`, and `java:app`.

The `java:global` JNDI namespace is the portable way of finding remote enterprise beans using JNDI lookups. JNDI addresses are of the following form:

```
java:global[/application name]/module name/enterprise bean name[/interface name]
```

Application name and module name default to the name of the application and module minus the file extension. Application names are only required if the application is packaged within an EAR. The interface name is only required if the enterprise bean implements more than one business interface.

The `java:module` namespace is used to lookup local enterprise beans within the same module. JNDI addresses using the `java:module` namespace are of the following form:

```
java:module/enterprise bean name/[interface name]
```

The interface name is only required if the enterprise bean implements more than one business interface.

The `java:app` namespace is used to lookup local enterprise beans packaged within the same application. That is, the enterprise bean is packaged within an EAR file containing multiple Java EE modules. JNDI addresses using the `java:app` namespace are of the following form:

```
java:app[/module name]/enterprise bean name[/interface name]
```

The module name is optional. The interface name is only required if the enterprise bean implements more than one business interface.

#### EXAMPLE 16-1 JNDI Address of an Enterprise Bean Packaged Within a WAR File

If an enterprise bean, `MyBean`, is packaged in within the web application archive `myApp.war`, the module name is `myApp`. The portable JNDI name is:

```
java:module/MyBean
```

An equivalent JNDI name using the `java:global` namespace is:

```
java:global/myApp/MyBean
```

## Deciding on Remote or Local Access

When you design a Java EE application, one of the first decisions you make is the type of client access allowed by the enterprise beans: remote, local, or web service.

Whether to allow local or remote access depends on the following factors.

- **Tight or loose coupling of related beans:** Tightly coupled beans depend on one another. For example, if a session bean that processes orders calls a session bean that emails a confirmation message to the customer, these beans are tightly coupled. Tightly coupled beans are good candidates for local access. Because they fit together as a logical unit, they typically call each other often and would benefit from the increased performance that is possible with local access.
- **Type of client:** If an enterprise bean is accessed by application clients, then it should allow remote access. In a production environment, these clients almost always run on different machines than the Enterprise Server. If an enterprise bean's clients are web components or other enterprise beans, then the type of access depends on how you want to distribute your components.
- **Component distribution:** Java EE applications are scalable because their server-side components can be distributed across multiple machines. In a distributed application, for example, the web components may run on a different server than do the enterprise beans they access. In this distributed scenario, the enterprise beans should allow remote access.

- **Performance:** Due to factors such as network latency, remote calls may be slower than local calls. On the other hand, if you distribute components among different servers, you may improve the application's overall performance. Both of these statements are generalizations; actual performance can vary in different operational environments. Nevertheless, you should keep in mind how your application design might affect performance.

If you aren't sure which type of access an enterprise bean should have, choose remote access. This decision gives you more flexibility. In the future you can distribute your components to accommodate the growing demands on your application.

Although it is uncommon, it is possible for an enterprise bean to allow both remote and local access. If this is the case, either the business interface of the bean must be explicitly designated as a business interface by being decorated with the `@Remote` or `@Local` annotations, or the bean class must explicitly designate the business interfaces by using the `@Remote` and `@Local` annotations. The same business interface cannot be both a local and remote business interface.

## Local Clients

A local client has these characteristics:

- It must run in the same application as the enterprise bean it accesses.
- It can be a web component or another enterprise bean.
- To the local client, the location of the enterprise bean it accesses is not transparent.

The no-interface view of an enterprise bean is a local view. The public methods of the enterprise bean implementation class are exposed to local clients that access the no-interface view of the enterprise bean. Enterprise beans that use the no-interface view do not implement a business interface.

The *local business interface* defines the bean's business and lifecycle methods. If the bean's business interface is not decorated with `@Local` or `@Remote`, and the bean class does not specify the interface using `@Local` or `@Remote`, the business interface is by default a local interface.

To build an enterprise bean that allows only local access, you may, but are not required to do one of the following:

- Create an enterprise bean implementation class that does not implement a business interface, indicating that the bean exposes a no-interface view to clients. For example:
 

```
@Session
public class MyBean { ... }
```
- Annotate the business interface of the enterprise bean as a `@Local` interface. For example:
 

```
@Local
public interface InterfaceName { ... }
```
- Specify the interface by decorating the bean class with `@Local` and specify the interface name. For example:

```
@Local(InterfaceName.class)
public class BeanName implements InterfaceName { ... }
```

## Accessing Local Enterprise Beans Using the No-Interface View

Client access to an enterprise bean that exposes a local, no-interface view is accomplished either through dependency injection or JNDI lookup.

Clients *do not* use the new operator to obtain a new instance of an enterprise bean that uses a no-interface view.

### EXAMPLE 16-2 Injecting an Enterprise Bean Using the No-Interface View

To obtain a reference to the no-interface view of an enterprise bean through dependency injection, use the `javax.ejb.EJB` annotation and specify the enterprise bean's implementation class.

```
@EJB
ExampleBean exampleBean;
```

### EXAMPLE 16-3 Looking Up an Enterprise Bean Using the No-Interface View

To obtain a reference to the no-interface view of an enterprise bean using JNDI lookup, use the `javax.naming.InitialContext` interface's `lookup` method.

```
ExampleBean exampleBean = (ExampleBean)
    InitialContext.lookup("java:module/ExampleBean");
```

## Accessing Local Enterprise Beans That Implement Business Interfaces

Client access to enterprise beans that implement local business interfaces is accomplished using either dependency injection or JNDI lookup.

### EXAMPLE 16-4 Injecting an Enterprise Bean's Local Business Interface

To obtain a reference to the local business interface of an enterprise bean through dependency injection, use the `javax.ejb.EJB` annotation and specify the enterprise bean's local business interface name.

```
@EJB
Example example;
```

### EXAMPLE 16-5 Looking Up a Local Enterprise Bean Using JNDI

To obtain a reference to a local business interface of an enterprise bean using JNDI lookup, use the `javax.naming.InitialContext` interface's `lookup` method.

EXAMPLE 16-5 Looking Up a Local Enterprise Bean Using JNDI (Continued)

```
ExampleLocal example = (ExampleLocal)
    InitialContext.lookup("java:module/ExampleLocal");
```

## Remote Clients

A remote client of an enterprise bean has the following traits:

- It can run on a different machine and a different Java virtual machine (JVM) than the enterprise bean it accesses. (It is not required to run on a different JVM.)
- It can be a web component, an application client, or another enterprise bean.
- To a remote client, the location of the enterprise bean is transparent.
- The enterprise bean must implement a business interface. That is, remote clients *may not* access an enterprise bean using a no-interface view.

To create an enterprise bean that allows remote access, you must do one of the following:

- Decorate the business interface of the enterprise bean with the `@Remote` annotation:

```
@Remote
public interface InterfaceName { ... }
```

- Decorate the bean class with `@Remote`, specifying the business interface or interfaces:

```
@Remote(InterfaceName.class)
public class BeanName implements InterfaceName { ... }
```

The *remote interface* defines the business and life cycle methods that are specific to the bean. For example, the remote interface of a bean named `BankAccountBean` might have business methods named `deposit` and `credit`. [Figure 16-1](#) shows how the interface controls the client's view of an enterprise bean.

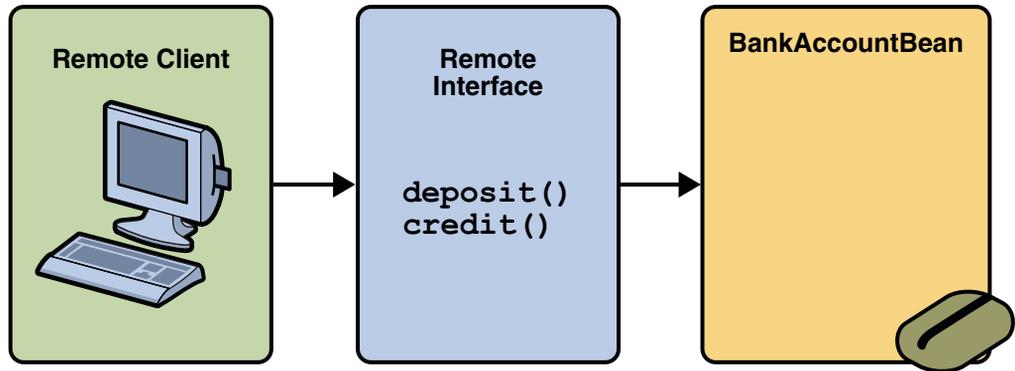


FIGURE 16-1 Interfaces for an Enterprise Bean with Remote Access

## Accessing Remote Enterprise Beans

Client access to an enterprise bean that implements a remote business interface is accomplished using either dependency injection or JNDI lookup.

### EXAMPLE 16-6 Injecting an Enterprise Bean's Remote Business Interface

To obtain a reference to the remote business interface of an enterprise bean through dependency injection, use the `javax.ejb.EJB` annotation and specify the enterprise bean's remote business interface name.

```
@EJB
Example example;
```

### EXAMPLE 16-7 Looking Up an Enterprise Bean's Remote Business Interface

To obtain a reference to a remote business interface of an enterprise bean using JNDI lookup, use the `javax.naming.InitialContext` interface's `lookup` method.

```
ExampleRemote example = (ExampleRemote)
    InitialContext.lookup("java:global/myApp/ExampleRemote");
```

## Web Service Clients

A web service client can access a Java EE application in two ways. First, the client can access a web service created with JAX-WS. (For more information on JAX-WS, see [Chapter 13, "Building Web Services with JAX-WS."](#)) Second, a web service client can invoke the business methods of a stateless session bean. Message beans cannot be accessed by web service clients.

Provided that it uses the correct protocols (SOAP, HTTP, WSDL), any web service client can access a stateless session bean, whether or not the client is written in the Java programming language. The client doesn't even "know" what technology implements the service: stateless session bean, JAX-WS, or some other technology. In addition, enterprise beans and web components can be clients of web services. This flexibility enables you to integrate Java EE applications with web services.

A web service client accesses a stateless session bean through the bean's web service endpoint implementation class. By default, all public methods in the bean class are accessible to web service clients. The `@WebMethod` annotation may be used to customize the behavior of web service methods. If the `@WebMethod` annotation is used to decorate the bean class's methods, only those methods decorated with `@WebMethod` are exposed to web service clients.

For a code sample, see [“A Web Service Example: `helloservice`” on page 324](#).

## Method Parameters and Access

The type of access affects the parameters of the bean methods that are called by clients. The following topics apply not only to method parameters but also to method return values.

### Isolation

The parameters of remote calls are more isolated than those of local calls. With remote calls, the client and bean operate on different copies of a parameter object. If the client changes the value of the object, the value of the copy in the bean does not change. This layer of isolation can help protect the bean if the client accidentally modifies the data.

In a local call, both the client and the bean can modify the same parameter object. In general, you should not rely on this side effect of local calls. Perhaps someday you will want to distribute your components, replacing the local calls with remote ones.

As with remote clients, web service clients operate on different copies of parameters than does the bean that implements the web service.

### Granularity of Accessed Data

Because remote calls are likely to be slower than local calls, the parameters in remote methods should be relatively coarse-grained. A coarse-grained object contains more data than a fine-grained one, so fewer access calls are required. For the same reason, the parameters of the methods called by web service clients should also be coarse-grained.

## The Contents of an Enterprise Bean

To develop an enterprise bean, you must provide the following files:

- **Enterprise bean class:** Implements the business methods of the enterprise bean and any life cycle callback methods.
- **Business Interfaces:** The business interface defines the business methods implemented by the enterprise bean class. A business interface is not required if the enterprise bean exposes a local, no-interface view.
- **Helper classes:** Other classes needed by the enterprise bean class, such as exception and utility classes.

Package the programming artifacts in the preceding list into either an EJB JAR file (a standalone module that stores the enterprise bean), or within a web application archive (WAR) module.

## Packaging Enterprise Beans In EJB JAR Modules

An EJB JAR file is portable and can be used for different applications.

To assemble a Java EE application, package one or more modules (such as EJB JAR files) into an EAR file, the archive file that holds the application. When deploying the EAR file that contains the enterprise bean's EJB JAR file, you also deploy the enterprise bean to the Enterprise Server. You can also deploy an EJB JAR that is not contained in an EAR file. [Figure 16–2](#) shows the contents of an EJB JAR file.

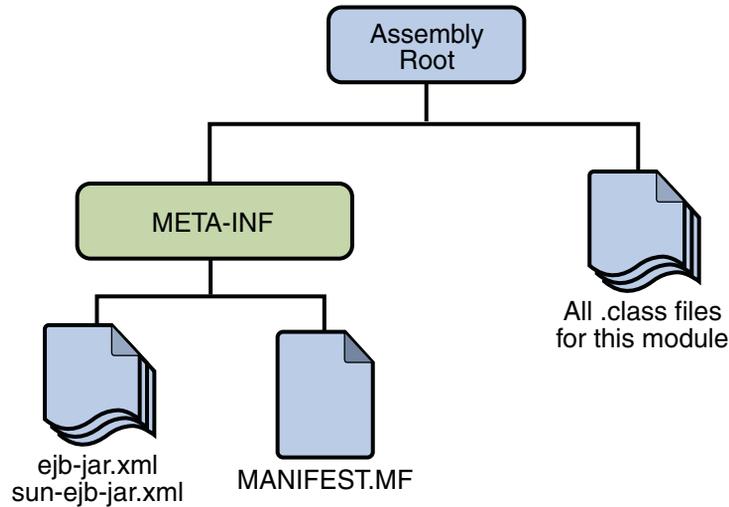


FIGURE 16-2 Structure of an Enterprise Bean JAR

## Packaging Enterprise Beans in WAR Modules

Enterprise beans often provide the business logic of a web application. In these cases, packaging the enterprise bean within the web application's WAR module simplifies deployment and application organization. Enterprise beans may be packaged within a WAR module as Java programming language class files or within a JAR file that is bundled within the WAR module.

To include enterprise bean class files in a WAR module, the class files should be in the `WEB-INF/classes` directory.

To include a JAR file that contains enterprise beans in a WAR module, add the JAR to the `WEB-INF/lib` directory of the WAR module.

WAR modules that contain enterprise beans do not require an `ejb-jar.xml` deployment descriptor. If the application uses `ejb-jar.xml`, it must be located in the WAR module's `WEB-INF` directory.

JAR files that contain enterprise bean classes packaged within a WAR module are not considered EJB JAR files, even if the bundled JAR file conforms to the format of an EJB JAR file. The enterprise beans contained within the JAR file are semantically equivalent to enterprise beans located in the WAR module's `WEB-INF/classes` directory, and the environment namespace of all the enterprise beans are scoped to the WAR module.

### EXAMPLE 16-8 Enterprise Beans Packaged In A WAR Module

Suppose a web application consisted of a shopping cart enterprise bean, a credit card processing enterprise bean, and a Java servlet front-end. The shopping cart bean exposes a local, no-interface view and is defined as follows:

**EXAMPLE 16-8** Enterprise Beans Packaged In A WAR Module *(Continued)*

```
package com.example.cart;

@Stateless
public class CartBean { ... }
```

The credit card processing bean is packaged within its own JAR file, `cc.jar`. It exposes a local, no-interface view and is defined as follows:

```
package com.example.cc;

@Stateless
public class CreditCardBean { ... }
```

The servlet, `com.example.web.StoreServlet` handles the web front-end and uses both `CartBean` and `CreditCardBean`. The WAR module layout for this application looks as follows:

```
WEB-INF/classes/com/example/cart/CartBean.class
WEB-INF/classes/com/example/web/StoreServlet
WEB-INF/lib/cc.jar
WEB-INF/ejb-jar.xml
WEB-INF/web.xml
```

## Naming Conventions for Enterprise Beans

Because enterprise beans are composed of multiple parts, it's useful to follow a naming convention for your applications. [Table 16-2](#) summarizes the conventions for the example beans in this tutorial.

**TABLE 16-2** Naming Conventions for Enterprise Beans

Item	Syntax	Example
Enterprise bean name	<i>name</i> Bean	AccountBean
Enterprise bean class	<i>name</i> Bean	AccountBean
Business interface	<i>name</i>	Account

## The Life Cycles of Enterprise Beans

An enterprise bean goes through various stages during its lifetime, or life cycle. Each type of enterprise bean (stateful session, stateless session, or message-driven) has a different life cycle.

The descriptions that follow refer to methods that are explained along with the code examples in the next two chapters. If you are new to enterprise beans, you should skip this section and run the code examples first.

### The Life Cycle of a Stateful Session Bean

Figure 16–3 illustrates the stages that a session bean passes through during its lifetime. The client initiates the life cycle by obtaining a reference to a stateful session bean. The container performs any dependency injection and then invokes the method annotated with `@PostConstruct`, if any. The bean is now ready to have its business methods invoked by the client.

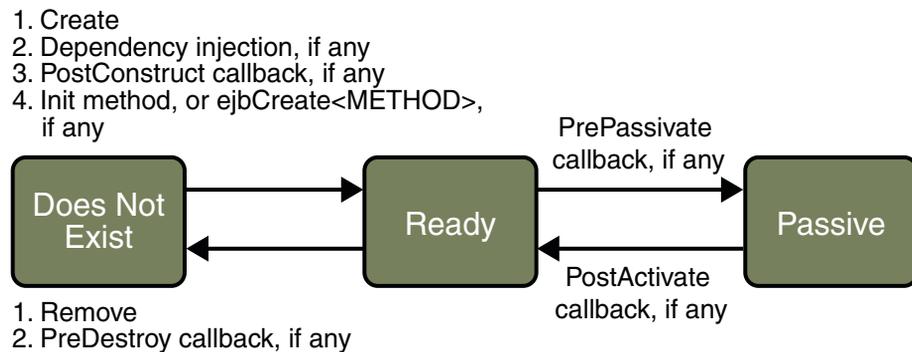


FIGURE 16–3 Life Cycle of a Stateful Session Bean

While in the ready stage, the EJB container may decide to deactivate, or *passivate*, the bean by moving it from memory to secondary storage. (Typically, the EJB container uses a least-recently-used algorithm to select a bean for passivation.) The EJB container invokes the method annotated `@PrePassivate`, if any, immediately before passivating it. If a client invokes a business method on the bean while it is in the passive stage, the EJB container activates the bean, calls the method annotated `@PostActivate`, if any, and then moves it to the ready stage.

At the end of the life cycle, the client invokes a method annotated `@Remove`, and the EJB container calls the method annotated `@PreDestroy`, if any. The bean's instance is then ready for garbage collection.

Your code controls the invocation of only one lifecycle method: the method annotated `@Remove`. All other methods in [Figure 16–3](#) are invoked by the EJB container. See [Chapter 32, “Resource Connections,”](#) for more information.

## The Lifecycle of a Stateless Session Bean

Because a stateless session bean is never passivated, its life cycle has only two stages: nonexistent and ready for the invocation of business methods. [Figure 16–4](#) illustrates the stages of a stateless session bean.

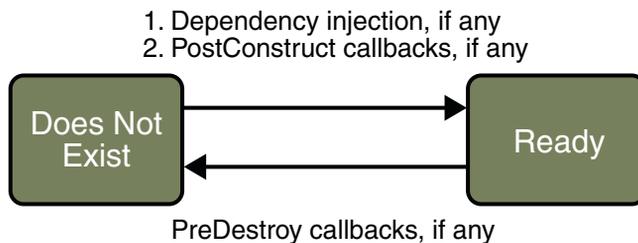


FIGURE 16–4 Lifecycle of a Stateless Session Bean

The EJB container typically creates and maintains a pool of stateless session beans, beginning the stateless session bean’s lifecycle. The container performs any dependency injection and then invokes the method annotated `@PostConstruct`, if it exists. The bean is now ready to have its business methods invoked by a client.

At the end of the lifecycle, the EJB container calls the method annotated `@PreDestroy`, if it exists. The bean’s instance is then ready for garbage collection.

## The Lifecycle of a Singleton Session Bean

Like a stateless session bean, a singleton session bean is never passivated and has only two stages: nonexistent and ready for the invocation of business methods.

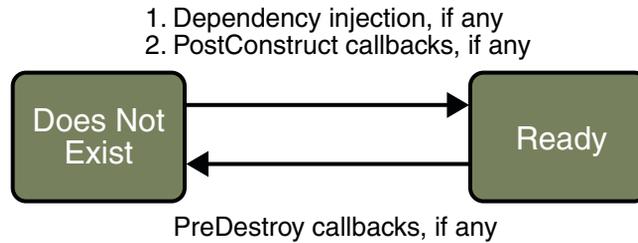


FIGURE 16-5 Lifecycle of a Singleton Session Bean

The EJB container initiates the singleton session bean lifecycle by creating the singleton instance. This occurs upon application deployment if the singleton is annotated with the `@Startup` annotation. The container performs any dependency injection and then invokes the method annotated `@PostConstruct`, if it exists. The singleton session bean is now ready to have its business methods invoked by the client.

At the end of the lifecycle, the EJB container calls the method annotated `@PreDestroy`, if it exists. The singleton session bean is now ready for garbage collection.

## The Lifecycle of a Message-Driven Bean

Figure 16-6 illustrates the stages in the life cycle of a message-driven bean.

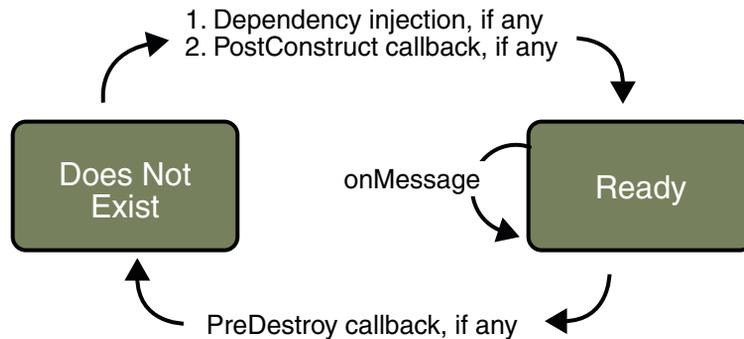


FIGURE 16-6 Life Cycle of a Message-Driven Bean

The EJB container usually creates a pool of message-driven bean instances. For each instance, the EJB container performs these tasks:

1. If the message-driven bean uses dependency injection, the container injects these references before instantiating the instance.
2. The container calls the method annotated `@PostConstruct`, if any.

Like a stateless session bean, a message-driven bean is never passivated, and it has only two states: nonexistent and ready to receive messages.

At the end of the life cycle, the container calls the method annotated `@PreDestroy`, if any. The bean's instance is then ready for garbage collection.

## Further Information about Enterprise Beans

For more information on Enterprise JavaBeans technology, see:

- Enterprise JavaBeans 3.1 specification:  
<http://java.sun.com/products/ejb/docs.html>
- The Enterprise JavaBeans web site:  
<http://java.sun.com/products/ejb>

# Getting Started with Enterprise Beans

---

This chapter shows how to develop, deploy, and run a simple Java EE application named `converter`. The purpose of `converter` is to calculate currency conversions between Japanese yen and Eurodollars. `converter` consists of an enterprise bean, which performs the calculations, and two types of clients: an application client and a web client.

Here's an overview of the steps you'll follow in this chapter:

1. Create the enterprise bean: `ConverterBean`.
2. Create the web client.
3. Deploy `converter` onto the server.
4. Using a browser, run the web client.

Before proceeding, make sure that you've done the following:

- Read [Chapter 1, "Overview."](#)
- Become familiar with enterprise beans (see [Chapter 16, "Enterprise Beans"](#)).
- Started the server (see ["Starting and Stopping the Enterprise Server"](#) on page 63).

## Creating the Enterprise Bean

The enterprise bean in our example is a stateless session bean called `ConverterBean`. The source code for `ConverterBean` is in the `tut-install/examples/ejb/converter/src/java/` directory.

Creating `ConverterBean` requires these steps:

1. Coding the bean's implementation class (the source code is provided)
2. Compiling the source code with the Ant tool

## Coding the Enterprise Bean

The enterprise bean in this example needs the following code:

- Enterprise bean class

### Coding the Enterprise Bean Class

The enterprise bean class for this example is called `ConverterBean`. This class implements two business methods (`dollarToYen` and `yenToEuro`). Because the enterprise bean class doesn't implement a business interface, the enterprise bean exposes a local, no-interface view. The public methods in the enterprise bean class are available to clients that obtain a reference to `ConverterBean`. The source code for the `ConverterBean` class follows.

```
package com.sun.tutorial.javaee.ejb;

import java.math.BigDecimal;
import javax.ejb.*;

@Stateless
public class ConverterBean {
    private BigDecimal yenRate = new BigDecimal("115.3100");
    private BigDecimal euroRate = new BigDecimal("0.0071");

    public BigDecimal dollarToYen(BigDecimal dollars) {
        BigDecimal result = dollars.multiply(yenRate);
        return result.setScale(2, BigDecimal.ROUND_UP);
    }

    public BigDecimal yenToEuro(BigDecimal yen) {
        BigDecimal result = yen.multiply(euroRate);
        return result.setScale(2, BigDecimal.ROUND_UP);
    }
}
```

Note the `@Stateless` annotation decorating the enterprise bean class. This lets the container know that `ConverterBean` is a stateless session bean.

## Creating the converter Web Client

The web client is contained in the servlet class

```
tut-install/examples/ejb/converter/src/java/converter/web/ConverterServlet.java. A
```

Java servlet is a web component that responds to HTTP requests.

## Coding the converter Web Client

The `ConverterServlet` class uses dependency injection to obtain a reference to `ConverterBean`. The `javax.ejb.EJB` annotation is added to the declaration of the private member variable `converterBean`, which is of type `ConverterBean`. `ConverterBean` exposes a local, no-interface view, so the enterprise bean implementation class is the variable type.

```
@WebServlet
public class ConverterServlet extends HttpServlet {
    @EJB
    ConverterBean converterBean;
    ...
}
```

When the user enters an amount to be converted to Yen and Euro, the amount is retrieved from the request parameters, then the `ConverterBean.dollarToYen` and `ConverterBean.yenToEuro` methods are called.

```
...
try {
    String amount = request.getParameter("amount");
    if (amount != null && amount.length() > 0) {
        // convert the amount to a BigDecimal from the request parameter
        BigDecimal d = new BigDecimal(amount);
        // call the ConverterBean.dollarToYen() method to get the amount
        // in Yen
        BigDecimal yenAmount = converter.dollarToYen(d);

        // call the ConverterBean.yenToEuro() method to get the amount
        // in Euros
        BigDecimal euroAmount = converter.yenToEuro(yenAmount);
        ...
    }
    ...
}
```

The results are displayed to the user.

## Compiling, Packaging, and Running the converter Example

Now you are ready to compile the enterprise bean class (`ConverterBean.java`) and the servlet class (`ConverterServlet.java`), and package the compiled classes into a WAR file.

## Compiling, Packaging, and Running the converter Example in NetBeans IDE

Follow these instructions to build and package the converter example in NetBeans IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to *tut-install/examples/ejb/*.
3. Select the converter folder.
4. Select the Open as Main Project and Open Required Projects check boxes.
5. Click Open Project.
6. In the Projects tab, right-click the converter project and select Run. A web browser window will open the URL `http://localhost:8080/converter`

## Compiling, Packaging, and Running the converter Example Using Ant

To compile and package converter using Ant, do the following:

1. In a terminal window, go to this directory:

```
tut-install/examples/ejb/converter/
```

2. Type the following command:

```
ant all
```

3. Open a web browser to the following URL:

```
http://localhost:8080/converter
```

This command calls the default task, which compiles the source files for the enterprise bean and the servlet, placing the class files in the `build` subdirectory (not the `src` directory) of the project. The default task packages the project into a WAR module: `converter.war`. For more information about the Ant tool, see [“Building the Examples” on page 65](#).

---

**Note** – When compiling the code, the preceding ant task includes the Java EE API JAR files in the classpath. These JARs reside in the `modules` directory of your Enterprise Server installation. If you plan to use other tools to compile the source code for Java EE components, make sure that the classpath includes the Java EE API JAR files.

---

After entering `100` in the input field and clicking Submit, you should see the screen shown in [Figure 17-1](#).



FIGURE 17-1 converter Web Client

## Modifying the Java EE Application

The Enterprise Server supports iterative development. Whenever you make a change to a Java EE application, you must redeploy the application.

### Modifying a Class File

To modify a class file in an enterprise bean, you change the source code, recompile it, and redeploy the application. For example, if you want to change the exchange rate in the `dollarToYen` business method of the `ConverterBean` class, you would follow these steps.

1. Edit `ConverterBean.java` and save the file.
2. Recompile `ConverterBean.java` in NetBeans IDE by right-clicking the `converter` project and selecting `Run`.

This recompiles the `ConverterBean.java` file, replaces the old class file in the build directory, and redeploys the application to Enterprise Server.

3. Recompile `ConverterBean.java` using Ant.
  - a. In a terminal window, go to the `tut-install/examples/ejb/converter/` subdirectory.
  - b. Type the following command:

```
ant all
```

This command repackages, deploys, and runs the application.

To modify `ConverterServlet` the procedure is the same as that described in the preceding steps.



# Running the Enterprise Bean Examples

---

Session beans provide a simple but powerful way to encapsulate business logic within an application. They can be accessed from remote Java clients, web service clients, and from components running in the same server.

In [Chapter 17, “Getting Started with Enterprise Beans,”](#) you built a stateless session bean named `ConverterBean`. This chapter examines the source code of four more session beans:

- `CartBean`: a stateful session bean that is accessed by a remote client
- `CounterBean`: a singleton session bean.
- `HelloServiceBean`: a stateless session bean that implements a web service
- `TimerSessionBean`: a stateless session bean that sets a timer

## The cart Example

The cart example represents a shopping cart in an online bookstore, and uses a stateful session bean to manage the operations of the shopping cart. The bean’s client can add a book to the cart, remove a book, or retrieve the cart’s contents. To assemble `cart`, you need the following code:

- Session bean class (`CartBean`)
- Remote business interface (`Cart`)

All session beans require a session bean class. All enterprise beans that permit remote access must have a remote business interface. To meet the needs of a specific application, an enterprise bean may also need some helper classes. The `CartBean` session bean uses two helper classes (`BookException` and `IdVerifier`) which are discussed in the section [“Helper Classes” on page 312](#).

The source code for this example is in the `tut-install/examples/ejb/cart/` directory.

## The Business Interface

The Cart business interface is a plain Java interface that defines all the business methods implemented in the bean class. If the bean class implements a single interface, that interface is assumed to be the business interface. The business interface is a local interface unless it is annotated with the `javax.ejb.Remote` annotation; the `javax.ejb.Local` annotation is optional in this case.

The bean class may implement more than one interface. If the bean class implements more than one interface, either the business interfaces must be explicitly annotated either `@Local` or `@Remote`, or the business interfaces must be specified by decorating the bean class with `@Local` or `@Remote`. However, the following interfaces are excluded when determining if the bean class implements more than one interface:

- `java.io.Serializable`
- `java.io.Externalizable`
- Any of the interfaces defined by the `javax.ejb` package

The source code for the Cart business interface follows:

```
package com.sun.tutorial.javaee.ejb;

import java.util.List;
import javax.ejb.Remote;

@Remote
public interface Cart {
    public void initialize(String person) throws BookException;
    public void initialize(String person, String id)
        throws BookException;
    public void addBook(String title);
    public void removeBook(String title) throws BookException;
    public List<String> getContents();
    public void remove();
}
```

## Session Bean Class

The session bean class for this example is called `CartBean`. Like any stateful session bean, the `CartBean` class must meet these requirements:

- The class is annotated `@Stateful`.
- The class implements the business methods defined in the business interface.

Stateful session beans also may:

- Implement the business interface, a plain Java interface. It is good practice to implement the bean's business interface.
- Implement any optional life cycle callback methods, annotated `@PostConstruct`, `@PreDestroy`, `@PostActivate`, and `@PrePassivate`.
- Implement any optional business methods annotated `@Remove`.

The source code for the `CartBean` class follows.

```
package com.sun.tutorial.javaee.ejb;

import java.util.ArrayList;
import java.util.List;
import javax.ejb.Remove;
import javax.ejb.Stateful;

@Stateful
public class CartBean implements Cart {
    String customerName;
    String customerId;
    List<String> contents;

    public void initialize(String person) throws BookException {
        if (person == null) {
            throw new BookException("Null person not allowed.");
        } else {
            customerName = person;
        }

        customerId = "0";
        contents = new ArrayList<String>();
    }

    public void initialize(String person, String id)
        throws BookException {
        if (person == null) {
            throw new BookException("Null person not allowed.");
        } else {

            customerName = person;
        }

        IdVerifier idChecker = new IdVerifier();

        if (idChecker.validate(id)) {
            customerId = id;
        } else {
```

```
        throw new BookException("Invalid id: " + id);
    }

    contents = new ArrayList<String>();
}

public void addBook(String title) {
    contents.add(title);
}

public void removeBook(String title) throws BookException {
    boolean result = contents.remove(title);
    if (result == false) {
        throw new BookException(title + " not in cart.");
    }
}

public List<String> getContents() {
    return contents;
}

@Remove
public void remove() {
    contents = null;
}
}
```

## Lifecycle Callback Methods

Methods in the bean class may be declared as a lifecycle callback method by annotating the method with the following annotations:

- `javax.annotation.PostConstruct`
- `javax.annotation.PreDestroy`
- `javax.ejb.PostActivate`
- `javax.ejb.PrePassivate`

Lifecycle callback methods must return `void` and have no parameters.

`@PostConstruct` methods are invoked by the container on newly constructed bean instances after all dependency injection has completed and before the first business method is invoked on the enterprise bean.

`@PreDestroy` methods are invoked after any method annotated `@Remove` has completed, and before the container removes the enterprise bean instance.

`@PostActivate` methods are invoked by the container after the container moves the bean from secondary storage to active status.

`@PrePassivate` methods are invoked by the container before the container passivates the enterprise bean, meaning the container temporarily removes the bean from the environment and saves it to secondary storage.

## Business Methods

The primary purpose of a session bean is to run business tasks for the client. The client invokes business methods on the object reference it gets from dependency injection or JNDI lookup. From the client's perspective, the business methods appear to run locally, but they actually run remotely in the session bean. The following code snippet shows how the `CartClient` program invokes the business methods:

```
cart.create("Duke DeEarl", "123");
...
cart.addBook("Bel Canto");
...
List<String> bookList = cart.getContents();
...
cart.removeBook("Gravity's Rainbow");
```

The `CartBean` class implements the business methods in the following code:

```
public void addBook(String title) {
    contents.addElement(title);
}

public void removeBook(String title) throws BookException {
    boolean result = contents.remove(title);
    if (result == false) {
        throw new BookException(title + "not in cart.");
    }
}

public List<String> getContents() {
    return contents;
}
```

The signature of a business method must conform to these rules:

- The method name must not begin with `ejb` to avoid conflicts with callback methods defined by the EJB architecture. For example, you cannot call a business method `ejbCreate` or `ejbActivate`.
- The access control modifier must be `public`.
- If the bean allows remote access through a remote business interface, the arguments and return types must be legal types for the Java RMI API.
- If the bean is a web service endpoint, the arguments and return types for the methods annotated `@WebMethod` must be legal types for JAX-WS.

- The modifier must not be `static` or `final`.

The `throws` clause can include exceptions that you define for your application. The `removeBook` method, for example, throws a `BookException` if the book is not in the cart.

To indicate a system-level problem, such as the inability to connect to a database, a business method should throw a `javax.ejb.EJBException`. The container will not wrap application exceptions such as `BookException`. Because `EJBException` is a subclass of `RuntimeException`, you do not need to include it in the `throws` clause of the business method.

## The Remove Method

Business methods annotated with `javax.ejb.Remove` in the stateful session bean class can be invoked by enterprise bean clients to remove the bean instance. The container will remove the enterprise bean after a `@Remove` method completes, either normally or abnormally.

In `CartBean`, the `remove` method is a `@Remove` method:

```
@Remove
public void remove() {
    contents = null;
}
```

## Helper Classes

The `CartBean` session bean has two helper classes: `BookException` and `IdVerifier`. The `BookException` is thrown by the `removeBook` method, and the `IdVerifier` validates the `customerId` in one of the `create` methods. Helper classes may reside in an EJB JAR file that contains the enterprise bean class, a WAR file if the enterprise bean is packaged within a WAR, or in an EAR that contains an EJB JAR or a WAR file that contains an enterprise bean.

## Building, Packaging, Deploying, and Running the cart Example

You can build, package, deploy, and run the `cart` application using either NetBeans IDE or the Ant tool.

## Building, Packaging, and Deploying the cart Example Using NetBeans IDE

Follow these instructions to build, package, and deploy the cart example to your Application Server instance using the NetBeans IDE IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to *tut-install/examples/ejb/*.
3. Select the cart folder.
4. Select the Open as Main Project and Open Required Projects check boxes.
5. Click Open Project Folder.
6. In the Projects tab, right-click the cart project and select Deploy Project.

This builds and packages the application into *cart.ear*, located in *tut-install/examples/ejb/cart/dist/*, and deploys this EAR file to your Application Server instance.

## Running the cart Application Client Using NetBeans IDE

To run cart's application client, select Run→Run Main Project. You will see the output of the application client in the Output pane:

```
...
Retrieving book title from cart: Infinite Jest
Retrieving book title from cart: Bel Canto
Retrieving book title from cart: Kafka on the Shore
Removing "Gravity's Rainbow" from cart.
Caught a BookException: "Gravity's Rainbow" not in cart.
Java Result: 1
run-cart-app-client:
run-nb:
BUILD SUCCESSFUL (total time: 14 seconds)
```

## Building, Packaging, and Deploying the cart Example Using Ant

Now you are ready to compile the remote interface (*Cart.java*), the home interface (*CartHome.java*), the enterprise bean class (*CartBean.java*), the client class (*CartClient.java*), and the helper classes (*BookException.java* and *IdVerifier.java*).

1. In a terminal window, go to this directory:

```
tut-install/examples/ejb/cart/
```

2. Type the following command:

```
ant
```

This command calls the `default` target, which builds and packages the application into an EAR file, *cart.ear*, located in the `dist` directory.

3. Type the following command:

```
ant deploy
```

`cart.ear` will be deployed to the Application Server.

## Running the cart Application Client Using Ant

When you run the client, the application client container injects any component references declared in the application client class, in this case the reference to the Cart enterprise bean. To run the application client, perform the following steps.

1. In a terminal window, go to this directory:

```
tut-install/examples/ejb/cart/
```

2. Type the following command:

```
ant run
```

This task will retrieve the application client JAR, `cartClient.jar` and run the application client. `cartClient.jar` contains the application client class, the helper class `BookException`, and the Cart business interface.

This is the equivalent of running:

```
appclient -client cartClient.jar
```

3. In the terminal window, the client displays these lines:

```
[echo] running application client container.  
[exec] Retrieving book title from cart: Infinite Jest  
[exec] Retrieving book title from cart: Bel Canto  
[exec] Retrieving book title from cart: Kafka on the Shore  
[exec] Removing "Gravity's Rainbow" from cart.  
[exec] Caught a BookException: "Gravity's Rainbow" not in cart.  
[exec] Result: 1
```

## The all Task

As a convenience, the `all` task will build, package, deploy, and run the application. To do this, enter the following command:

```
ant all
```

## Undeploying the cart Example

To undeploy `cart.ear` using NetBeans IDE:

1. Click the Runtime tab.
2. Expand the Servers node and locate the Application Server instance to which you deployed `cart`.
3. Expand your Application Server instance node, then Applications→Enterprise Applications.
4. Right-click `cart` and select Undeploy.

To undeploy `cart.ear` using Ant, enter the following command:

```
ant undeploy
```

## A Singleton Session Bean Example: counter

The `counter` example demonstrates how to create a singleton session bean.

### Creating a Singleton Session Bean

The `javax.ejb.Singleton` annotation is used to specify that the enterprise bean implementation class is a singleton session bean.

```
@Singleton
public class SingletonBean { ... }
```

### Initializing Singleton Session Beans

The EJB container is responsible for determining when to initialize a singleton session bean instance unless the singleton session bean implementation class is annotated with the `javax.ejb.Startup` annotation. This is sometimes called *eager initialization*. In this case, the EJB container must initialize the singleton session bean upon application startup. The singleton session bean is initialized before the EJB container delivers client requests to any enterprise beans in the application. This allows the singleton session bean to perform, for example, application startup tasks.

#### EXAMPLE 18-1 An Eagerly Initialized Singleton Session Bean

The following singleton session bean stores the status of an application, and is eagerly initialized:

**EXAMPLE 18-1** An Eagerly Initialized Singleton Session Bean *(Continued)*

```
@Startup
@Singleton
public class StatusBean {
    private String status;

    @PostConstruct
    void init {
        status = "Ready";
    }
    ...
}
```

Sometimes multiple singleton session beans are used to initialize data for an application, and therefore must be initialized in a specific order. In these cases, use the `javax.ejb.DependsOn` annotation to declare the startup dependencies of the singleton session bean. The `@DependsOn` annotation's value attribute is one or more strings that specify the name of the target singleton session bean. If more than one dependent singleton bean is specified in `@DependsOn`, the order that they are listed is not necessarily the order that the EJB container will initialize the target singleton session beans.

**EXAMPLE 18-2** Specifying the Ordering Of Singleton Session Bean Initialization

The following singleton session bean, `PrimaryBean` should be started up first:

```
@Singleton
public class PrimaryBean { ... }
```

`SecondaryBean` depends on `PrimaryBean`:

```
@Singleton
@DependsOn("PrimaryBean")
public class SecondaryBean { ... }
```

This guarantees that the EJB container will initialize `PrimaryBean` before `SecondaryBean`.

**EXAMPLE 18-3** Specifying Multiple Dependent Singleton Session Beans

The following singleton session bean, `TertiaryBean`, depends on `PrimaryBean` and `SecondaryBean`:

```
@Singleton
@DependsOn("PrimaryBean", "SecondaryBean")
public class TertiaryBean { ... }
```

**EXAMPLE 18-3** Specifying Multiple Dependent Singleton Session Beans (Continued)

SecondaryBean explicitly requires PrimaryBean to be initialized before it is initialized (through its own @DependsOn annotation). In this case, the EJB container will first initialize PrimaryBean, then SecondaryBean, and finally TertiaryBean.

If, however, SecondaryBean did not explicitly depend on PrimaryBean, the EJB container may initialize either PrimaryBean or SecondaryBean first. That is, the EJB container could initialize the singletons in the following order: SecondaryBean, PrimaryBean, TertiaryBean.

## Managing Concurrent Access in a Singleton Session Bean

Singleton session beans are designed for *concurrent access*, or situations where many clients need to access a single instance of a session bean at the same time. A singleton's client only needs a reference to a singleton in order to invoke any business methods exposed by the singleton, and doesn't need to worry about any other clients that may be simultaneously invoking business methods on the same singleton.

When creating a singleton session bean there are two ways of controlling concurrent access to the singleton's business methods: *container-managed concurrency* and *bean-managed concurrency*.

The `javax.ejb.ConcurrencyManagement` annotation is used to specify either container-managed or bean-managed concurrency for the singleton. `@ConcurrencyManagement` requires a type attribute to be set, one of `javax.ejb.ConcurrencyManagementType.CONTAINER` or `javax.ejb.ConcurrencyManagementType.BEAN`. If no `@ConcurrencyManagement` annotation is present on the singleton implementation class, the EJB container default of container-managed concurrency is used.

### Container-Managed Concurrency

If a singleton uses container-managed concurrency, the EJB container controls client access to the business methods of the singleton. The `javax.ejb.Lock` annotation and a `javax.ejb.LockType` type are used to specify the access level of the singleton's business methods or `@Timeout` methods.

Annotate a singleton's business or timeout method using `@Lock(READ)` if the method can be concurrently accessed, or shared, with many clients. Annotate the business or timeout method with `@Lock(WRITE)` if the singleton session bean should be locked to other clients while a client is calling that method. Typically, the `@Lock(WRITE)` annotation is used when clients are modifying the state of the singleton.

Annotating a singleton class with `@Lock` specifies that all the business methods and any timeout methods of the singleton will use the specified lock type unless they explicitly set the lock type with a method-level `@Lock` annotation. If no `@Lock` annotation is present on the singleton class, the default lock type of `@Lock(WRITE)` is applied to all business and timeout methods.

**EXAMPLE 18-4** Specifying Container-Managed Concurrency in a Singleton Session Bean

The following example shows how to use the `@ConcurrencyManagement`, `@Lock(READ)`, and `@Lock(WRITE)` annotations for a singleton that uses container-managed concurrency.

Although by default singletons use container-managed concurrency, the `@ConcurrencyManagement(CONTAINER)` annotation may be added at the class level of the singleton to explicitly set the concurrency management type.

```
@ConcurrencyManagement (CONTAINER)
@Singleton
public class ExampleSingletonBean {
    private String state;

    @Lock (READ)
    public String getState() {
        return state;
    }

    @Lock (WRITE)
    public void setState(String newState) {
        state = newState;
    }
}
```

The `getState` method can be accessed by many clients at the same time, because it is annotated with `@Lock(READ)`. When the `setState` method is called, however, all the methods in `ExampleSingletonBean` will be locked to other clients because `setState` is annotated with `@Lock(WRITE)`. This prevents two clients from attempting to simultaneously change the state variable of `ExampleSingletonBean`.

**EXAMPLE 18-5** Using Class- and Method-Level `@Lock` Annotations in a Singleton Session Bean

The `getData` and `getStatus` methods in the following singleton are of type `READ`, and the `setStatus` method is of type `WRITE`:

```
@Singleton
@Lock(READ)
public class SharedSingletonBean {
    private String data;
    private String status;

    public String getData() {
        return data;
    }

    public String getStatus() {
```

**EXAMPLE 18-5** Using Class- and Method-Level `@Lock` Annotations in a Singleton Session Bean  
(Continued)

```

        return status;
    }

    @Lock(WRITE)
    public void setStatus(String newStatus) {
        status = newStatus;
    }
}

```

If a method is of locking type `WRITE`, client access to all the singleton's methods are blocked until the current client finishes its method call, or an access timeout occurs. When an access timeout occurs, the EJB container throws a `javax.ejb.ConcurrentAccessTimeoutException`. The `javax.ejb.AccessTimeout` annotation is used to specify the number of milliseconds before an access timeout occurs. If `@AccessTimeout` is added at the class level of a singleton, it specifies the access timeout value for all methods in the singleton unless a method explicitly overrides the default with its own `@AccessTimeout` annotation.

The `@AccessTimeout` annotation can be applied to both `@Lock(READ)` and `@Lock(WRITE)` methods.

`@AccessTimeout` has one required element, `value`, and one optional element, `timeUnit`. By default, the `value` is specified in milliseconds. To change the `value` unit, set `timeUnit` to one of the `java.util.concurrent.TimeUnit` constants: `MICROSECONDS`, `MILLISECONDS`, `MICROSECONDS`, or `SECONDS`.

**EXAMPLE 18-6** Setting the Access Timeout in a Singleton

The following singleton has a default access timeout value of 120,000 milliseconds, or 2 minutes. The `doTediousOperation` method overrides the default access timeout and sets the value to 360,000 milliseconds, or 6 minutes.

```

@Singleton
@AccessTimeout(value=120000)
public class StatusSingletonBean {
    private String status;

    @Lock(WRITE)
    public void setStatus(String new Status) {
        status = newStatus;
    }

    @Lock(WRITE)
    @AccessTimeout(value=360000)
    public void doTediousOperation {

```

**EXAMPLE 18-6** Setting the Access Timeout in a Singleton *(Continued)*

```
    ...  
    }  
}
```

**EXAMPLE 18-7** Setting the Access Timeout in a Singleton in Seconds

The following singleton has a default access timeout value of 60 seconds, specified using the `TimeUnit.SECONDS` constant.

```
@Singleton  
@AccessTimeout(value=60, timeUnit=SECONDS)  
public class StatusSingletonBean { ... }
```

## Bean-Managed Concurrency

Singletons that use bean-managed concurrency allow full concurrent access to all the business and timeout methods in the singleton. The developer of the singleton is responsible for ensuring that the state of the singleton is synchronized across all clients. Developers who create singletons with bean-managed concurrency are allowed to use the Java programming language synchronization primitives like `synchronization` and `volatile` to prevent errors during concurrent access.

**EXAMPLE 18-8** Specifying Bean-Managed Concurrency in a Singleton Session Bean

Add a `@ConcurrencyManagement` annotation at the class level of the singleton to specify bean-managed concurrency.

```
@ConcurrencyManagement(BEAN)  
@Singleton  
public class AnotherSingletonBean { ... }
```

## Handling Errors in a Singleton Session Bean

If a singleton session bean encounters an error when it is initialized by the EJB container, that singleton instance will be destroyed.

Unlike other enterprise beans, once a singleton session bean instance is initialized it is not destroyed if the singleton's business or lifecycle methods cause system exceptions. This ensures that the same singleton instance is used throughout the application lifecycle.

## The Architecture of the counter Example

The counter example consists of a singleton session bean, `CounterBean`, and a JavaServer Faces Facelets web front-end.

CounterBean is a simple singleton with one method, `getHits`, that returns an integer representing the number of times a web page has been accessed. Here is the code of CounterBean:

```
package counter.ejb;

import javax.ejb.Singleton;

/**
 *
 * @author ian
 * CounterBean is a simple singleton session bean that records the number
 * of hits to a web page.
 */
@Singleton
public class CounterBean {
    private int hits = 1;

    // Increment and return the number of hits
    public int getHits() {
        return hits++;
    }
}
```

The `@Singleton` annotation marks CounterBean as a singleton session bean. CounterBean uses a local, no-interface view.

CounterBean uses the EJB container's default metadata values for singletons to simplify the coding of the singleton implementation class. There is no `@ConcurrencyManagement` annotation on the class, so the default of container-managed concurrency access is applied. There is no `@Lock` annotation on the class or business method, so the default of `@Lock(WRITE)` is applied to the only business method, `getHits`. The following version of CounterBean is functionally equivalent to the version above:

```
package counter.ejb;

import javax.ejb.Singleton;
import javax.ejb.ConcurrencyManagement;
import static javax.ejb.ConcurrencyManagementType.CONTAINER;
import javax.ejb.Lock;
import javax.ejb.LockType.WRITE;

/**
 *
 * @author ian
 * CounterBean is a simple singleton session bean that records the number
 * of hits to a web page.
 */
```

```
@Singleton
@ConcurrencyManagement(CONTAINER)
public class CounterBean {
    private int hits = 1;

    // Increment and return the number of hits
    @Lock(WRITE)
    public int getHits() {
        return hits++;
    }
}
```

The web-front end of counter consists of a JSF managed bean, `Count.java`, that is used by the Facelets XHTML files `template.xhtml` and `template-client.xhtml`. The `Count` JSF managed bean obtains a reference to `CounterBean` through dependency injection. `Count` defines a `hitCount` JavaBeans property. When the `getHitCount` getter method is called from the XHTML files, `CounterBean`'s `getHits` method is called to return the current number of page hits.

Here's the `Count` managed bean class:

```
@ManagedBean
@SessionScoped
public class Count {
    @EJB
    private CounterBean counterBean;

    private int hitCount;

    public Count() {
        this.hitCount = 0;
    }

    public int getHitCount() {
        hitCount = counterBean.getHits();
        return hitCount;
    }

    public void setHitCount(int newHits) {
        this.hitCount = newHits;
    }
}
```

The `template.xhtml` and `template-client.xhtml` files are used to render a Facelets view that displays the number of hits to that view. The `template-client.xhtml` file uses an expression language statement, `#{count.hitCount}`, to access the `hitCount` property of the `Count` managed bean. Here is the content of `template-client.xhtml`:

```

<?xml version='1.0' encoding='UTF-8' ?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:ui="http://java.sun.com/jsf/facelets"
    xmlns:h="http://java.sun.com/jsf/html">
  <body>

    This text above will not be displayed.

    <ui:composition template="/template.xhtml">

      This text will not be displayed.

      <ui:define name="title">
        This page has been accessed #{count.hitCount} time(s).
      </ui:define>

      This text will also not be displayed.

      <ui:define name="body">
        Hooray!
      </ui:define>

      This text will not be displayed.

    </ui:composition>

    This text below will also not be displayed.

  </body>
</html>

```

Follow these instructions to build, package, and deploy the `cart` example to your Application Server instance using the NetBeans IDE IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to `tut-install/examples/ejb/`.
3. Select the `cart` folder.
4. Select the Open as Main Project and Open Required Projects check boxes.
5. Click Open Project Folder.
6. In the Projects tab, right-click the `cart` project and select Deploy Project.

## Building, Deploying, and Running the counter Example

The `counter` example application can be built, deployed, and run using NetBeans IDE or Ant.

## ▼ Building, Deploying, and Running the counter Example in NetBeans IDE

- 1 In NetBeans IDE, select File→Open Project.
- 2 In the Open Project dialog, navigate to *tut-install/examples/ejb/*.
- 3 Select the counter folder.
- 4 Select the Open as Main Project check box.
- 5 Click Open Project Folder.
- 6 In the Projects tab, right-click the counter project and select Run.  
A web browser will open the URL `http://localhost:8080/counter` that displays the number of hits.
- 7 Click the browser's Refresh button to see the hit count increment.

## ▼ Building, Deploying, and Running the counter Example Using Ant

- 1 In a terminal, navigate to *tut-install/examples/ejb/counter*.
- 2 Enter the following command:  

```
ant all
```

This will build and deploy counter to your Enterprise Server instance.
- 3 In a web browser, enter the following URL: `http://localhost:8080/counter`.
- 4 Click the browser's Refresh button to see the hit count increment.

## A Web Service Example: helloService

This example demonstrates a simple web service that generates a response based on information received from the client. `HelloServiceBean` is a stateless session bean that implements a single method, `sayHello`. This method matches the `sayHello` method invoked by the client described in “[A Simple JAX-WS Client](#)” on page 242.

## The Web Service Endpoint Implementation Class

HelloServiceBean is the endpoint implementation class. The endpoint implementation class is typically the primary programming artifact for enterprise bean web service endpoints. The web service endpoint implementation class has the following requirements:

- The class must be annotated with either the `javax.jws.WebService` or `javax.jws.WebServiceProvider` annotations.
- The implementing class may explicitly reference an SEI through the `endpointInterface` element of the `@WebService` annotation, but is not required to do so. If no `endpointInterface` is specified in `@WebService`, an SEI is implicitly defined for the implementing class.
- The business methods of the implementing class must be public, and must not be declared `static` or `final`.
- Business methods that are exposed to web service clients must be annotated with `javax.jws.WebMethod`.
- Business methods that are exposed to web service clients must have JAXB-compatible parameters and return types. See [JAXB default data type bindings \(http://java.sun.com/javaee/5/docs/tutorial/doc/bnazq.html#bnazs\)](http://java.sun.com/javaee/5/docs/tutorial/doc/bnazq.html#bnazs).
- The implementing class must not be declared `final` and must not be `abstract`.
- The implementing class must have a default public constructor.
- The endpoint class must be annotated `@Stateless`.
- The implementing class must not define the `finalize` method.
- The implementing class may use the `javax.annotation.PostConstruct` or `javax.annotation.PreDestroy` annotations on its methods for lifecycle event callbacks.
 

The `@PostConstruct` method is called by the container before the implementing class begins responding to web service clients.

The `@PreDestroy` method is called by the container before the endpoint is removed from operation.

## Stateless Session Bean Implementation Class

The HelloServiceBean class implements the `sayHello` method, which is annotated `@WebMethod`. The source code for the HelloServiceBean class follows:

```
package com.sun.tutorial.javaee.ejb;

import javax.ejb.Stateless;
import javax.jws.WebMethod;
import javax.jws.WebService;
```

```
@Stateless
@WebService
public class HelloServiceBean {
    private String message = "Hello, ";

    public void HelloServiceBean() {}

    @WebMethod
    public String sayHello(String name) {
        return message + name + ".";
    }
}
```

## Building, Packaging, Deploying, and Testing the helloworldservice Example

You can build, package, and deploy the helloworldservice example using either NetBeans IDE or Ant. You can then use the Admin Console to test the web service endpoint methods.

### Building, Packaging, and Deploying the helloworldservice Example Using NetBeans IDE

Follow these instructions to build, package, and deploy the helloworldservice example to your Application Server instance using the NetBeans IDE IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to *tut-install/examples/ejb/*.
3. Select the helloworldservice folder.
4. Select the Open as Main Project and Open Required Projects check boxes.
5. Click Open Project Folder.
6. In the Projects tab, right-click the helloworldservice project and select Deploy Project.

This builds and packages to application into helloworldservice.ear, located in *tut-install/examples/ejb/helloworldservice/dist*, and deploys this ear file to your Application Server instance.

## Building, Packaging, and Deploying the hello-service Example Using Ant

Follow these instructions to build, package, and deploy the hello-service example to your Application Server instance using Ant.

1. In a terminal window, go to the `tut-install/examples/ejb/hello-service/` directory.
2. To build hello-service, type the following command:

```
ant
```

This runs the default task, which compiles the source files and packages the application into a JAR file located at `tut-install/examples/ejb/hello-service/dist/hello-service.jar`.

3. To deploy hello-service, type the following command:

```
ant deploy
```

Upon deployment, the Application Server generates additional artifacts required for web service invocation, including the WSDL file.

## Testing the Service without a Client

1. The Application Server Admin Console allows you to test the methods of a web service endpoint. To test the `sayHello` method of `HelloServiceBean`, do the following: Open the Admin Console by opening the following URL in a web browser:

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

2. Enter the admin username and password to log in to the Admin Console.
3. Click Web Services in the left pane of the Admin Console.
4. Click `hello-service`.
5. Click Test.
6. Under Methods, enter a name as the parameter to the `sayHello` method.
7. Click the `sayHello` button.

This will take you to the `sayHello` Method invocation page.

8. Under Method returned, you'll see the response from the endpoint.

## Using the Timer Service

Applications that model business work flows often rely on timed notifications. The timer service of the enterprise bean container enables you to schedule timed notifications for all types of enterprise beans except for stateful session beans. You can schedule a timed notification to occur according to a calendar schedule, at a specific time, after a duration of time, or at timed intervals. For example, you could set timers to go off at 10:30 AM on May 23, in 30 days, or every 12 hours.

There are two types of enterprise bean timers: *programmable timers* and *automatic timers*. Programmable timers are set by explicitly calling one of the timer creation methods of the `TimerService` interface. Automatic timers are created upon the successful deployment of an enterprise bean that contains a method annotated with the `java.ejb.Schedule` or `java.ejb.Schedules` annotations.

## Creating Calendar-Based Timer Expressions

Timers can be set according to a calendar-based schedule, expressed using a syntax similar to the UNIX `cron` utility. Both programmable and automatic timers can use calendar-based timer expressions.

TABLE 18-1 Calendar-Based Timer Attributes

Attribute	Description	Allowable Values	Default Value	Examples
second	One or more seconds within a minute.	0 to 59	0	second="30"
minute	One or more minutes within an hour.	0 to 59	0	minute="15"
hour	One or more hours within a day.	0 to 23	0	hour="13"
dayOfWeek	One or more days within a week.	0 to 7 <sup>1</sup> Sun, Mon, Tue, Wed, Thu, Fri, Sat	*	dayOfWeek="3" dayOfWeek="Mon"

<sup>1</sup> Both 0 and 7 refer to Sunday.

TABLE 18-1 Calendar-Based Timer Attributes (Continued)

Attribute	Description	Allowable Values	Default Value	Examples
dayOfMonth	One or more days within a month.	1 to 31 - 7 to -1 <sup>2</sup> Last [1st, 2nd, 3rd, 4th, 5th, Last] [Sun, Mon, Tue, Wed, Thu, Fri, Sat]	*	dayOfMonth="15" dayOfMonth="- 3" dayOfMonth="Last" dayOfMonth="2nd Fri"
month	One or more months within a year.	1 to 12 Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec	*	month="7" month="July"
year	A particular calendar year.	A four-digit calendar year.	*	year="2010"

<sup>2</sup> A negative number means the xth day or days before the end of the month.

## Specifying Multiple Values in Calendar Expressions

You can specify multiple values in calendar expressions in the following ways:

- “Using Wildcards in Calendar Expressions” on page 329
- “Specifying a List of Values” on page 329
- “Specifying a Range of Values” on page 330
- “Specifying Intervals” on page 330

## Using Wildcards in Calendar Expressions

Setting an attribute to an asterisk symbol (\*) represents all allowable values for the attribute.

### EXAMPLE 18-9 Calendar Expressions with Wildcards

The following expression represents every minute:

```
minute="*"
```

The following expression represents every day of the week:

```
dayOfWeek="*"
```

## Specifying a List of Values

To specify two or more values for an attribute, use a comma (,) to separate the values. A range of values are allowed as part of a list. Wildcards and intervals, however, are not allowed.

Duplicates within a list are ignored.

**EXAMPLE 18-10** Calendar Expressions with a List of Values

The following expression sets the day of the week to Tuesday and Thursday:

```
dayOfWeek="Tue, Thu"
```

The following expression represents 4:00 AM, every hour from 9:00 AM to 5:00 PM using a range, and 10:00 PM:

```
hour="4,9-17,20"
```

## Specifying a Range of Values

Use a dash character (-) to specify an inclusive range of values for an attribute. Members of a range cannot be wildcards, lists, or intervals. If the range is of the form  $x-x$ , it is equivalent to the single-valued expression  $x$ . If the range is of the form  $x-y$  and  $x$  is greater than  $y$ , it is equivalent to the expression  $x$  - *maximum value*, *minimum value* -  $y$ . That is, the expression begins at  $x$ , rolls-over to the beginning of the allowable values, and continues up to  $y$ .

**EXAMPLE 18-11** Calendar Expressions Using Ranges

The following expression represents 9:00 AM to 5:00 PM:

```
hour="9-17"
```

The following expression represents Friday through Monday:

```
dayOfWeek="5-1"
```

The following expression represents the 25th day of the month to the end of the month, and the beginning of the month to the 5th day of the month:

```
dayOfMonth="25-5"
```

It is equivalent to the following expression:

```
dayOfMonth="25-Last,1-5"
```

## Specifying Intervals

The forward slash (/) constrains an attribute to a starting point and an interval. It is used to specify every  $N$  seconds, minutes, or hours within the minute, hour or day. For an expression of the form  $x/y$ ,  $x$  represents the starting point and  $y$  represents the interval. The wildcard character may be used in the  $x$  position of an interval, and is equivalent to setting  $x$  to  $\emptyset$ .

Intervals may only be set for second, minute, and hour attributes.

**EXAMPLE 18-12** Calendar Expressions Using Intervals

The following expression represents every 10 minutes within the hour:

```
minute="*/10"
```

It is equivalent to:

```
minute="0,10,20,30,40,50"
```

The following expression represents every two hours starting at noon:

```
hour="12/2"
```

## Programmatic Timers

When a programmatic timer expires (goes off), the container calls the method annotated `@Timeout` in the bean's implementation class. The `@Timeout` method contains the business logic that handles the timed event.

### The Timeout Method

Methods annotated `@Timeout` in the enterprise bean class must return `void` and optionally take a `javax.ejb.Timer` object as the only parameter. They may not throw application exceptions.

```
@Timeout
public void timeout(Timer timer) {
    System.out.println("TimerBean: timeout occurred");
}
```

### Creating Programmatic Timers

To create a timer, the bean invokes one of the create methods of the `TimerService` interface. The create methods of `TimerService` allow single-action, interval, or calendar-based timers to be created.

For single-action or interval timers, the expiration of the timer can be expressed either as a duration or an absolute time. The duration is expressed as a the number of milliseconds before a timeout event is triggered. To specify an absolute time, create a `java.util.Date` object and pass it to either the `TimerService.createSingleActionTimer` or `TimerService.createTimer` methods.

**EXAMPLE 18-13** Setting a Programmatic Timer Based On a Duration

The following code sets a programmatic timer that will expire in one minute (6000 milliseconds):

**EXAMPLE 18-13** Setting a Programmatic Timer Based On a Duration *(Continued)*

```
long duration = 6000;  
Timer timer = timerService.createSingleActionTimer(duration, new TimerConfig());
```

**EXAMPLE 18-14** Setting a Programmatic Timer Based On an Absolute Time

The following code sets a programmatic timer that will expire at 12:05 PM on May 1st, 2010, specified as a `java.util.Date`:

```
SimpleDateFormat formatter = new SimpleDateFormat("MM/dd/yyyy 'at' HH:mm");  
Date date = formatter.parse("05/01/2010 at 12:05");  
Timer timer = timerService.createSingleActionTimer(date, new TimerConfig());
```

For calendar-based timers, the expiration of the timer is expressed as a `javax.ejb.ScheduleExpression` object, passed as a parameter to the `TimerService.createCalendarTimer` method. The `ScheduleExpression` class represents calendar-based timer expressions, and has methods that correspond to the attributes described in [“Creating Calendar-Based Timer Expressions” on page 328](#).

**EXAMPLE 18-15** Using `ScheduleExpression` to Set a Timer

The following code creates a programmatic timer using the `ScheduleExpression` helper class:

```
ScheduleExpression schedule = new ScheduleExpression();  
schedule.dayOfWeek("Mon");  
schedule.hour("12-17, 23");  
Timer timer = timerService.createCalendarTimer(schedule);
```

For details on the method signatures, see the `TimerService` API documentation at <http://java.sun.com/javaee/6/docs/api/javax/ejb/TimerService.html>.

The bean described in [“The `timersession` Example” on page 335](#) creates a timer as follows:

```
Timer timer = timerService.createTimer(intervalDuration,  
    "Created new programmatic timer");
```

In the `timersession` example, `createTimer` is invoked in a business method, which is called by a client.

Timers are persistent by default. If the server is shut down (or even crashes), persistent timers are saved and will become active again when the server is restarted. If a persistent timer expires while the server is down, the container will call the `@Timeout` method when the server is restarted.

Non-persistent programmatic timers are created by calling `TimerConfig.setPersistent(false)` and passing the `TimerConfig` object to one of the timer creation methods.

The `Date` and `long` parameters of the `createTimer` methods represent time with the resolution of milliseconds. However, because the timer service is not intended for real-time applications, a callback to the `@Timeout` method might not occur with millisecond precision. The timer service is for business applications, which typically measure time in hours, days, or longer durations.

## Automatic Timers

Automatic timers are created by the EJB container when an enterprise bean that contains methods annotated with the `@Schedule` or `@Schedules` annotations is deployed. An enterprise bean can have multiple automatic timeout methods, unlike a programmatic timer where there can only be one method annotated with the `@Timeout` annotation in the enterprise bean class.

Automatic timers can be configured through annotations or through the `ejb-jar.xml` deployment descriptor.

### The `@Schedule` and `@Schedules` Annotations

Adding a `@Schedule` annotation on an enterprise bean marks that method as a timeout method according to the calendar schedule specified in the attributes of `@Schedule`.

The `@Schedule` annotation has elements that correspond to the calendar expressions detailed in [“Creating Calendar-Based Timer Expressions” on page 328](#) and the `persistent`, `info`, and `timezone` elements.

The optional `persistent` element takes a boolean value, and is used to specify whether the automatic timer should survive a server restart or crash. By default, all automatic timers are persistent.

The optional `timezone` element is used to optionally specify that the automatic timer is associated with a particular time zone. If set, this element will evaluate all timer expressions in relation to the specified time zone, regardless of the time zone in which the EJB container is running. By default, all automatic timers set are in relation to the default time zone of the server.

The optional `info` element is used to set an informational description of the timer. A timer's information can be retrieved later using `Timer.getInfo`.

#### EXAMPLE 18-16 Setting an Automatic Timer Using `@Schedule`

The following timeout method uses `@Schedule` to set a timer that will expire every Sunday at midnight:

```
@Schedule(dayOfWeek="Sun", hour="0")
public void cleanupWeekData() { ... }
```

The `@Schedules` annotation is used to specify multiple calendar-based timer expressions for a given timeout method.

**EXAMPLE 18-17** Setting Multiple Automatic Timers for a Timeout Method Using `@Schedules`

The following timeout method uses the `@Schedules` annotation to set multiple calendar-based timer expressions. The first expression sets a timer to expire on the last day of every month. The second expression sets a timer to expire every Friday at 11:00 PM.

```
@Schedules ({
    @Schedule(dayOfMonth="Last"),
    @Schedule(dayOfWeek="Fri", hour="23")
})
public void doPeriodicCleanup() { ... }
```

## Canceling and Saving Timers

Timers can be canceled by the following events:

- When a single-event timer expires, the EJB container calls the associated timeout method and then cancels the timer.
- When the bean invokes the `cancel` method of the `Timer` interface, the container cancels the timer.

If a method is invoked on a canceled timer, the container throws the `javax.ejb.NoSuchObjectLocalException`.

To save a `Timer` object for future reference, invoke its `getHandle` method and store the `TimerHandle` object in a database. (A `TimerHandle` object is serializable.) To re-instantiate the `Timer` object, retrieve the handle from the database and invoke `getTimer` on the handle. A `TimerHandle` object cannot be passed as an argument of a method defined in a remote or web service interface. In other words, remote clients and web service clients cannot access a bean's `TimerHandle` object. Local clients, however, do not have this restriction.

## Getting Timer Information

In addition to defining the `cancel` and `getHandle` methods, the `Timer` interface defines methods for obtaining information about timers:

```
public long getTimeRemaining();
public java.util.Date getNextTimeout();
public java.io.Serializable getInfo();
```

The `getInfo` method returns the object that was the last parameter of the `createTimer` invocation. For example, in the `createTimer` code snippet of the preceding section, this information parameter is a `String` object with the value `created timer`.

To retrieve all of a bean's active timers, call the `getTimers` method of the `TimerService` interface. The `getTimers` method returns a collection of `Timer` objects.

## Transactions and Timers

An enterprise bean usually creates a timer within a transaction. If this transaction is rolled back, the timer creation is also rolled back. Similarly, if a bean cancels a timer within a transaction that gets rolled back, the timer cancellation is rolled back. In this case, the timer's duration is reset as if the cancellation had never occurred.

In beans that use container-managed transactions, the `@Timeout` method usually has the `Required` or `RequiresNew` transaction attribute to preserve transaction integrity. With these attributes, the EJB container begins the new transaction before calling the `@Timeout` method. If the transaction is rolled back, the container will call the `@Timeout` method at least one more time.

## The `timersession` Example

The source code for this example is in the `tut-install/examples/ejb/timersession/src/java/` directory.

`TimerSessionBean` is a singleton session bean that shows how to set both an automatic timer and a programmatic timer. In the source code listing of `TimerSessionBean` that follows, the `setTimer` and `@Timeout` methods are used to set a programmatic timer. A `TimerService` instance is injected by the container when the bean is created. Because it's a business method, `setTimer` is exposed to the local, no-interface view of `TimerSessionBean` and can be invoked by the client. In this example, the client invokes `setTimer` with an interval duration of 30,000 milliseconds. The `setTimer` method creates a new timer by invoking the `createTimer` method of `TimerService`. Now that the timer is set, the EJB container will invoke the `programmaticTimeout` method of `TimerSessionBean` when the timer expires, in about 30 seconds.

```
...
    public void setTimer(long intervalDuration) {
        logger.info("Setting a programmatic timeout for " +
            intervalDuration + " milliseconds from now.");
        Timer timer = timerService.createTimer(intervalDuration,
            "Created new programmatic timer");
    }

    @Timeout
    public void programmaticTimeout(Timer timer) {
        this.setLastProgrammaticTimeout(new Date());
        logger.info("Programmatic timeout occurred.");
    }
}
```

```
    }  
    ...
```

TimerSessionBean also has an automatic timer and timeout method, `automaticTimeout`. The automatic timer is set to expire every 3 minutes, and is set using a calendar-based timer expression in the `@Schedule` annotation.

```
    ...  
    @Schedule(minute="*/3", hour="*")  
    public void automaticTimeout() {  
        this.setLastAutomaticTimeout(new Date());  
        logger.info("Automatic timeout occurred");  
    }  
    ...
```

TimerSessionBean also has two business methods, `getLastProgrammaticTimeout` and `getLastAutomaticTimeout`. Clients call these methods to get the date and time of the last timeout for the programmatic timer and automatic timer, respectively.

Here's the source code for the TimerSessionBean class:

```
package timersession.ejb;  
  
import java.util.Date;  
import java.util.logging.Logger;  
import javax.annotation.Resource;  
import javax.ejb.Schedule;  
import javax.ejb.Stateless;  
import javax.ejb.Timeout;  
import javax.ejb.Timer;  
import javax.ejb.TimerService;  
  
@Singleton  
public class TimerSessionBean {  
    @Resource  
    TimerService timerService;  
  
    private Date lastProgrammaticTimeout;  
    private Date lastAutomaticTimeout;  
  
    private Logger logger = Logger  
        .getLogger("com.sun.tutorial.javaee.ejb.timersession.TimerSessionBean");  
  
    public void setTimer(long intervalDuration) {  
        logger.info("Setting a programmatic timeout for " +  
            intervalDuration + " milliseconds from now.");  
        Timer timer = timerService.createTimer(intervalDuration,  
            ...
```

```

        "Created new programmatic timer");
    }

    @Timeout
    public void programmaticTimeout(Timer timer) {
        this.setLastProgrammaticTimeout(new Date());
        logger.info("Programmatic timeout occurred.");
    }

    @Schedule(minute="*/3", hour="*")
    public void automaticTimeout() {
        this.setLastAutomaticTimeout(new Date());
        logger.info("Automatic timeout occurred.");
    }

    public String getLastProgrammaticTimeout() {
        if (lastProgrammaticTimeout != null) {
            return lastProgrammaticTimeout.toString();
        } else {
            return "never";
        }
    }

    public void setLastProgrammaticTimeout(Date lastTimeout) {
        this.lastProgrammaticTimeout = lastTimeout;
    }

    public String getLastAutomaticTimeout() {
        if (lastAutomaticTimeout != null) {
            return lastAutomaticTimeout.toString();
        } else {
            return "never";
        }
    }

    public void setLastAutomaticTimeout(Date lastAutomaticTimeout) {
        this.lastAutomaticTimeout = lastAutomaticTimeout;
    }
}

```

---

**Note** – Enterprise Server has a default minimum timeout value of 1000 milliseconds, or 1 second. If you need to set the timeout value lower than 1000 milliseconds, change the value of the `minimum-delivery-interval-in-millis` element in `domain-dir/config/domain.xml`. Due to virtual machine constraints, the lowest practical value for `minimum-delivery-interval-in-millis` is around 10 milliseconds.

---

## Building, Packaging, Deploying, and Running the `timersession` Example

You can build, package, deploy, and run the `timersession` example using either NetBeans IDE or Ant.

### ▼ Building, Packaging, Deploying, and Running the `timersession` Example Using NetBeans IDE

Follow these instructions to build, package, and deploy the `timersession` example to your Enterprise Server instance using the NetBeans IDE IDE.

- 1 In NetBeans IDE, select **File**→**Open Project**.
- 2 In the **Open Project** dialog, navigate to `tut-install/examples/ejb/`.
- 3 Select the `timersession` folder.
- 4 Select the **Open as Main Project** check box.
- 5 Click **Open Project Folder**.
- 6 Select **Run**→**Run Main Project**.

This builds and packages the application into `timersession.war`, located in `tut-install/examples/ejb/timersession/dist/`, deploys this WAR file to your Enterprise Server instance, and then runs the web client.

### ▼ Building, Packaging, and Deploying the `timersession` Example Using Ant

Follow these instructions to build, package, and deploy the `timersession` example to your Application Server instance using Ant.

- 1 In a terminal window, go to the `tut-install/examples/ejb/timersession/` directory.
- 2 To build `timersession`, type the following command:

```
ant build
```

This runs the default task, which compiles the source files and packages the application into a WAR file located at `tut-install/examples/ejb/timersession/dist/timersession.war`.

- 3 To deploy the application, type the following command:

```
ant deploy
```

## ▼ Running the Web Client

- 1 **Open a web browser to `http://localhost:8080/timersession`.**
- 2 **Click the Set Timer button to set a programmatic timer.**
- 3 **Wait for a while and click the browser's Refresh button.**

You will see the date and time of the last programmatic and automatic timeouts.

You can also see the messages that are logged when a timeout occurs by opening the `server.log` file located in `domain-dir/server/logs/`.

## Handling Exceptions

The exceptions thrown by enterprise beans fall into two categories: system and application.

A *system exception* indicates a problem with the services that support an application. Examples of these problems include the following: a connection to an external resource cannot be obtained or an injected resource cannot be found. If your enterprise bean encounters a system-level problem, it should throw a `javax.ejb.EJBException`. Because the `EJBException` is a subclass of the `RuntimeException`, you do not have to specify it in the `throws` clause of the method declaration. If a system exception is thrown, the EJB container might destroy the bean instance. Therefore, a system exception cannot be handled by the bean's client program; it requires intervention by a system administrator.

An *application exception* signals an error in the business logic of an enterprise bean. Application exceptions are typically exceptions that you've coded yourself, such as the `BookException` thrown by the business methods of the `CartBean` example. When an enterprise bean throws an application exception, the container does not wrap it in another exception. The client should be able to handle any application exception it receives.

If a system exception occurs within a transaction, the EJB container rolls back the transaction. However, if an application exception is thrown within a transaction, the container does not roll back the transaction.



# A Message-Driven Bean Example

---

Message-driven beans can implement any messaging type. Most commonly, they implement the Java Message Service (JMS) technology. The example in this chapter uses JMS technology, so you should be familiar with basic JMS concepts such as queues and messages. To learn about these concepts, see [Chapter 33, “Java Message Service Concepts.”](#)

This chapter describes the source code of a simple message-driven bean example. Before proceeding, you should read the basic conceptual information in the section [“What Is a Message-Driven Bean?”](#) on page 285 as well as [“Using Message-Driven Beans to Receive Messages Asynchronously”](#) on page 640 in Chapter 33, [“Java Message Service Concepts.”](#)

## simplemessage **Example Application Overview**

The simplemessage application has the following components:

- SimpleMessageClient: An application client that sends several messages to a queue
- SimpleMessageEJB: A message-driven bean that asynchronously receives and processes the messages that are sent to the queue

[Figure 19-1](#) illustrates the structure of this application. The application client sends messages to the queue, which was created administratively using the Administration Console. The JMS provider (in this case, the Enterprise Server) delivers the messages to the instances of the message-driven bean, which then processes the messages.

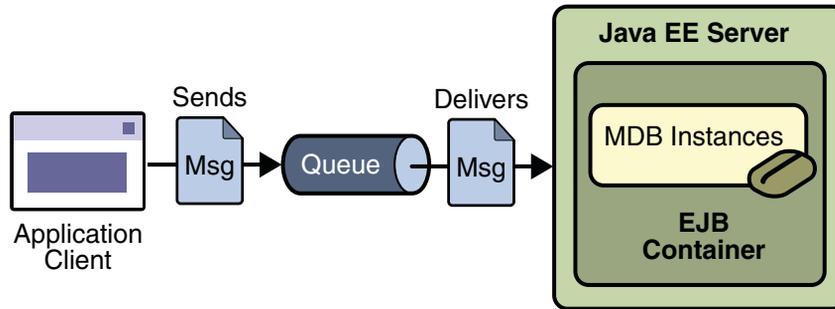


FIGURE 19-1 The simplemessage Application

The source code for this application is in the *tut-install/examples/ejb/simplemessage/* directory.

## The simplemessage Application Client

The `SimpleMessageClient` sends messages to the queue that the `SimpleMessageBean` listens to. The client starts by injecting the connection factory and queue resources:

```

@Resource(mappedName="jms/ConnectionFactory")
private static ConnectionFactory connectionFactory;

@Resource(mappedName="jms/Queue")
private static Queue queue;
  
```

Next, the client creates the connection, session, and message producer:

```

connection = connectionFactory.createConnection();
session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE);
messageProducer = session.createProducer(queue);
  
```

Finally, the client sends several messages to the queue:

```

message = session.createTextMessage();

for (int i = 0; i < NUM_MSGS; i++) {
    message.setText("This is message " + (i + 1));
    System.out.println("Sending message: " + message.getText());
    messageProducer.send(message);
}
  
```

## The Message-Driven Bean Class

The code for the `SimpleMessageBean` class illustrates the requirements of a message-driven bean class:

- It must be annotated with the `@MessageDriven` annotation if it does not use a deployment descriptor.
- The class must be defined as `public`.
- The class cannot be defined as `abstract` or `final`.
- It must contain a public constructor with no arguments.
- It must not define the `finalize` method.

It is recommended, but not required, that a message-driven bean class implement the message listener interface for the message type it supports. A bean that supports the JMS API implements the `javax.jms.MessageListener` interface.

Unlike session beans and entities, message-driven beans do not have the remote or local interfaces that define client access. Client components do not locate message-driven beans and invoke methods on them. Although message-driven beans do not have business methods, they may contain helper methods that are invoked internally by the `onMessage` method.

For the Enterprise Server, the `@MessageDriven` annotation typically contains a `mappedName` element that specifies the JNDI name of the destination from which the bean will consume messages. For complex message-driven beans, there can also be an `activationconfig` element containing `@ActivationConfigProperty` annotations used by the bean.

A message-driven bean can also inject a `MessageDrivenContext` resource. Commonly you use this resource to call the `setRollbackOnly` method to handle exceptions for a bean that uses container-managed transactions.

Therefore, the first few lines of the `SimpleMessageBean` class look like this:

```
@MessageDriven(mappedName="jms/Queue", activationConfig = {
    @ActivationConfigProperty(propertyName = "acknowledgeMode",
        propertyValue = "Auto-acknowledge"),
    @ActivationConfigProperty(propertyName = "destinationType",
        propertyValue = "javax.jms.Queue")
})
public class SimpleMessageBean implements MessageListener {
    @Resource
    private MessageDrivenContext mdc;
    ...
}
```

NetBeans IDE typically creates a message-driven bean with a default set of `@ActivationConfigProperty` settings. You can delete those you do not need, or add others. [Table 19–1](#) lists commonly used properties.

TABLE 19-1 @ActivationConfigProperty Settings for Message-Driven Beans

Property Name	Description
acknowledgeMode	Acknowledgment mode; see <a href="#">“Controlling Message Acknowledgment” on page 629</a> for information
destinationType	Either <code>javax.jms.Queue</code> or <code>javax.jms.Topic</code>
subscriptionDurability	For durable subscribers, set to <code>Durable</code> ; see <a href="#">“Creating Durable Subscriptions” on page 633</a> for information
clientId	For durable subscribers, the client ID for the connection
subscriptionName	For durable subscribers, the name of the subscription
messageSelector	A string that filters messages; see <a href="#">“JMS Message Selectors” on page 624</a> for information, and see <a href="#">“An Application That Uses the JMS API with a Session Bean” on page 680</a> for an example
addressList	Remote system or systems to communicate with; see <a href="#">“An Application Example That Consumes Messages from a Remote Server” on page 695</a> for an example

## The onMessage Method

When the queue receives a message, the EJB container invokes the message listener method or methods. For a bean that uses JMS, this is the `onMessage` method of the `MessageListener` interface.

A message listener method must follow these rules:

- The method must be declared as `public`.
- The method must not be declared as `final` or `static`.

The `onMessage` method is called by the bean’s container when a message has arrived for the bean to service. This method contains the business logic that handles the processing of the message. It is the message-driven bean’s responsibility to parse the message and perform the necessary business logic.

The `onMessage` method has a single argument: the incoming message.

The signature of the `onMessage` method must follow these rules:

- The return type must be `void`.
- The method must have a single argument of type `javax.jms.Message`.

In the `SimpleMessageBean` class, the `onMessage` method casts the incoming message to a `TextMessage` and displays the text:

```
public void onMessage(Message inMessage) {
    TextMessage msg = null;

    try {
        if (inMessage instanceof TextMessage) {
            msg = (TextMessage) inMessage;
            logger.info("MESSAGE BEAN: Message received: " +
                msg.getText());
        } else {
            logger.warning("Message of wrong type: " +
                inMessage.getClass().getName());
        }
    } catch (JMSEException e) {
        e.printStackTrace();
        mdc.setRollbackOnly();
    } catch (Throwable te) {
        te.printStackTrace();
    }
}
```

## Packaging, Deploying, and Running the simpleMessage Example

To package, deploy and run this example, go to the *tut-install/examples/ejb/simplemessage/* directory.

### Creating the Administered Objects for the simpleMessage Example

This example requires the following:

- A JMS connection factory resource
- A JMS destination resource

If you have run the simple JMS examples in [Chapter 33, “Java Message Service Concepts,”](#) and have not deleted the resources, you already have these resources and do not need to perform these steps.

You can use Ant targets to create the resources. The Ant targets, which are defined in the `build.xml` file for this example, use the `asadmin` command. To create the resources needed for this example, use the following commands:

```
ant create-cf
ant create-queue
```

These commands do the following:

- Create a connection factory resource named `.jms/ConnectionFactory`
- Create a destination resource named `.jms/Queue`

The Ant targets for these commands refer to other targets that are defined in the `tut-install/examples/bp-project/app-server-ant.xml` file.

## Building, Deploying, and Running the `simplemessage` Application Using NetBeans IDE

To build, deploy, and run the application using NetBeans IDE, do the following:

1. In NetBeans IDE, choose Open Project from the File menu.
2. In the Open Project dialog, navigate to `tut-install/examples/ejb/`.
3. Select the `simplemessage` folder.
4. Select the Open as Main Project check box and the Open Required Projects check box.
5. Click Open Project.
6. Right-click the `simplemessage` project and choose Build.

This task packages the application client and the message-driven bean, then creates a file named `simplemessage.ear` in the `dist` directory.

7. Right-click the project and choose Run.

This command deploys the project, returns a JAR file named `simplemessageClient.jar`, and then executes it.

The output of the application client in the Output pane looks like this (preceded by application client container output):

```
Sending message: This is message 1
Sending message: This is message 2
Sending message: This is message 3
To see if the bean received the messages,
  check <install_dir>/domains/domain1/logs/server.log.
```

The output from the message-driven bean appears in the server log (`domain-dir/logs/server.log`), wrapped in logging information.

```
MESSAGE BEAN: Message received: This is message 1
MESSAGE BEAN: Message received: This is message 2
MESSAGE BEAN: Message received: This is message 3
```

The received messages may appear in a different order from the order in which they were sent.

Undeploy the application after you finish running the client. To undeploy the application, follow these steps:

1. Click the Services tab.
2. Expand the Servers node.
3. Expand the GlassFish v3 Domain node.
4. Expand the Applications node.

You may need to right-click the node and choose Refresh to see the contents of the Applications node.

5. Right-click `simplemessage` and choose Undeploy.

To remove the generated files, click the Projects tab, then right-click the `simplemessage` project and choose Clean.

## Building, Deploying, and Running the `simplemessage` Application Using Ant

To build the application using Ant, do the following:

1. Go to the following directory:

```
tut-install/examples/ejb/simplemessage/
```

2. To compile the source files and package the application, use the following command:

```
ant
```

This target packages the application client and the message-driven bean, then creates a file named `simplemessage.ear` in the `dist` directory.

By using resource injection and annotations, you avoid having to create deployment descriptor files for the message-driven bean and application client. You need to use deployment descriptors only if you want to override the values specified in the annotated source files.

To deploy the application and run the client using Ant, use the following command:

```
ant run
```

Ignore the message that states that the application is deployed at a URL.

The output in the terminal window looks like this (preceded by application client container output):

```
Sending message: This is message 1  
Sending message: This is message 2  
Sending message: This is message 3
```

To see if the bean received the messages, check `<install_dir>/domains/domain1/logs/server.log`.

In the server log file, the following lines appear, wrapped in logging information:

```
MESSAGE BEAN: Message received: This is message 1
MESSAGE BEAN: Message received: This is message 2
MESSAGE BEAN: Message received: This is message 3
```

The received messages may appear in a different order from the order in which they were sent.

Undeploy the application after you finish running the client. Use the following command:

```
ant undeploy
```

To remove the generated files, use the following command in the `simplemessage`, `simplemessage-app-client`, and `simplemessage-ejb` directories:

```
ant clean
```

## Removing the Administered Objects for the simplemessage Example

After you run the example, you can use the following Ant targets to delete the connection factory and queue:

```
ant delete-cf
ant delete-queue
```

PART V

Contexts and Dependency Injection for the  
Java EE Platform

Part Five explores Contexts and Dependency Injection for the Java EE Platform.



# Introduction to Contexts and Dependency Injection for the Java™ EE Platform

---

Contexts and Dependency Injection for the Java EE Platform (CDI) is one of several Java EE 6 features that help to knit together the web tier and the transactional tier of the Java EE platform. It is a set of services that, used together, makes it easy for developers to use enterprise beans along with JavaServer™ Faces technology in web applications. Designed for use with stateful objects, it also has many broader uses, allowing developers a great deal of flexibility to integrate different kinds of components in a loosely coupled but type-safe way.

CDI is specified by JSR-299, which was formerly known as Web Beans. Related specifications that CDI uses include the following:

- JSR-330, Dependency Injection for Java
- The Managed Beans specification that is an offshoot of the Java EE 6 platform specification (JSR-316)

This chapter covers the following topics.

- “Overview of Contexts and Dependency Injection for the Java EE Platform” on page 352
- “About Beans” on page 353
- “Beans as Injectable Objects” on page 354
- “Using Qualifiers” on page 355
- “Injecting Beans” on page 356
- “Using Scopes” on page 356
- “Giving Beans EL Names” on page 358
- “Adding Setter and Getter Methods” on page 358
- “Using a Managed Bean in a Facelets Page” on page 359
- “Injecting Objects by Using Producer Methods” on page 360
- “Configuring a CDI Application” on page 361
- “Further Information” on page 361

# Overview of Contexts and Dependency Injection for the Java EE Platform

The two most fundamental services provided by CDI are as follows:

- **Contexts:** The ability to bind the lifecycle and interactions of stateful components to well-defined but extensible *lifecycle contexts*
- **Dependency injection:** The ability to inject components into an application in a type-safe way, including the ability to choose at deployment time which implementation of a particular interface to inject

In addition, CDI provides the following services:

- Integration with the Unified Expression Language (EL), which allows any component to be used directly within a JavaServer Faces page (or a JavaServer Pages™ page)
- The ability to decorate injected components
- The ability to associate interceptors with components using type-safe interceptor bindings
- An event notification model
- A web conversation scope in addition to the three standard scopes (request, session, and application) defined by the Java Servlet specification
- A complete Service Provider Interface (SPI) that allows third-party frameworks to integrate cleanly in the Java EE 6 environment

A major theme of CDI is loose coupling. CDI does the following:

- It decouples the server and the client by means of well-defined types and qualifiers, so that the server implementation may vary.
- It decouples the lifecycles of collaborating components by doing the following:
  - Making components contextual, with automatic lifecycle management
  - Allowing stateful components to interact like services, purely by message passing
- It completely decouples message producers from consumers, by means of events.
- It decouples orthogonal concerns by means of Java EE interceptors.

Along with loose coupling, CDI provides strong typing, as follows:

- It eliminates lookup using string-based names for wiring and correlations, so that the compiler will detect typing errors.
- It allows the use of declarative Java annotations to specify everything, largely eliminating the need for XML deployment descriptors, and making it easy to provide tools that introspect the code and understand the dependency structure at development time.

## About Beans

CDI redefines the concept of a *bean* beyond its use in other Java technologies, such as the JavaBeans™ and Enterprise JavaBeans™ (EJB™) technologies. In CDI, a bean is a source of contextual objects that defines application state and/or logic. A Java EE component is a bean if the lifecycle of its instances may be managed by the container according to the lifecycle context model defined in the CDI specification.

More specifically, a bean has the following attributes:

- A (nonempty) set of bean types
- A (nonempty) set of qualifiers (see “Using Qualifiers” on page 355)
- A scope (see “Using Scopes” on page 356)
- Optionally, a bean EL name (see “Giving Beans EL Names” on page 358)
- A set of interceptor bindings
- A bean implementation

A bean type defines a client-visible type of the bean. Almost any Java type may be a bean type of a bean:

- A bean type may be an interface, a concrete class or an abstract class, and may be declared `final` or have `final` methods.
- A bean type may be a parameterized type with actual type parameters and type variables.
- A bean type may be an array type. Two array types are considered identical only if the element type is identical.
- A bean type may be a primitive type. Primitive types are considered to be identical to their corresponding wrapper types in `java.lang`.
- A bean type may be a raw type.

## About Managed Beans

A *managed bean* is a bean that is implemented by a Java class, which is called its *bean class*. A top-level Java class is a managed bean if it is defined to be a managed bean by any other Java EE technology specification (for example, the JavaServer Faces technology specification), or if it meets all of the following conditions:

- It is not a non-static inner class.
- It is a concrete class, or is annotated `@Decorator`.
- It is not annotated with an EJB component-defining annotation or declared as an EJB bean class in `ejb-jar.xml`.
- It has an appropriate constructor. That is, one of the following is the case:
  - The class has a constructor with no parameters.

- The class declares a constructor annotated `@Inject`.

No special declaration, such as an annotation, is required to define a managed bean.

## Beans as Injectable Objects

The concept of *injection* has been part of Java technology for some time. Since the Java EE 5 platform was introduced, annotations have made it possible to inject resources and some other kinds of objects into container-managed objects. CDI makes it possible to inject more kinds of objects, and to inject them into objects that are not container-managed.

The following kinds of objects can be injected:

- (Almost) any Java class
- Session beans
- Java EE resources (data sources, Java Message Service topics, queues, and connection factories, and the like)
- Persistence contexts (JPA `EntityManager` objects)
- Producer fields, and objects returned by producer methods
- Web service references
- Remote enterprise bean references

For example, suppose you created a simple Java class with a method that returns a string:

```
package greetings;

public class Greeting {
    public String greet(String name) {
        return "Hello, " + name + ".";
    }
}
```

This class becomes a bean that you can then inject into another class. This bean is not exposed to the EL in this form. [“Giving Beans EL Names” on page 358](#) explains how you can make a bean accessible via the EL.

## Using Qualifiers

You can use qualifiers to provide different implementations of a particular bean type. A qualifier is an annotation that you apply to a bean. A qualifier type is a Java annotation defined as `@Target({METHOD, FIELD, PARAMETER, TYPE})` and `@Retention(RUNTIME)`.

For example, you could declare an `@Informal` qualifier type and apply it to another class that extends the `Greeting` class. To declare this qualifier type, you would use the following code:

```
package greetings;

import static java.lang.annotation.ElementType.FIELD;
import static java.lang.annotation.ElementType.METHOD;
import static java.lang.annotation.ElementType.PARAMETER;
import static java.lang.annotation.ElementType.TYPE;
import static java.lang.annotation.RetentionPolicy.RUNTIME;

import java.lang.annotation.Retention;
import java.lang.annotation.Target;

import javax.inject.Qualifier;

@Qualifier
@Retention(RUNTIME)
@Target({TYPE, METHOD, FIELD, PARAMETER})
public @interface Informal {}
```

You can then define a bean class that extends the `Greeting` class and uses this qualifier:

```
package greetings;

@Informal
public class InformalGreeting extends Greeting {
    public String greet(String name) {
        return "Hi, " + name + "!";
    }
}
```

Both implementations of the bean can now be used in the application.

If you define a bean with no qualifier, it automatically has the qualifier `@Default`. The unannotated `Greeting` class could be declared as follows:

```
package greetings;

import javax.enterprise.inject.Default;

@Default
```

```
public class Greeting {
    public String greet(String name) {
        return "Hello, " + name + ".";
    }
}
```

## Injecting Beans

In order to use the beans you create, you inject them into yet another bean that can then be used by an application, such as a JavaServer Faces application. For example, you might create a bean called `Printer` into which you would inject one of the `Greeting` beans:

```
import javax.inject.Inject;

public class Printer {

    @Inject Greeting greeting;
    ...
}
```

This code injects the `@DefaultGreeting` implementation into the bean, while the following code injects the `@Informal` implementation:

```
import javax.inject.Inject;

public class Printer {

    @Inject @Informal Greeting greeting;
    ...
}
```

More is needed for the complete picture of this bean. Its use of scope needs to be understood. In addition, for a JavaServer Faces application, the bean needs to be accessible through the unified expression language (unified EL).

## Using Scopes

For a web application to use a bean that injects another bean class, the bean needs to be able to hold state over the duration of the user's interaction with the application. The way to define this state is to give the bean a scope. You can give an object any of the scopes described in [Table 20-1](#), depending on how you are using it.

TABLE 20-1 Scopes

Scope	Annotation	Duration
Request scope	@RequestScoped	A user's interaction with a web application in a single HTTP request
Session scope	@SessionScoped	A user's interaction with a web application across multiple HTTP requests
Application scope	@ApplicationScoped	Shared state across all users' interactions with a web application
Dependent scope	@Dependent	The default scope if none is specified; it means that an object exists to serve exactly one client (bean), and has the same lifecycle as that client (bean)
Conversation scope	@ConversationScoped	A user's interaction with a JavaServer Faces application, within explicit boundaries controlled by the developer that extend the scope across multiple invocations of the JavaServer Faces life cycle. All long-running conversations are scoped to a particular HTTP servlet session and may not cross session boundaries.

The first three scopes were originally defined by the Java Servlet specification. The last two are defined by JSR-299.

You can also define and implement custom scopes, but that is an advanced topic. Custom scopes are likely to be used by those who implement and extend the CDI specification.

A scope gives an object a well-defined lifecycle context. A scoped object can be automatically created when it is needed and automatically destroyed when the context in which it was created ends. Moreover, its state is automatically shared by any clients that execute in the same context.

JavaBeans components, and Java EE components such as servlets and enterprise beans, do not by definition have a well-defined scope. These components are either:

- Singletons, such as Enterprise JavaBeans singleton beans, whose state is shared between all clients
- Stateless objects, such as servlets and stateless session beans, which do not contain client-visible state
- Objects that must be explicitly created and destroyed by their client, such as JavaBeans components and stateful session beans, whose state is shared by explicit reference passing between clients

If, however, you create a Java EE component that is a managed bean, it becomes a scoped object, which exists in a well-defined lifecycle context.

The web application for the `Printer` bean will use a simple request and response mechanism, so the managed bean can be annotated as follows:

```
import javax.inject.Inject;
import javax.enterprise.context.RequestScoped;

@RequestScoped
public class Printer {

    @Inject @Informal Greeting greeting;
    ...
}
```

Beans that use session, application, or conversation scope must be serializable, but beans that use request scope do not have to be serializable.

## Giving Beans EL Names

To make a bean accessible through the unified expression language (EL), use the `@Named` built-in qualifier.

```
import javax.inject.Inject;
import javax.enterprise.context.RequestScoped;
import javax.inject.Named;

@Named
@RequestScoped
public class Printer {

    @Inject @Informal Greeting greeting;
    ...
}
```

The `@Named` qualifier allows you to access the bean using the bean name, with the first letter in lowercase. For example, a Facelets page would refer to the bean as `printer`.

You can specify an argument to the `@Named` qualifier to use a non-default name:

```
@Named("MyPrinter")
```

With this annotation, the Facelets page would refer to the bean as `MyPrinter`.

## Adding Setter and Getter Methods

To make the state of the managed bean accessible, you need to add setter and getter methods for that state. The `createSalutation` method calls the bean's `greet` method, and the `getSalutation` method retrieves the result.

Once the setter and getter methods have been added, the bean is complete. The final code looks like this:

```
package greetings;

import javax.inject.Inject;
import javax.enterprise.context.RequestScoped;
import javax.inject.Named;

@Named
@RequestScoped
public class Printer {

    @Inject @Informal Greeting greeting;

    private String name;
    private String salutation;

    public void createSalutation() {
        this.salutation = greeting.greet(name);
    }

    public String getSalutation() {
        return salutation;
    }
    public String setName(String name) {
        this.name = name;
    }

    public String getName() {
        return name;
    }
}
```

## Using a Managed Bean in a Facelets Page

To use the managed bean in a Facelets page, you typically create a form that uses user interface elements to call its methods and display their results. This example provides a button that asks the user to type a name, retrieves the salutation, then displays the text in a paragraph below the button:

```
<h:form id="greetme">
  <p><h:outputLabel value="Enter your name: " for="name"/>
    <h:inputText id="name" value="#{printer.name}"/></p>
  <p><h:commandButton value="Say Hello" action="#{printer.createSalutation}"/></p>
  <p><h:outputText value="#{printer.salutation}"/></p>
</h:form>
```

## Injecting Objects by Using Producer Methods

Producer methods provide a way to inject objects that are not beans, objects whose values may vary at run time, and objects that require custom initialization.

For example, if you want to initialize a numeric value defined by a qualifier named `@MaxNumber`, you can define the value in a managed bean and then define a producer method, `getMaxNumber`, for it:

```
private int maxNumber = 100;
...
@Produces @MaxNumber int getMaxNumber() {
    return maxNumber;
}
```

When you inject the object in another managed bean, the container automatically invokes the producer method, initializing the value to 100:

```
@Inject @MaxNumber private int maxNumber;
```

If the value can vary at run time, the process is slightly different. For example, the following code defines a producer method that generates a random number defined by a qualifier called `@Random`:

```
private java.util.Random random =
    new java.util.Random( System.currentTimeMillis() );

java.util.Random getRandom() {
    return random;
}

@Produces @Random int next() {
    return getRandom().nextInt(maxNumber);
}
```

When you inject this object in another managed bean, you declare a *contextual instance* of the object:

```
@Inject @Random Instance<Integer> randomInt;
```

You then call the `get` method of the `Instance`:

```
this.number = randomInt.get();
```

## Configuring a CDI Application

An application that uses CDI must have a file named `beans.xml`. The file can be completely empty (it has content only in certain limited situations), but it must be present. For a web application, the `beans.xml` file can be in either the `WEB-INF` directory or the `WEB-INF/classes/META-INF` directory. For EJB modules or JAR files, the `beans.xml` file must be in the `META-INF` directory.

## Further Information

For more information about Contexts and Dependency Injection for the Java EE Platform, see:

- Contexts and Dependency Injection for the Java EE Platform specification:  
<http://jcp.org/en/jsr/detail?id=299>
- An introduction to Contexts and Dependency Injection for the Java EE Platform:  
<http://docs.jboss.org/webbeans/reference/1.0.0.PREVIEW1/en-US/html/>
- Dependency Injection for Java specification:  
<http://jcp.org/en/jsr/detail?id=330>



## Running the Basic Contexts and Dependency Injection Examples

---

This chapter describes in detail how to build and run simple examples that use Contexts and Dependency Injection for the Java™ EE Platform (CDI).

The examples are in the following directory:

*tut-install/examples/cdi/*

To build and run the examples, you will do the following:

1. Use NetBeans IDE or the Ant tool to compile and package the example.
2. Use NetBeans IDE or the Ant tool to deploy the example.
3. Run the example in a web browser.

Each example has a `build.xml` file that refers to files in the following directory:

*tut-install/examples/bp-project/*

See [Chapter 2, “Using the Tutorial Examples,”](#) for basic information on installing, building, and running the examples.

### The `simplegreeting` Example

The `simplegreeting` example illustrates some of the most basic features of CDI: scopes, qualifiers, bean injection, and accessing a managed bean in a JavaServer™ Faces application.

When you run the example, you click a button that presents either a formal or an informal greeting, depending on how you edited one of the classes.

The example includes four source files, a Facelets page and template, and configuration files.

## The simplegreeting Source Files

The four source files for the simplegreeting example include the following:

- The default Greeting class, shown in [“Beans as Injectable Objects” on page 354](#)
- The `@Informal` qualifier interface definition and the `InformalGreeting` class that implements the interface, both shown in [“Using Qualifiers” on page 355](#)
- The `Printer` managed bean class that injects one of the two interfaces, shown in full in [“Adding Setter and Getter Methods” on page 358](#)

The source files are located in the following directory:

*tut-install/examples/cdi/simplegreeting/src/java/greetings*

## The Facelets Template and Page

To use the managed bean in a simple Facelets application, you can use a very simple template file and `index.xhtml` page.

The template page, `template.xhtml`, looks like this:

```
<?xml version='1.0' encoding='UTF-8' ?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:h="http://java.sun.com/jsf/html"
    xmlns:ui="http://java.sun.com/jsf/facelets">
  <h:head>
    <meta http-equiv="Content-Type" content="text/html; charset=UTF-8"/>
    <link href="resources/css/default.css" rel="stylesheet" type="text/css"/>
    <title><ui:insert name="title">Default Title</ui:insert></title>
  </h:head>

  <body>
    <div id="container">
      <div id="header">
        <h2><ui:insert name="head">Head</ui:insert></h2>
      </div>

      <div id="space">
        <p></p>
      </div>

      <div id="content">
        <ui:insert name="content"/>
      </div>
    </div>
  </body>
</html>
```

```

    </div>
  </body>
</html>

```

To create the Facelets page, you can redefine the title and head, then add a small form to the content:

```

<?xml version='1.0' encoding='UTF-8' ?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:ui="http://java.sun.com/jsf/facelets"
    xmlns:h="http://java.sun.com/jsf/html"
    xmlns:f="http://java.sun.com/jsf/core">
  <ui:composition template="/template.xhtml">

    <ui:define name="title">Simple Greeting</ui:define>
    <ui:define name="head">Simple Greeting</ui:define>
    <ui:define name="content">
      <h:form id="greetme">
        <p><h:outputLabel value="Enter your name: " for="name"/>
          <h:inputText id="name" value="#{printer.name}"/></p>
        <p><h:commandButton value="Say Hello" action="#{printer.createSalutation}"/></p>
        <p><h:outputText value="#{printer.salutation}"/> </p>
      </h:form>
    </ui:define>

  </ui:composition>
</html>

```

The form asks the user to type a name. The button is labeled “Say Hello,” and the action defined for it is to call the `createSalutation` method of the `Printer` managed bean. This method in turn calls the `greet` method of the defined `Greeting` class.

The output text for the form is the value of the greeting returned by the setter method. Depending on whether the default or the `@Informal` version of the greeting is injected, this is either

Hello, *name*.

or

Hi, *name*!

where *name* is the name typed by the user.

The Facelets page and template are located in the following directory:

`tut-install/examples/cdi/simplegreeting/web`

The simple CSS file that is used by the Facelets page is in the following location:

*tut-install/examples/cdi/simplegreeting/web/resources/css/default.css*

## Configuration Files

You must create an empty `beans.xml` file to indicate to Sun GlassFish™ Enterprise Server v3 that your application is a CDI application. This file can have content in some situations, but not in simple applications like this one.

Your application also needs the basic web application deployment descriptors `web.xml` and `sun-web.xml`.

These configuration files are located in the following directory:

*tut-install/examples/cdi/simplegreeting/web/WEB-INF*

## Building, Packaging, Deploying, and Running the simplegreeting Example

You can build, package, deploy, and run the simplegreeting application using either NetBeans IDE or the Ant tool.

### ▼ Building and Deploying the simplegreeting Example Using NetBeans IDE

Follow these instructions to build, package, and deploy the simplegreeting example to the Enterprise Server using NetBeans IDE.

This procedure builds the application into the *tut-install/examples/cdi/simplegreeting/build/web* directory and deploys the contents of this directory to the Enterprise Server.

- 1 In NetBeans IDE, select **File**→**Open Project**.
- 2 In the **Open Project** dialog, navigate to *tut-install/examples/cdi/*.
- 3 Select the **simplegreeting** folder.
- 4 Select the **Open as Main Project** checkbox.
- 5 Click **Open Project Folder**.

**6 (Optional) To modify the `Printer.java` file, perform these steps:**

- a. Expand the **Source Packages** node.
- b. Expand the **greetings** node.
- c. **Double-click** the `Printer.java` file.
- d. In the edit pane, comment out the `@Informal` annotation:

```
//@Informal
@Inject
Greeting greeting;
```

- e. Save the file.

**7 In the Projects tab, right-click the `simplegreeting` project and select **Deploy**.****▼ Building, Packaging, and Deploying the `simplegreetingExample` Using Ant**

Follow these instructions to build, package, and deploy the `simplegreeting` example to the Enterprise Server using Ant.

This procedure builds the application into the `tut-install/examples/cdi/simplegreeting/build/web` directory and deploys the contents of this directory to the Enterprise Server.

This procedure builds and packages the application into `simplegreeting.war`, located in `tut-install/examples/cdi/simplegreeting/dist/`, and deploys this WAR file to the Enterprise Server.

**1 In a terminal window, go to this directory:**

```
tut-install/examples/cdi/simplegreeting/
```

**2 Type the following command:**

```
ant
```

This command calls the default target, which builds and packages the application into a WAR file, `simplegreeting.war`, located in the `dist` directory.

**3 Type the following command:**

```
ant deploy
```

`simplegreeting.war` will be deployed to the Enterprise Server.

## ▼ Running the `simplegreeting` Example

### 1 In a web browser, type the following URL:

`http://localhost:8080/simplegreeting`

The Simple Greeting page opens.

### 2 Type a name in the text field.

For example, suppose you type Duke.

### 3 Click the Say Hello button.

If you did not modify the `Printer.java` file, the following text string appears below the button:

Hi, Duke!

If you commented out the `@Informal` annotation in the `Printer.java` file, the following text string appears below the button:

Hello, Duke.

Figure 21–1 shows what the application looks like if you did not modify the `Printer.java` file.



FIGURE 21–1 Simple Greeting Application

## The guessnumber Example

The `guessnumber` example is somewhat more complex than the `simplegreeting` example. It illustrates the use of producer methods and of session and application scope.

The example is a game in which you try to guess a number in fewer than 10 attempts. It is similar to the `guessNumber` example described in [Chapter 5, “Introduction to Facelets,”](#) except that you can keep guessing until you get the right answer or until you use up your 10 attempts.

The example includes four source files, a Facelets page and template, and configuration files. The configuration files and the template are the same as those used for the `simplegreeting` example.

## The guessnumber Source Files

The four source files for the `guessnumber` example include the following:

- The `@MaxNumber` qualifier interface
- The `@Random` qualifier interface
- The `Generator` managed bean, which defines producer methods
- The `UserNumberBean` managed bean

The source files are located in the following directory:

```
tut-install/examples/cdi/guessnumber/src/java/guessnumber
```

### The `@MaxNumber` and `@Random` Qualifier Interfaces

The `@MaxNumber` qualifier interface is defined as follows.

```
package guessnumber;

import static java.lang.annotation.ElementType.FIELD;
import static java.lang.annotation.ElementType.METHOD;
import static java.lang.annotation.ElementType.PARAMETER;
import static java.lang.annotation.ElementType.TYPE;
import static java.lang.annotation.RetentionPolicy.RUNTIME;

import java.lang.annotation.Documented;
import java.lang.annotation.Retention;
import java.lang.annotation.Target;

import javax.inject.Qualifier;

@Target( { TYPE, METHOD, PARAMETER, FIELD })
@Retention(RUNTIME)
@Documented
@Qualifier
public @interface MaxNumber {

}
```

The `@Random` qualifier interface is defined as follows.

```
package guessnumber;

import static java.lang.annotation.ElementType.FIELD;
```

```
import static java.lang.annotation.ElementType.METHOD;
import static java.lang.annotation.ElementType.PARAMETER;
import static java.lang.annotation.ElementType.TYPE;
import static java.lang.annotation.RetentionPolicy.RUNTIME;

import java.lang.annotation.Documented;
import java.lang.annotation.Retention;
import java.lang.annotation.Target;

import javax.inject.Qualifier;

@Target( { TYPE, METHOD, PARAMETER, FIELD })
@Retention(RUNTIME)
@Documented
@Qualifier
public @interface Random {

}
```

## The Generator Managed Bean

The Generator managed bean contains the two producer methods for the application. The bean has the `@ApplicationScoped` annotation to specify that its context extends for the duration of the user's interaction with the application.

```
package guessnumber;

import java.io.Serializable;

import javax.enterprise.context.ApplicationScoped;
import javax.enterprise.inject.Produces;

@ApplicationScoped
public class Generator implements Serializable {

    private static final long serialVersionUID = -7213673465118041882L;

    private java.util.Random random =
        new java.util.Random( System.currentTimeMillis() );

    private int maxNumber = 100;

    java.util.Random getRandom() {
        return random;
    }

    @Produces @Random int next() {
        return getRandom().nextInt(maxNumber);
    }
}
```

```

    }

    @Produces @MaxNumber int getMaxNumber() {
        return maxNumber;
    }
}

```

## The UserNumberBean Managed Bean

The UserNumberBean managed bean, the backing bean for the JavaServer Faces application, provides the basic logic for the game. It implements setter and getter methods for the bean fields and injects the two qualifier objects. It provides a reset method that allows you to begin a new game after you complete one. It also provides a check method that determines whether the user has guessed the number, and a validateNumberRange method that determines whether the user's input is correct.

```

package guessnumber;

import java.io.Serializable;

import javax.annotation.PostConstruct;
import javax.enterprise.context.SessionScoped;
import javax.enterprise.inject.Instance;
import javax.inject.Inject;
import javax.inject.Named;
import javax.faces.application.FacesMessage;
import javax.faces.component.UIComponent;
import javax.faces.component.UIInput;
import javax.faces.context.FacesContext;

@Named
@SessionScoped
public class UserNumberBean implements Serializable {

    private static final long serialVersionUID = 1L;
    private int number;
    private Integer userNumber;
    private int minimum;
    private int remainingGuesses;

    @MaxNumber
    @Inject
    private int maxNumber;

    private int maximum;

    @Random

```

```
@Inject
Instance<Integer> randomInt;

public UserNumberBean() {
}

public int getNumber() {
    return number;
}

public void setUserNumber(Integer user_number) {
    userNumber = user_number;
}

public Integer getUserNumber() {
    return userNumber;
}

public int getMaximum() {
    return (this.maximum);
}

public void setMaximum(int maximum) {
    this.maximum = maximum;
}

public int getMinimum() {
    return (this.minimum);
}

public void setMinimum(int minimum) {
    this.minimum = minimum;
}

public int getRemainingGuesses() {
    return remainingGuesses;
}

public String check() throws InterruptedException {
    if (userNumber > number) {
        maximum = userNumber - 1;
    }
    if (userNumber < number) {
        minimum = userNumber + 1;
    }
    if (userNumber == number) {
        FacesContext.getCurrentInstance().addMessage(null,
            new FacesMessage("Correct!"));
    }
}
```

```

    }
    remainingGuesses--;
    return null;
}

@PostConstruct
public void reset() {
    this.minimum = 0;
    this.userNumber = 0;
    this.remainingGuesses = 10;
    this.maximum = maxNumber;
    this.number = randomInt.get();
}

public void validateNumberRange(FacesContext context, UIComponent toValidate,
                                Object value) {
    if (remainingGuesses <= 0) {
        FacesMessage message = new FacesMessage("No guesses left!");
        context.addMessage(toValidate.getClientId(context), message);
        ((UIInput) toValidate).setValid(false);
        return;
    }
    int input = (Integer) value;

    if (input < minimum || input > maximum) {
        ((UIInput) toValidate).setValid(false);

        FacesMessage message = new FacesMessage("Invalid guess");
        context.addMessage(toValidate.getClientId(context), message);
    }
}
}
}

```

## The Facelets Page

This example uses the same template that the `simplegreeting` example uses. The `index.xhtml` file, however, is more complex.

```

<?xml version='1.0' encoding='UTF-8' ?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
<html xmlns="http://www.w3.org/1999/xhtml"
    xmlns:ui="http://java.sun.com/jsf/facelets"
    xmlns:h="http://java.sun.com/jsf/html"
    xmlns:f="http://java.sun.com/jsf/core">
<ui:composition template="/template.xhtml">

```

```

<ui:define name="title">Guess My Number</ui:define>
<ui:define name="head">Guess My Number</ui:define>
<ui:define name="content">
  <h:form id="GuessMain">
    <div style="color: black; font-size: 24px;">
      <p>I'm thinking of a number between
      <span style="color: blue">#{userNumberBean.minimum}</span> and
      <span style="color: blue">#{userNumberBean.maximum}</span>. You have
      <span style="color: blue">#{userNumberBean.remainingGuesses}</span>
      guesses.</p>
    </div>
    <h:panelGrid border="0" columns="5" style="font-size: 18px;">
      Number:
      <h:inputText id="inputGuess"
        value="#{userNumberBean.userNumber}"
        required="true" size="3"
        disabled="#{userNumberBean.number eq userNumberBean.userNumber}"
        validator="#{userNumberBean.validateNumberRange}">
      </h:inputText>
      <h:commandButton id="GuessButton" value="Guess"
        action="#{userNumberBean.check}"
        disabled="#{userNumberBean.number eq userNumberBean.userNumber}"/>
      <h:commandButton id="RestartButton" value="Reset"
        action="#{userNumberBean.reset}"
        immediate="true" />
      <h:outputText id="Higher" value="Higher!"
        rendered="#{userNumberBean.number gt userNumberBean.userNumber and userNumberBean.userNumber ne 0}"
        style="color: red"/>
      <h:outputText id="Lower" value="Lower!"
        rendered="#{userNumberBean.number lt userNumberBean.userNumber and userNumberBean.userNumber ne 0}"
        style="color: red"/>
    </h:panelGrid>
    <div style="color: red; font-size: 14px;">
      <h:messages id="messages" globalOnly="false"/>
    </div>
  </h:form>
</ui:define>

</ui:composition>
</html>

```

The Facelets page presents the user with the minimum and maximum values and the number of guesses remaining. The user's interaction with the game takes place within the `panelGrid` table, which contains an input field, Guess and Reset buttons, and a text field that appears if the guess is higher or lower than the correct number. Every time the user clicks the Guess button, the `userNumberBean.check` method is called to reset the maximum or minimum value or, if the guess is correct, to generate a `FacesMessage` to that effect. The `userNumberBean.validateNumberRange` method determines whether each guess is valid.

## Building, Packaging, Deploying, and Running the guessnumber Example

You can build, package, deploy, and run the `guessnumber` application using either NetBeans IDE or the Ant tool.

### ▼ Building and Deploying the `guessnumber` Example Using NetBeans IDE

Follow these instructions to build, package, and deploy the `guessnumber` example to the Enterprise Server using NetBeans IDE.

This procedure builds the application into the `tut-install/examples/cdi/guessnumber/build/web` directory and deploys the contents of this directory to the Enterprise Server.

- 1 In NetBeans IDE, select **File**→**Open Project**.
- 2 In the **Open Project** dialog, navigate to `tut-install/examples/cdi/`.
- 3 Select the `guessnumber` folder.
- 4 Select the **Open as Main Project** checkbox.
- 5 Click **Open Project Folder**.
- 6 In the **Projects** tab, right-click the `guessnumber` project and select **Deploy**.

### ▼ Building, Packaging, and Deploying the `guessnumber` Example Using Ant

Follow these instructions to build, package, and deploy the `guessnumber` example to the Enterprise Server using Ant.

This procedure builds the application into the `tut-install/examples/cdi/guessnumber/build/web` directory and deploys the contents of this directory to the Enterprise Server.

This procedure builds and packages the application into `guessnumber.war`, located in `tut-install/examples/cdi/guessnumber/dist/`, and deploys this WAR file to the Enterprise Server.

- 1 In a terminal window, go to this directory:  
`tut-install/examples/cdi/guessnumber/`

**2 Type the following command:**

`ant`

This command calls the default target, which builds and packages the application into a WAR file, `guessnumber.war`, located in the `dist` directory.

**3 Type the following command:**

`ant deploy`

The `guessnumber.war` file will be deployed to the Enterprise Server.

▼ **Running the guessnumber Example**

**1 In a web browser, type the following URL:**

`http://localhost:8080/guessnumber`

The Guess My Number page opens, as shown in [Figure 21-2](#).

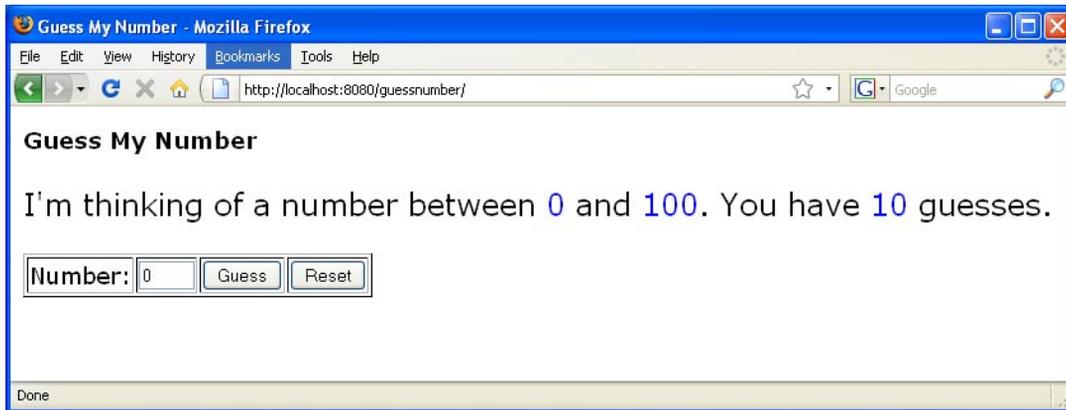


FIGURE 21-2 Guess My Number Example

**2 Type a number in the Number text field, then click Guess.**

The minimum and maximum values are modified, along with the remaining number of guesses.

**3 Keep guessing numbers until you get the right answer or run out of guesses.**

If you get the right answer, the input field and Guess button are grayed out, as shown in [Figure 21-3](#).

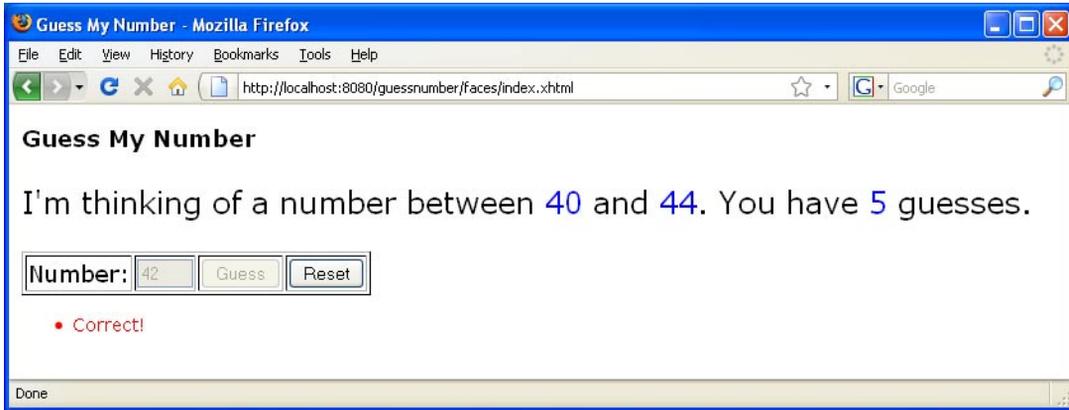


FIGURE 21-3 Guess My Number at End of Game

- 4 Click the Reset button to play the game again with a new random number.





PART VI

# Persistence

Part Six explores the Java Persistence API.



# Introduction to the Java Persistence API

---

The Java Persistence API provides an object/relational mapping facility to Java developers for managing relational data in Java applications. Java Persistence consists of four areas:

- The Java Persistence API
- The query language
- The Java Persistence Criteria API
- Object/relational mapping metadata

## Entities

An entity is a lightweight persistence domain object. Typically an entity represents a table in a relational database, and each entity instance corresponds to a row in that table. The primary programming artifact of an entity is the entity class, although entities can use helper classes.

The persistent state of an entity is represented either through persistent fields or persistent properties. These fields or properties use object/relational mapping annotations to map the entities and entity relationships to the relational data in the underlying data store.

## Requirements for Entity Classes

An entity class must follow these requirements:

- The class must be annotated with the `javax.persistence.Entity` annotation.
- The class must have a public or protected, no-argument constructor. The class may have other constructors.
- The class must not be declared `final`. No methods or persistent instance variables must be declared `final`.
- If an entity instance is passed by value as a detached object, such as through a session bean's remote business interface, the class must implement the `Serializable` interface.

- Entities may extend both entity and non-entity classes, and non-entity classes may extend entity classes.
- Persistent instance variables must be declared private, protected, or package-private, and can only be accessed directly by the entity class's methods. Clients must access the entity's state through accessor or business methods.

## Persistent Fields and Properties in Entity Classes

The persistent state of an entity can be accessed either through the entity's instance variables or through JavaBeans-style properties. The fields or properties must be of the following Java language types:

- Java primitive types
- `java.lang.String`
- Other serializable types including:
  - Wrappers of Java primitive types
  - `java.math.BigInteger`
  - `java.math.BigDecimal`
  - `java.util.Date`
  - `java.util.Calendar`
  - `java.sql.Date`
  - `java.sql.Time`
  - `java.sql.Timestamp`
  - User-defined serializable types
  - `byte[]`
  - `Byte[]`
  - `char[]`
  - `Character[]`
- Enumerated types
- Other entities and/or collections of entities
- Embeddable classes

Entities may use persistent fields, persistent properties, or a combination of both. If the mapping annotations are applied to the entity's instance variables, the entity uses persistent fields. If the mapping annotations are applied to the entity's getter methods for JavaBeans-style properties, the entity uses persistent properties.

### Persistent Fields

If the entity class uses persistent fields, the Persistence runtime accesses entity class instance variables directly. All fields not annotated `javax.persistence.Transient` or not marked as `JavaTransient` will be persisted to the data store. The object/relational mapping annotations must be applied to the instance variables.

## Persistent Properties

If the entity uses persistent properties, the entity must follow the method conventions of JavaBeans components. JavaBeans-style properties use getter and setter methods that are typically named after the entity class's instance variable names. For every persistent property *property* of type *Type* of the entity, there is a getter method `getProperty` and setter method `setProperty`. If the property is a boolean, you may use `isProperty` instead of `getProperty`. For example, if a `Customer` entity uses persistent properties, and has a private instance variable called `firstName`, the class defines a `getFirstName` and `setFirstName` method for retrieving and setting the state of the `firstName` instance variable.

The method signature for single-valued persistent properties are as follows:

```
Type getProperty()
void setProperty(Type type)
```

The object/relational mapping annotations for persistent properties must be applied to the getter methods. Mapping annotations cannot be applied to fields or properties annotated `@Transient` or marked `transient`.

## Using Collections in Entity Fields and Properties

Collection-valued persistent fields and properties must use the supported Java collection interfaces regardless of whether the entity uses persistent fields or properties. The following collection interfaces may be used:

- `java.util.Collection`
- `java.util.Set`
- `java.util.List`
- `java.util.Map`

If the entity class uses persistent fields, the type in the above method signatures must be one of these collection types. Generic variants of these collection types may also be used. For example, if the `Customer` entity has a persistent property that contains a set of phone numbers, it would have the following methods:

```
Set<PhoneNumber> getPhoneNumbers() { ... }
void setPhoneNumbers(Set<PhoneNumber>) { ... }
```

If a field or property of an entity consists of a collection of basic types or embeddable classes, use the `javax.persistence.ElementCollection` annotation on the field or property.

`@ElementCollection` has two attributes: `targetClass` and `fetch`. The `targetClass` attribute specifies the class name of the basic or embeddable class, and is optional if the field or property is defined using Java programming language generics. The optional `fetch` attribute is used to specify whether the collection should be retrieved lazily or eagerly, using the `javax.persistence.FetchType` constants of either `LAZY` or `EAGER`, respectively. By default, the collection will be fetched lazily.

**EXAMPLE 22-1** Specifying a Collection of Basic Types Using `@ElementCollection`

The following entity, `Person`, has a persistent field `nicknames` that is a collection of `String` classes that will be fetched eagerly. The `targetClass` element is not required because it uses generics to define the field.

```
@Entity
public class Person {
    ...
    @ElementCollection(fetch=EAGER)
    protected Set<String> nickname = new HashSet();
    ...
}
```

## Using Map Collections in Entities

Collections of entity elements and relationships may be represented by `java.util.Map` collections. A `Map` consists of a key and value.

When using `Map` elements or relationships, the following rules apply:

- The `Map` key or value may be a basic Java programming language type, an embeddable class, or an entity.
- When the `Map` value is an embeddable class or basic type, use the `@ElementCollection` annotation.
- When the `Map` value is an entity use the `@OneToMany` or `@ManyToMany` annotation.
- Only use the `Map` type on one side of a bidirectional relationship.

If the key type of a `Map` is a Java programming language basic type, use the `javax.persistence.MapKeyColumn` annotation to set the column mapping for the key. By default, the name attribute of `@MapKeyColumn` is of the form *RELATIONSHIP FIELD/PROPERTY NAME\_KEY*. For example, if the referencing relationship field name is `image`, the default name attribute is `IMAGE_KEY`.

If the key type of a `Map` is an entity, use the `javax.persistence.MapKeyJoinColumn` annotation. If the multiple columns are needed to set the mapping, use the `javax.persistence.MapKeyJoinColumns` annotation to include multiple `@MapKeyJoinColumn` annotations. If no `@MapKeyJoinColumn` is present, the mapping column name is by default set to *RELATIONSHIP FIELD/PROPERTY NAME\_KEY*. For example, if the relationship field name is `employee`, the default name attribute is `EMPLOYEE_KEY`.

If Java programming language generic types are not used in the relationship field or property, the key class must be explicitly set using the `javax.persistence.MapKeyClass` annotation.

If the `Map` key is the primary key, or a persistent field or property of the entity that is the `Map` value, use the `javax.persistence.MapKey` annotation. The `@MapKeyClass` and `@MapKey` annotations cannot be used on the same field or property.

If `Map` value is a Java programming language basic type or an embeddable class, it will be mapped as a collection table in the underlying database. If generic types are not used, the `@ElementCollection` annotation's `targetClass` attribute must be set to the type of the `Map` value.

If the `Map` value is an entity, and part of a many-to-many or one-to-many unidirectional relationship, it will be mapped as a join table in the underlying database. A unidirectional one-to-many relationship that uses a `Map` may also be mapped using the `@JoinColumn` annotation.

If the entity is part of a one-to-many/many-to-one bidirectional relationship, it will be mapped in the table of the entity that represents the value of the `Map`. If generic types are not used, the `targetEntity` attribute of the `@OneToMany` and `@ManyToMany` annotations must be set to the type of the `Map` value.

## Primary Keys in Entities

Each entity has a unique object identifier. A customer entity, for example, might be identified by a customer number. The unique identifier, or *primary key*, enables clients to locate a particular entity instance. Every entity must have a primary key. An entity may have either a simple or a composite primary key.

Simple primary keys use the `javax.persistence.Id` annotation to denote the primary key property or field.

Composite primary keys are used when a primary key consists of more than one attribute, which corresponds to a set of single persistent properties or fields. Composite primary keys must be defined in a primary key class. Composite primary keys are denoted using the `javax.persistence.EmbeddedId` and `javax.persistence.IdClass` annotations.

The primary key, or the property or field of a composite primary key, must be one of the following Java language types:

- Java primitive types
- Java primitive wrapper types
- `java.lang.String`
- `java.util.Date` (the temporal type should be `DATE`)
- `java.sql.Date`
- `java.math.BigDecimal`
- `java.math.BigInteger`

Floating point types should never be used in primary keys. If you use a generated primary key, only integral types will be portable.

## Primary Key Classes

A primary key class must meet these requirements:

- The access control modifier of the class must be `public`.
- The properties of the primary key class must be `public` or `protected` if property-based access is used.
- The class must have a public default constructor.
- The class must implement the `hashCode()` and `equals(Object other)` methods.
- The class must be serializable.
- A composite primary key must be represented and mapped to multiple fields or properties of the entity class, or must be represented and mapped as an embeddable class.
- If the class is mapped to multiple fields or properties of the entity class, the names and types of the primary key fields or properties in the primary key class must match those of the entity class.

The following primary key class is a composite key, the `orderId` and `itemId` fields together uniquely identify an entity.

```
public final class LineItemKey implements Serializable {
    public Integer orderId;
    public int itemId;

    public LineItemKey() {}

    public LineItemKey(Integer orderId, int itemId) {
        this.orderId = orderId;
        this.itemId = itemId;
    }

    public boolean equals(Object otherObj) {
        if (this == otherObj) {
            return true;
        }
        if (!(otherObj instanceof LineItemKey)) {
            return false;
        }
        LineItemKey other = (LineItemKey) otherObj;
        return (
            (orderId==null?other.orderId==null:orderId.equals
            (other.orderId)
            )
            &&
            (itemId == other.itemId)
        );
    }
}
```

```
public int hashCode() {
    return (
        (orderId==null?0:orderId.hashCode())
        ^
        ((int) itemId)
    );
}

public String toString() {
    return "" + orderId + "-" + itemId;
}
}
```

## Multiplicity in Entity Relationships

There are four types of multiplicities: one-to-one, one-to-many, many-to-one, and many-to-many.

*One-to-one:* Each entity instance is related to a single instance of another entity. For example, to model a physical warehouse in which each storage bin contains a single widget, `StorageBin` and `Widget` would have a one-to-one relationship. One-to-one relationships use the `javax.persistence.OneToOne` annotation on the corresponding persistent property or field.

*One-to-many:* An entity instance can be related to multiple instances of the other entities. A sales order, for example, can have multiple line items. In the order application, `Order` would have a one-to-many relationship with `LineItem`. One-to-many relationships use the `javax.persistence.OneToMany` annotation on the corresponding persistent property or field.

*Many-to-one:* Multiple instances of an entity can be related to a single instance of the other entity. This multiplicity is the opposite of a one-to-many relationship. In the example just mentioned, from the perspective of `LineItem` the relationship to `Order` is many-to-one. Many-to-one relationships use the `javax.persistence.ManyToOne` annotation on the corresponding persistent property or field.

*Many-to-many:* The entity instances can be related to multiple instances of each other. For example, in college each course has many students, and every student may take several courses. Therefore, in an enrollment application, `Course` and `Student` would have a many-to-many relationship. Many-to-many relationships use the `javax.persistence.ManyToMany` annotation on the corresponding persistent property or field.

## Direction in Entity Relationships

The direction of a relationship can be either bidirectional or unidirectional. A bidirectional relationship has both an owning side and an inverse side. A unidirectional relationship has only an owning side. The owning side of a relationship determines how the Persistence runtime makes updates to the relationship in the database.

### Bidirectional Relationships

In a *bidirectional* relationship, each entity has a relationship field or property that refers to the other entity. Through the relationship field or property, an entity class's code can access its related object. If an entity has a related field, then the entity is said to “know” about its related object. For example, if `Order` knows what `LineItem` instances it has and if `LineItem` knows what `Order` it belongs to, then they have a bidirectional relationship.

Bidirectional relationships must follow these rules:

- The inverse side of a bidirectional relationship must refer to its owning side by using the `mappedBy` element of the `@OneToOne`, `@OneToMany`, or `@ManyToMany` annotation. The `mappedBy` element designates the property or field in the entity that is the owner of the relationship.
- The many side of many-to-one bidirectional relationships must not define the `mappedBy` element. The many side is always the owning side of the relationship.
- For one-to-one bidirectional relationships, the owning side corresponds to the side that contains the corresponding foreign key.
- For many-to-many bidirectional relationships either side may be the owning side.

### Unidirectional Relationships

In a *unidirectional* relationship, only one entity has a relationship field or property that refers to the other. For example, `LineItem` would have a relationship field that identifies `Product`, but `Product` would not have a relationship field or property for `LineItem`. In other words, `LineItem` knows about `Product`, but `Product` doesn't know which `LineItem` instances refer to it.

### Queries and Relationship Direction

Java Persistence query language and Criteria API queries often navigate across relationships. The direction of a relationship determines whether a query can navigate from one entity to another. For example, a query can navigate from `LineItem` to `Product` but cannot navigate in the opposite direction. For `Order` and `LineItem`, a query could navigate in both directions, because these two entities have a bidirectional relationship.

## Cascade Operations and Relationships

Entities that use relationships often have dependencies on the existence of the other entity in the relationship. For example, a line item is part of an order, and if the order is deleted, then the line item should also be deleted. This is called a cascade delete relationship.

The `javax.persistence.CascadeType` enumerated type defines the cascade operations that are applied in the cascade element of the relationship annotations.

TABLE 22-1 Cascade Operations For Entities

Cascade Operation	Description
ALL	All cascade operations will be applied to the parent entity's related entity. All is equivalent to specifying <code>cascade={DETACH, MERGE, PERSIST, REFRESH, REMOVE}</code>
DETACH	If the parent entity is detached from the persistence context, the related entity will also be detached.
MERGE	If the parent entity is merged into the persistence context, the related entity will also be merged.
PERSIST	If the parent entity is persisted into the persistence context, the related entity will also be persisted.
REFRESH	If the parent entity is refreshed in the current persistence context, the related entity will also be refreshed.
REMOVE	If the parent entity is removed from the current persistence context, the related entity will also be removed.

Cascade delete relationships are specified using the `cascade=REMOVE` element specification for `@OneToOne` and `@OneToMany` relationships. For example:

```
@OneToMany(cascade=REMOVE, mappedBy="customer")
public Set<Order> getOrders() { return orders; }
```

## Orphan Removal in Relationships

When a target entity in one-to-one or one-to-many relationship is removed from the relationship, it is often desirable to cascade the remove operation to the target entity. Such target entities are considered “orphans,” and the `orphanRemoval` attribute can be used to specify that orphaned entities should be removed. For example, if an order has many line items, and one of the line items is removed from the order, the removed line item is considered an orphan. If `orphanRemoval` is set to `true`, the line item entity will be deleted when the line item is removed from the order.

The `orphanRemoval` attribute in `@OneToMany` and `@OneToOne` takes a boolean value, and is by default `false`.

**EXAMPLE 22-2** Enabling Orphan Removal in `@OneToMany` Relationship

The following example will cascade the remove operation to the orphaned customer entity when it is removed from the relationship.

```
@OneToMany(mappedBy="customer", orphanRemoval=true")
public List<Order> getOrders() { ... }
```

## Embeddable Classes in Entities

Embeddable classes are used to represent the state of an entity, but don't have a persistent identity of their own, unlike entity classes. Instances of an embeddable class share the identity of the entity that owns it. Embeddable classes only exist as the state of another entity. An entity may have single-valued or collection-valued embeddable class attributes.

Embeddable classes have the same rules as entity classes, except that they are annotated with the `javax.persistence.Embeddable` annotation instead of `@Entity`.

**EXAMPLE 22-3** Embeddable Class Example

The following embeddable class, `ZipCode` has two fields, `zip` and `plusFour`.

```
@Embeddable
public class ZipCode {
    String zip;
    String plusFour;
    ...
}
```

This embeddable class is used by the `Address` entity.

```
@Entity
public class Address {
    @Id
    protected long id
    String street1;
    String street2;
    String city;
    String province;
    @Embedded
    ZipCode zipCode;
    String country;
    ...
}
```

**EXAMPLE 22-3** Embeddable Class Example (Continued)

```
}

```

Entities that own embeddable classes as part of their persistent state may annotate the field or property with the `javax.persistence.Embedded` annotation, but are not required to do so.

Embeddable classes may themselves use other embeddable classes to represent their state. They may also contain collections of basic Java programming language types, or other embeddable classes. Embeddable classes may also contain relationships to other entities or collections of entities. If the embeddable class has such a relationship, the relationship is from the target entity or collection of entities to the entity that owns the embeddable class.

## Entity Inheritance

Entities support class inheritance, polymorphic associations, and polymorphic queries. They can extend non-entity classes, and non-entity classes can extend entity classes. Entity classes can be both abstract and concrete.

The `roster` example application demonstrates entity inheritance, and is described in “[Entity Inheritance in the roster Application](#)” on page 416.

### Abstract Entities

An abstract class may be declared an entity by decorating the class with `@Entity`. Abstract entities differ from concrete entities only in that they cannot be instantiated.

Abstract entities can be queried just like concrete entities. If an abstract entity is the target of a query, the query operates on all the concrete subclasses of the abstract entity.

```
@Entity
public abstract class Employee {
    @Id
    protected Integer employeeId;
    ...
}
@Entity
public class FullTimeEmployee extends Employee {
    protected Integer salary;
    ...
}
@Entity
public class PartTimeEmployee extends Employee {
    protected Float hourlyWage;
}

```

## Mapped Superclasses

Entities may inherit from superclasses that contain persistent state and mapping information, but are not entities. That is, the superclass is not decorated with the `@Entity` annotation, and is not mapped as an entity by the Java Persistence provider. These superclasses are most often used when you have state and mapping information common to multiple entity classes.

Mapped superclasses are specified by decorating the class with the `javax.persistence.MappedSuperclass` annotation.

```
@MappedSuperclass
public class Employee {
    @Id
    protected Integer employeeId;
    ...
}
@Entity
public class FullTimeEmployee extends Employee {
    protected Integer salary;
    ...
}
@Entity
public class PartTimeEmployee extends Employee {
    protected Float hourlyWage;
    ...
}
```

Mapped superclasses cannot be queried, and can't be used in `EntityManager` or `Query` operations. You must use entity subclasses of the mapped superclass in `EntityManager` or `Query` operations. Mapped superclasses can't be targets of entity relationships. Mapped superclasses can be abstract or concrete.

Mapped superclasses do not have any corresponding tables in the underlying datastore. Entities that inherit from the mapped superclass define the table mappings. For instance, in the code sample above the underlying tables would be `FULLTIMEEMPLOYEE` and `PARTTIMEEMPLOYEE`, but there is no `EMPLOYEE` table.

## Non-Entity Superclasses

Entities may have non-entity superclasses, and these superclasses can be either abstract or concrete. The state of non-entity superclasses is non-persistent, and any state inherited from the non-entity superclass by an entity class is non-persistent. Non-entity superclasses may not be used in `EntityManager` or `Query` operations. Any mapping or relationship annotations in non-entity superclasses are ignored.

## Entity Inheritance Mapping Strategies

You can configure how the Java Persistence provider maps inherited entities to the underlying datastore by decorating the root class of the hierarchy with the `javax.persistence.Inheritance` annotation. There are three mapping strategies that are used to map the entity data to the underlying database:

- A single table per class hierarchy
- A table per concrete entity class
- A “join” strategy, where fields or properties that are specific to a subclass are mapped to a different table than the fields or properties that are common to the parent class

The strategy is configured by setting the `strategy` element of `@Inheritance` to one of the options defined in the `javax.persistence.InheritanceType` enumerated type:

```
public enum InheritanceType {
    SINGLE_TABLE,
    JOINED,
    TABLE_PER_CLASS
};
```

The default strategy is `InheritanceType.SINGLE_TABLE`, and is used if the `@Inheritance` annotation is not specified on the root class of the entity hierarchy.

### The Single Table per Class Hierarchy Strategy

With this strategy, which corresponds to the default `InheritanceType.SINGLE_TABLE`, all classes in the hierarchy are mapped to a single table in the database. This table has a *discriminator column*, a column that contains a value that identifies the subclass to which the instance represented by the row belongs.

The discriminator column can be specified by using the `javax.persistence.DiscriminatorColumn` annotation on the root of the entity class hierarchy.

TABLE 22-2 @DiscriminatorColumn Elements

Type	Name	Description
String	name	The name of the column in the table to be used as the discriminator column. The default is <code>DTYPE</code> . This element is optional.

TABLE 22-2 @DiscriminatorColumn Elements (Continued)

Type	Name	Description
DiscriminatorType	discriminatorType	The type of the column to be used as a discriminator column. The default is <code>DiscriminatorType.STRING</code> . This element is optional.
String	columnDefinition	The SQL fragment to use when creating the discriminator column. The default is generated by the Persistence provider, and is implementation-specific. This element is optional.
String	length	The column length for <code>String</code> -based discriminator types. This element is ignored for non- <code>String</code> discriminator types. The default is 31. This element is optional.

The `javax.persistence.DiscriminatorType` enumerated type is used to set the type of the discriminator column in the database by setting the `discriminatorType` element of `@DiscriminatorColumn` to one of the defined types. `DiscriminatorType` is defined as:

```
public enum DiscriminatorType {
    STRING,
    CHAR,
    INTEGER
};
```

If `@DiscriminatorColumn` is not specified on the root of the entity hierarchy and a discriminator column is required, the Persistence provider assumes a default column name of `DTYPE`, and column type of `DiscriminatorType.STRING`.

The `javax.persistence.DiscriminatorValue` annotation may be used to set the value entered into the discriminator column for each entity in a class hierarchy. You may only decorate concrete entity classes with `@DiscriminatorValue`.

If `@DiscriminatorValue` is not specified on an entity in a class hierarchy that uses a discriminator column, the Persistence provider will provide a default, implementation-specific value. If the `discriminatorType` element of `@DiscriminatorColumn` is `DiscriminatorType.STRING`, the default value is the name of the entity.

This strategy provides good support for polymorphic relationships between entities and queries that cover the entire entity class hierarchy. However, it requires the columns that contain the state of subclasses to be nullable.

## The Table per Concrete Class Strategy

In this strategy, which corresponds to `InheritanceType.TABLE_PER_CLASS`, each concrete class is mapped to a separate table in the database. All fields or properties in the class, including inherited fields or properties, are mapped to columns in the class's table in the database.

This strategy provides poor support for polymorphic relationships, and usually requires either SQL `UNION` queries or separate SQL queries for each subclass for queries that cover the entire entity class hierarchy.

Support for this strategy is optional, and may not be supported by all Java Persistence API providers. The default Java Persistence API provider in the Enterprise Server does not support this strategy.

## The Joined Subclass Strategy

In this strategy, which corresponds to `InheritanceType.JOINED`, the root of the class hierarchy is represented by a single table, and each subclass has a separate table that only contains those fields specific to that subclass. That is, the subclass table does not contain columns for inherited fields or properties. The subclass table also has a column or columns that represent its primary key, which is a foreign key to the primary key of the superclass table.

This strategy provides good support for polymorphic relationships, but requires one or more join operations to be performed when instantiating entity subclasses. This may result in poor performance for extensive class hierarchies. Similarly, queries that cover the entire class hierarchy require join operations between the subclass tables, resulting in decreased performance.

Some Java Persistence API providers, including the default provider in the Enterprise Server, require a discriminator column in the table that corresponds to the root entity when using the joined subclass strategy. If you are not using automatic table creation in your application, make sure the database table is set up correctly for the discriminator column defaults, or use the `@DiscriminatorColumn` annotation to match your database schema. For information on discriminator columns, see [“The Single Table per Class Hierarchy Strategy” on page 393](#).

# Managing Entities

Entities are managed by the entity manager. The entity manager is represented by `javax.persistence.EntityManager` instances. Each `EntityManager` instance is associated with a persistence context. A persistence context defines the scope under which particular entity instances are created, persisted, and removed.

## The Persistence Context

A persistence context is a set of managed entity instances that exist in a particular data store. The `EntityManager` interface defines the methods that are used to interact with the persistence context.

## The EntityManager Interface

The `EntityManager` API creates and removes persistent entity instances, finds entities by the entity's primary key, and allows queries to be run on entities.

### Container-Managed Entity Managers

With a *container-managed entity manager*, an `EntityManager` instance's persistence context is automatically propagated by the container to all application components that use the `EntityManager` instance within a single Java Transaction Architecture (JTA) transaction.

JTA transactions usually involve calls across application components. To complete a JTA transaction, these components usually need access to a single persistence context. This occurs when an `EntityManager` is injected into the application components by means of the `javax.persistence.PersistenceContext` annotation. The persistence context is automatically propagated with the current JTA transaction, and `EntityManager` references that are mapped to the same persistence unit provide access to the persistence context within that transaction. By automatically propagating the persistence context, application components don't need to pass references to `EntityManager` instances to each other in order to make changes within a single transaction. The Java EE container manages the life cycle of container-managed entity managers.

To obtain an `EntityManager` instance, inject the entity manager into the application component:

```
@PersistenceContext
EntityManager em;
```

### Application-Managed Entity Managers

With *application-managed entity managers*, on the other hand, the persistence context is not propagated to application components, and the life cycle of `EntityManager` instances is managed by the application.

Application-managed entity managers are used when applications need to access a persistence context that is not propagated with the JTA transaction across `EntityManager` instances in a particular persistence unit. In this case, each `EntityManager` creates a new, isolated persistence context. The `EntityManager`, and its associated persistence context, is created and destroyed explicitly by the application. They are also used when directly injecting `EntityManager`

instances can't be done because `EntityManager` instances are not thread-safe. `EntityManagerFactory` instances are thread-safe.

Applications create `EntityManager` instances in this case by using the `createEntityManager` method of `javax.persistence.EntityManagerFactory`.

To obtain an `EntityManager` instance, you first must obtain an `EntityManagerFactory` instance by injecting it into the application component by means of the `javax.persistence.PersistenceUnit` annotation:

```
@PersistenceUnit
EntityManagerFactory emf;
```

Then, obtain an `EntityManager` from the `EntityManagerFactory` instance:

```
EntityManager em = emf.createEntityManager();
```

Application-managed entity managers don't automatically propagate the JTA transaction context. Such applications need to manually gain access to the JTA transaction manager and add transaction demarcation information when performing entity operations. The `javax.transaction.UserTransaction` interface defines methods to begin, commit, and rollback transactions. Inject an instance of `UserTransaction` by creating an instance variable annotated with `@Resource`.

```
@Resource
UserTransaction utx;
```

To begin a transaction, call the `UserTransaction.begin` method. When all the entity operations are complete, call the `UserTransaction.commit` method to commit the transaction. The `UserTransaction.rollback` method is used to roll back the current transaction.

#### EXAMPLE 22-4 User-Managed Transactions

The following example shows how to manage transactions in an application that uses an application-managed entity manager.

```
@PersistenceContext
EntityManagerFactory emf;
EntityManager em;
@Resource
UserTransaction utx;
...
em = emf.createEntityManager();
try {
    utx.begin();
    em.persist(SomeEntity);
    em.merge(AnotherEntity);
}
```

**EXAMPLE 22-4** User-Managed Transactions *(Continued)*

```
    em.remove(ThirdEntity);
    utx.commit();
} catch (Exception e) {
    utx.rollback();
}
```

## Finding Entities Using the EntityManager

The `EntityManager.find` method is used to look up entities in the data store by the entity's primary key.

```
@PersistenceContext
EntityManager em;
public void enterOrder(int custID, Order newOrder) {
    Customer cust = em.find(Customer.class, custID);
    cust.getOrders().add(newOrder);
    newOrder.setCustomer(cust);
}
```

## Managing an Entity Instance's Life Cycle

You manage entity instances by invoking operations on the entity by means of an `EntityManager` instance. Entity instances are in one of four states: new, managed, detached, or removed.

New entity instances have no persistent identity and are not yet associated with a persistence context.

Managed entity instances have a persistent identity and are associated with a persistence context.

Detached entity instances have a persistent identity and are not currently associated with a persistence context.

Removed entity instances have a persistent identity, are associated with a persistent context, and are scheduled for removal from the data store.

## Persisting Entity Instances

New entity instances become managed and persistent either by invoking the `persist` method, or by a cascading `persist` operation invoked from related entities that have the `cascade=PERSIST` or `cascade=ALL` elements set in the relationship annotation. This means the entity's data is stored to the database when the transaction associated with the `persist` operation is completed. If the entity is already managed, the `persist` operation is ignored, although the `persist` operation will cascade to related entities that have the `cascade` element

set to `PERSIST` or `ALL` in the relationship annotation. If `persist` is called on a removed entity instance, it becomes managed. If the entity is detached, `persist` will throw an `IllegalArgumentException`, or the transaction commit will fail.

```
@PersistenceContext
EntityManager em;
...
public LineItem createLineItem(Order order, Product product,
    int quantity) {
    LineItem li = new LineItem(order, product, quantity);
    order.getLineItems().add(li);
    em.persist(li);
    return li;
}
```

The `persist` operation is propagated to all entities related to the calling entity that have the cascade element set to `ALL` or `PERSIST` in the relationship annotation.

```
@OneToMany(cascade=ALL, mappedBy="order")
public Collection<LineItem> getLineItems() {
    return lineItems;
}
```

## Removing Entity Instances

Managed entity instances are removed by invoking the `remove` method, or by a cascading remove operation invoked from related entities that have the `cascade=REMOVE` or `cascade=ALL` elements set in the relationship annotation. If the `remove` method is invoked on a new entity, the `remove` operation is ignored, although `remove` will cascade to related entities that have the cascade element set to `REMOVE` or `ALL` in the relationship annotation. If `remove` is invoked on a detached entity it will throw an `IllegalArgumentException`, or the transaction commit will fail. If `remove` is invoked on an already removed entity, it will be ignored. The entity's data will be removed from the data store when the transaction is completed, or as a result of the `flush` operation.

```
public void removeOrder(Integer orderId) {
    try {
        Order order = em.find(Order.class, orderId);
        em.remove(order);
    }...
}
```

In this example, all `LineItem` entities associated with the order are also removed, as `Order.getLineItems` has `cascade=ALL` set in the relationship annotation.

## Synchronizing Entity Data to the Database

The state of persistent entities is synchronized to the database when the transaction with which the entity is associated commits. If a managed entity is in a bidirectional relationship with another managed entity, the data will be persisted based on the owning side of the relationship.

To force synchronization of the managed entity to the data store, invoke the `flush` method of the `EntityManager` instance. If the entity is related to another entity, and the relationship annotation has the `cascade` element set to `PERSIST` or `ALL`, the related entity's data will be synchronized with the data store when `flush` is called.

If the entity is removed, calling `flush` will remove the entity data from the data store.

## Persistence Units

A persistence unit defines a set of all entity classes that are managed by `EntityManager` instances in an application. This set of entity classes represents the data contained within a single data store.

Persistence units are defined by the `persistence.xml` configuration file. The JAR file or directory whose `META-INF` directory contains `persistence.xml` is called the root of the persistence unit. The scope of the persistence unit is determined by the persistence unit's root.

Each persistence unit must be identified with a name that is unique to the persistence unit's scope.

Persistent units can be packaged as part of a WAR or EJB JAR file, or can be packaged as a JAR file that can then be included in an WAR or EAR file.

If you package the persistent unit as a set of classes in an EJB JAR file, `persistence.xml` should be put in the EJB JAR's `META-INF` directory.

If you package the persistence unit as a set of classes in a WAR file, `persistence.xml` should be located in the WAR file's `WEB-INF/classes/META-INF` directory.

If you package the persistence unit in a JAR file that will be included in a WAR or EAR file, the JAR file should be located:

- In the `WEB-INF/lib` directory of a WAR.
- In the EAR file's library directory.

---

**Note** – In the Java Persistence API 1.0, JAR files could be located at the root of an EAR file as the root of the persistence unit. This is no longer supported. Portable applications should use the EAR file's library directory as the root of the persistence unit.

---

## The persistence.xml File

persistence.xml defines one or more persistence units. The following is an example persistence.xml file.

```
<persistence>
  <persistence-unit name="OrderManagement">
    <description>This unit manages orders and customers.
      It does not rely on any vendor-specific features and can
      therefore be deployed to any persistence provider.
    </description>
    <jta-data-source>jdbc/MyOrderDB</jta-data-source>
    <jar-file>MyOrderApp.jar</jar-file>
    <class>com.widgets.Order</class>
    <class>com.widgets.Customer</class>
  </persistence-unit>
</persistence>
```

This file defines a persistence unit named `OrderManagement`, which uses a JTA-aware data source `jdbc/MyOrderDB`. The `jar-file` and `class` elements specify managed persistence classes: entity classes, embeddable classes, and mapped superclasses. The `jar-file` element specifies JAR files that are visible to the packaged persistence unit that contain managed persistence classes, while the `class` element explicitly names managed persistence classes.

The `jta-data-source` (for JTA-aware data sources) and `non-jta-data-source` (non-JTA-aware data sources) elements specify the global JNDI name of the data source to be used by the container.

## Querying Entities

There are two methods of querying entities using the Java Persistence API. The Java Persistence query language (JPQL) is a simple, string-based language similar to SQL used to query entities and their relationships. See [Chapter 24, “The Java Persistence Query Language,”](#) for more information on the Java Persistence query language. The Criteria API is used to create type-safe queries using Java programming language APIs to query for entities and their relationships. See [Chapter 25, “Creating Queries Using the Criteria API,”](#) for more information on the Criteria API.

Each approach, JPQL and the Criteria API, has advantages and disadvantages.

JPQL queries are typically more concise compared to Criteria queries, just a few lines long. They also tend to be more readable compared to Criteria queries, and developers familiar with SQL will find it easy to learn the syntax of JPQL. JPQL named queries can be defined in the entity class using a Java programming language annotation, or in the application's deployment descriptor. JPQL queries are not type-safe, however, and require a cast when retrieving the query result from the entity manager. This means that type casting errors may not be caught at compile-time. JPQL queries don't support open-ended parameters.

Criteria queries allow you to define the query in the business tier of the application. While this is also possible using JPQL dynamic queries, Criteria queries provide better performance because JPQL dynamic queries must be parsed each time they are called. Criteria queries are type-safe, and therefore don't require casting like JPQL queries. The Criteria API is just another Java programming language API, and doesn't require developers to learn the syntax of another query language. Criteria queries are typically more verbose compared to JPQL queries, and require the developer to create several objects and perform operations on those objects before submitting the query to the entity manager.

## Running the Persistence Examples

---

This chapter describes how to use the Java Persistence API in different example applications. The material here focuses on the source code and settings of three examples. The first example called `order` is an application that uses a stateful session bean to manage entities related to an ordering system. The second example is `roster`, an application that manages a community sports system. This chapter assumes that you are familiar with the concepts detailed in [Chapter 22, “Introduction to the Java Persistence API.”](#)

### The `order` Application

The `order` application is a simple inventory and ordering application for maintaining a catalog of parts and placing an itemized order of those parts. It has entities that represent parts, vendors, orders, and line items. These entities are accessed using a stateful session bean that holds the business logic of the application. A simple singleton session bean creates the initial entities on application deployment. A Facelets web application manipulates the data, and displays data from the catalog.

The information contained in an order can be divided into different elements. What is the order number? What parts are included in the order? What parts make up that part? Who makes the part? What are the specifications for the part? Are there any schematics for the part? `order` is a simplified version of an ordering system that has all these elements.

The `order` application consists of a single WAR module that includes the enterprise bean classes, the entities, the support classes, and the Facelets XHTML and class files.

### Entity Relationships in the `order` Application

The `order` application demonstrates several types of entity relationships: one-to-many, many-to-one, one-to-one, unidirectional, and self-referential relationships.

## Self-Referential Relationships

A *self-referential* relationship is a relationship between relationship fields in the same entity. Part has a field `bomPart` that has a one-to-many relationship with the field `parts`, which is also in Part. That is, a part can be made up of many parts, and each of those parts has exactly one bill-of-material part.

The primary key for Part is a compound primary key, a combination of the `partNumber` and `revision` fields. It is mapped to the `PARTNUMBER` and `REVISION` columns in the `EJB_ORDER_PART` table.

```
...
@ManyToOne
@JoinColumns({
    @JoinColumn(name="BOMPARTNUMBER",
        referencedColumnName="PARTNUMBER"),
    @JoinColumn(name="BOMREVISION",
        referencedColumnName="REVISION")
})
public Part getBomPart() {
    return bomPart;
}
...
@OneToMany(mappedBy="bomPart")
public Collection<Part> getParts() {
    return parts;
}
...
```

## One-to-One Relationships

Part has a field, `vendorPart`, that has a one-to-one relationship with `VendorPart`'s `part` field. That is, each part has exactly one vendor part, and vice versa.

Here is the relationship mapping in Part:

```
@OneToOne(mappedBy="part")
public VendorPart getVendorPart() {
    return vendorPart;
}
```

Here is the relationship mapping in VendorPart:

```
@OneToOne
@JoinColumns({
    @JoinColumn(name="PARTNUMBER",
        referencedColumnName="PARTNUMBER"),
    @JoinColumn(name="PARTREVISION",
```

```

        referencedColumnName="REVISION")
    })
    public Part getPart() {
        return part;
    }
}

```

Note that, because `Part` uses a compound primary key, the `@JoinColumn` annotation is used to map the columns in the `PERSISTENCE_ORDER_VENDOR_PART` table to the columns in `PERSISTENCE_ORDER_PART`. `PERSISTENCE_ORDER_VENDOR_PART`'s `PARTREVISION` column refers to `PERSISTENCE_ORDER_PART`'s `REVISION` column.

## One-to-Many Relationship Mapped to Overlapping Primary and Foreign Keys

`Order` has a field, `lineItems`, that has a one-to-many relationship with `LineItem`'s field `order`. That is, each order has one or more line item.

`LineItem` uses a compound primary key that is made up of the `orderId` and `itemId` fields. This compound primary key maps to the `ORDERID` and `ITEMID` columns in the `PERSISTENCE_ORDER_LINEITEM` database table. `ORDERID` is a foreign key to the `ORDERID` column in the `PERSISTENCE_ORDER_ORDER` table. This means that the `ORDERID` column is mapped twice: once as a primary key field, `orderId`; and again as a relationship field, `order`.

Here's the relationship mapping in `Order`:

```

@OneToMany(cascade=ALL, mappedBy="order")
    public Collection<LineItem> getLineItems() {
        return lineItems;
    }
}

```

Here is the relationship mapping in `LineItem`:

```

@ManyToOne
    public Order getOrder() {
        return order;
    }
}

```

## Unidirectional Relationships

`LineItem` has a field, `vendorPart`, that has a unidirectional many-to-one relationship with `VendorPart`. That is, there is no field in the target entity in this relationship.

```

@ManyToOne
    public VendorPart getVendorPart() {
        return vendorPart;
    }
}

```

## Primary Keys in the order Application

The order application uses several types of primary keys: single-valued primary keys, compound primary keys, and generated primary keys.

### Generated Primary Keys

`VendorPart` uses a generated primary key value. That is, the application does not assign primary key values for the entities, but instead relies on the persistence provider to generate the primary key values. The `@GeneratedValue` annotation is used to specify that an entity will use a generated primary key.

In `VendorPart`, the following code specifies the settings for generating primary key values:

```
@TableGenerator(  
    name="vendorPartGen",  
    table="PERSISTENCE_ORDER_SEQUENCE_GENERATOR",  
    pkColumnName="GEN_KEY",  
    valueColumnName="GEN_VALUE",  
    pkColumnValue="VENDOR_PART_ID",  
    allocationSize=10)  
@Id  
@GeneratedValue(strategy=GenerationType.TABLE,  
    generator="vendorPartGen")  
public Long getVendorPartNumber() {  
    return vendorPartNumber;  
}
```

The `@TableGenerator` annotation is used in conjunction with `@GeneratedValue`'s `strategy=TABLE` element. That is, the strategy used to generate the primary keys is use a table in the database. `@TableGenerator` is used to configure the settings for the generator table. The `name` element sets the name of the generator, which is `vendorPartGen` in `VendorPart`.

The `EJB_ORDER_SEQUENCE_GENERATOR` table, which has two columns `GEN_KEY` and `GEN_VALUE`, will store the generated primary key values. This table could be used to generate other entity's primary keys, so the `pkColumnValue` element is set to `VENDOR_PART_ID` to distinguish this entity's generated primary keys from other entity's generated primary keys. The `allocationSize` element specifies the amount to increment when allocating primary key values. In this case, each `VendorPart`'s primary key will increment by 10.

The primary key field `vendorPartNumber` is of type `Long`, as the generated primary key's field must be an integral type.

### Compound Primary Keys

A compound primary key is made up of multiple fields and follows the requirements described in [“Primary Key Classes” on page 386](#). To use a compound primary key, you must create a wrapper class.

In order, two entities use compound primary keys: Part and LineItem.

Part uses the PartKey wrapper class. Part's primary key is a combination of the part number and the revision number. PartKey encapsulates this primary key.

LineItem uses the LineItemKey class. LineItem's primary key is a combination of the order number and the item number. LineItemKey encapsulates this primary key. This is the LineItemKey compound primary key wrapper class:

```
package order.entity;

public final class LineItemKey implements
    java.io.Serializable {

    private Integer orderId;
    private int itemId;

    public int hashCode() {
        return ((this.getOrderId()==null
            ?0:this.getOrderId().hashCode())
            ^ ((int) this.getItemId()));
    }

    public boolean equals(Object otherObj) {
        if (this == otherObj) {
            return true;
        }
        if (!(otherObj instanceof LineItemKey)) {
            return false;
        }
        LineItemKey other = (LineItemKey) otherObj;
        return ((this.getOrderId()==null
            ?other.orderId==null:this.getOrderId().equals
            (other.orderId)) && (this.getItemId ==
            other.itemId));
    }

    public String toString() {
        return "" + orderId + "-" + itemId;
    }
}
```

The @IdClass annotation is used to specify the primary key class in the entity class. In LineItem, @IdClass is used as follows:

```
@IdClass(order.entity.LineItemKey.class)
@Entity
...
```

```
public class LineItem {  
    ...  
}
```

The two fields in `LineItem` are tagged with the `@Id` annotation to mark those fields as part of the compound primary key:

```
@Id  
public int getItemId() {  
    return itemId;  
}  
...  
@Id  
@Column(name="ORDERID", nullable=false,  
        insertable=false, updatable=false)  
public Integer getOrderId() {  
    return orderId;  
}
```

For `orderId`, you also use the `@Column` annotation to specify the column name in the table, and that this column should not be inserted or updated, as it is an overlapping foreign key pointing at the `PERSISTENCE_ORDER_ORDER` table's `ORDERID` column (see [“One-to-Many Relationship Mapped to Overlapping Primary and Foreign Keys”](#) on page 405). That is, `orderId` will be set by the `Order` entity.

In `LineItem`'s constructor, the line item number (`LineItem.itemId`) is set using the `Order.getNextId` method.

```
public LineItem(Order order, int quantity, VendorPart  
    vendorPart) {  
    this.order = order;  
    this.itemId = order.getNextId();  
    this.orderId = order.getOrderId();  
    this.quantity = quantity;  
    this.vendorPart = vendorPart;  
}
```

`Order.getNextId` counts the number of current line items, adds one, and returns that number.

```
public int getNextId() {  
    return this.lineItems.size() + 1;  
}
```

`Part` doesn't require the `@Column` annotation on the two fields that comprise `Part`'s compound primary key. This is because `Part`'s compound primary key is not an overlapping primary key/foreign key.

```

@IdClass(order.entity.PartKey.class)
@Entity
...
public class Part {
...
    @Id
    public String getPartNumber() {
        return partNumber;
    }
...
    @Id
    public int getRevision() {
        return revision;
    }
...
}

```

## Entity Mapped to More Than One Database Table

Part's fields map to more than one database table: PERSISTENCE\_ORDER\_PART and PERSISTENCE\_ORDER\_PART\_DETAIL. The PERSISTENCE\_ORDER\_PART\_DETAIL table holds the specification and schematics for the part. The @SecondaryTable annotation is used to specify the secondary table.

```

...
@Entity
@Table(name="PERSISTENCE_ORDER_PART")
@SecondaryTable(name="PERSISTENCE_ORDER_PART_DETAIL", pkJoinColumns={
    @PrimaryKeyJoinColumn(name="PARTNUMBER",
        referencedColumnName="PARTNUMBER"),
    @PrimaryKeyJoinColumn(name="REVISION",
        referencedColumnName="REVISION")
})
public class Part {
...
}

```

PERSISTENCE\_ORDER\_PART\_DETAIL shares the same primary key values as PERSISTENCE\_ORDER\_PART. The pkJoinColumns element of @SecondaryTable is used to specify that PERSISTENCE\_ORDER\_PART\_DETAIL's primary key columns are foreign keys to PERSISTENCE\_ORDER\_PART. The @PrimaryKeyJoinColumn annotation sets the primary key column names and specifies which column in the primary table the column refers to. In this case, the primary key column names for both PERSISTENCE\_ORDER\_PART\_DETAIL and PERSISTENCE\_ORDER\_PART are the same: PARTNUMBER and REVISION, respectively.

## Cascade Operations in the order Application

Entities that have relationships to other entities often have dependencies on the existence of the other entity in the relationship. For example, a line item is part of an order, and if the order is deleted, then the line item should also be deleted. This is called a cascade delete relationship.

In order, there are two cascade delete dependencies in the entity relationships. If the `Order` to which a `LineItem` is related is deleted, then the `LineItem` should also be deleted. If the `Vendor` to which a `VendorPart` is related is deleted, then the `VendorPart` should also be deleted.

You specify the cascade operations for entity relationships by setting the cascade element in the inverse (non-owning) side of the relationship. The cascade element is set to `ALL` in the case of `Order.lineItems`. This means that all persistence operations (deletes, updates, and so on) are cascaded from orders to line items.

Here is the relationship mapping in `Order`:

```
@OneToMany(cascade=ALL, mappedBy="order")
public Collection<LineItem> getLineItems() {
    return lineItems;
}
```

Here is the relationship mapping in `LineItem`:

```
@ManyToOne
public Order getOrder() {
    return order;
}
```

## BLOB and CLOB Database Types in the order Application

The `PARTDETAIL` table in the database has a column, `DRAWING`, of type `BLOB`. `BLOB` stands for binary large objects, which are used for storing binary data such as an image. The `DRAWING` column is mapped to the field `Part.drawing` of type `java.io.Serializable`. The `@Lob` annotation is used to denote that the field is large object.

```
@Column(table="PERSISTENCE_ORDER_PART_DETAIL")
@Lob
public Serializable getDrawing() {
    return drawing;
}
```

`PERSISTENCE_ORDER_PART_DETAIL` also has a column, `SPECIFICATION`, of type `CLOB`. `CLOB` stands for character large objects, which are used to store string data too large to be stored in a

VARCHAR column. SPECIFICATION is mapped to the field `Part.specification` of type `java.lang.String`. The `@Lob` annotation is also used here to denote that the field is a large object.

```
@Column(table="PERSISTENCE_ORDER_PART_DETAIL")
@Lob
public String getSpecification() {
    return specification;
}
```

Both of these fields use the `@Column` annotation and set the `table` element to the secondary table.

## Temporal Types in the order Application

The `Order.lastUpdate` persistent property, which is of type `java.util.Date`, is mapped to the `PERSISTENCE_ORDER_ORDER.LASTUPDATE` database field, which is of the SQL type `TIMESTAMP`. To ensure the proper mapping between these types, you must use the `@Temporal` annotation with the proper temporal type specified in `@Temporal`'s element. `@Temporal`'s elements are of type `javax.persistence.TemporalType`. The possible values are:

- `DATE`, which maps to `java.sql.Date`
- `TIME`, which maps to `java.sql.Time`
- `TIMESTAMP`, which maps to `java.sql.Timestamp`

Here is the relevant section of `Order`:

```
@Temporal(TIMESTAMP)
public Date getLastUpdate() {
    return lastUpdate;
}
```

## Managing the order Application's Entities

The `RequestBean` stateful session bean contains the business logic and manages the entities of `order`.

`RequestBean` uses the `@PersistenceContext` annotation to retrieve an entity manager instance which is used to manage `order`'s entities in `RequestBean`'s business methods.

```
@PersistenceContext
private EntityManager em;
```

This `EntityManager` instance is a container-managed entity manager, so the container takes care of all the transactions involved in the managing `order`'s entities.

## Creating Entities

The `RequestBean.createPart` business method creates a new `Part` entity. The `EntityManager.persist` method is used to persist the newly created entity to the database.

```
Part part = new Part(partNumber,
    revision,
    description,
    revisionDate,
    specification,
    drawing);
em.persist(part);
```

The `ConfigBean` singleton session bean is used to initialize the data in order. `ConfigBean` is annotated with `@Startup`, which indicates that the EJB container should create `ConfigBean` when order is deployed. The `createData` method is annotated with `@PostConstruct`, and it creates the initial entities used by order by calling `RequestsBean`'s business methods.

## Finding Entities

The `RequestBean.getOrderPrice` business method returns the price of a given order, based on the `orderId`. The `EntityManager.find` method is used to retrieve the entity from the database.

```
Order order = em.find(Order.class, orderId);
```

The first argument of `EntityManager.find` is the entity class, and the second is the primary key.

## Setting Entity Relationships

The `RequestBean.createVendorPart` business method creates a `VendorPart` associated with a particular `Vendor`. The `EntityManager.persist` method is used to persist the newly created `VendorPart` entity to the database, and the `VendorPart.setVendor` and `Vendor.setVendorPart` methods are used to associate the `VendorPart` with the `Vendor`.

```
PartKey pkey = new PartKey();
pkey.partNumber = partNumber;
pkey.revision = revision;

Part part = em.find(Part.class, pkey);
VendorPart vendorPart = new VendorPart(description, price,
    part);
em.persist(vendorPart);

Vendor vendor = em.find(Vendor.class, vendorId);
vendor.addVendorPart(vendorPart);
vendorPart.setVendor(vendor);
```

## Using Queries

The `RequestBean.adjustOrderDiscount` business method updates the discount applied to all orders. It uses the `findAllOrders` named query, defined in `Order`:

```
@NamedQuery(
    name="findAllOrders",
    query="SELECT o FROM Order o"
)
```

The `EntityManager.createNamedQuery` method is used to run the query. Because the query returns a `List` of all the orders, the `Query.getResultList` method is used.

```
List orders = em.createNamedQuery(
    "findAllOrders")
    .getResultList();
```

The `RequestBean.getTotalPricePerVendor` business method returns the total price of all the parts for a particular vendor. It uses a named parameter, `id`, defined in the named query `findTotalVendorPartPricePerVendor` defined in `VendorPart`.

```
@NamedQuery(
    name="findTotalVendorPartPricePerVendor",
    query="SELECT SUM(vp.price) " +
    "FROM VendorPart vp " +
    "WHERE vp.vendor.vendorId = :id"
)
```

When running the query, the `Query.setParameter` method is used to set the named parameter `id` to the value of `vendorId`, the parameter to `RequestBean.getTotalPricePerVendor`.

```
return (Double) em.createNamedQuery(
    "findTotalVendorPartPricePerVendor")
    .setParameter("id", vendorId)
    .getSingleResult();
```

The `Query.getSingleResult` method is used for this query because the query returns a single value.

## Removing Entities

The `RequestBean.removeOrder` business method deletes a given order from the database. It uses the `EntityManager.remove` method to delete the entity from the database.

```
Order order = em.find(Order.class, orderId);
em.remove(order);
```

## Building and Running the order Application

This section describes how to build, package, deploy, and run the order application. To do this, you will create the database tables in the JavaDB server, then build, deploy, and run the example.

### Building, Packaging, Deploying, and Running order In NetBeans IDE

Follow these instructions to build, package, deploy, and run the order example to your Enterprise Server instance using NetBeans IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to *tut-install/examples/persistence/*.
3. Select the order folder.
4. Select the Open as Main Project check box.
5. Click Open Project.
6. In the Projects tab, right-click the order project and select Run.

NetBeans will open a web browser to <http://localhost:8080/order/>.

### Building, Packaging, Deploying, and Running order Using Ant

To build the application components of order, enter the following command:

```
ant
```

This runs the default task, which compiles the source files and packages the application into an WAR file located at *tut-install/examples/persistence/order/dist/order.war*.

To deploy the WAR, make sure the Enterprise Server is started, then enter the following command:

```
ant deploy
```

Open a web browser to <http://localhost:8080/order/> to create and update the order data.

### The all Task

As a convenience, the all task will build, package, deploy, and run the application. To do this, enter the following command:

```
ant all
```

### Undeploying order

To undeploy order.war, enter the following command:

```
ant undeploy
```

## The roster Application

The roster application maintains the team rosters for players in recreational sports leagues. The application has four components: Java Persistence API entities (`Player`, `Team`, and `League`), a stateful session bean (`RequestBean`), an application client (`RosterClient`), and three helper classes (`PlayerDetails`, `TeamDetails`, and `LeagueDetails`).

Functionally, roster is similar to the order application described earlier in this chapter with three new features that order does not have: many-to-many relationships, entity inheritance, and automatic table creation at deployment time.

### Relationships in the roster Application

A recreational sports system has the following relationships:

- A player can be on many teams.
- A team can have many players.
- A team is in exactly one league.
- A league has many teams.

In roster this is reflected by the following relationships between the `Player`, `Team`, and `League` entities:

- There is a many-to-many relationship between `Player` and `Team`.
- There is a many-to-one relationship between `Team` and `League`.

### The Many-To-Many Relationship in roster

The many-to-many relationship between `Player` and `Team` is specified by using the `@ManyToMany` annotation.

In `Team.java`, the `@ManyToMany` annotation decorates the `getPlayers` method:

```
@ManyToMany
@JoinTable(
    name="EJB_ROSTER_TEAM_PLAYER",
    joinColumns=
        @JoinColumn(name="TEAM_ID", referencedColumnName="ID"),
    inverseJoinColumns=
        @JoinColumn(name="PLAYER_ID", referencedColumnName="ID")
)
public Collection<Player> getPlayers() {
    return players;
}
```

The `@JoinTable` annotation is used to specify a table in the database that will associate player IDs with team IDs. The entity that specifies the `@JoinTable` is the owner of the relationship, so in this case the `Team` entity is the owner of the relationship with the `Player` entity. Because `roster` uses automatic table creation at deployment time, the container will create a join table in the database named `EJB_ROSTER_TEAM_PLAYER`.

`Player` is the inverse, or non-owning side of the relationship with `Team`. As one-to-one and many-to-one relationships, the non-owning side is marked by the `mappedBy` element in the relationship annotation. Because the relationship between `Player` and `Team` is bidirectional, the choice of which entity is the owner of the relationship is arbitrary.

In `Player.java`, the `@ManyToMany` annotation decorates the `getTeams` method:

```
@ManyToMany(mappedBy="players")
public Collection<Team> getTeams() {
    return teams;
}
```

## Entity Inheritance in the roster Application

The `roster` application demonstrates how to use entity inheritance, as described in [“Entity Inheritance” on page 391](#).

The `League` entity in `roster` is an abstract entity with two concrete subclasses: `SummerLeague` and `WinterLeague`. Because `League` is an abstract class it cannot be instantiated:

```
...
@Entity
@Table(name = "EJB_ROSTER_LEAGUE")
public abstract class League implements java.io.Serializable {
    ...
}
```

Instead, `SummerLeague` or `WinterLeague` are used by clients when creating a league. `SummerLeague` and `WinterLeague` inherit the persistent properties defined in `League`, and only add a constructor that verifies that the sport parameter matches the type of sport allowed in that seasonal league. For example, here is the `SummerLeague` entity:

```
...
@Entity
public class SummerLeague extends League
    implements java.io.Serializable {

    /** Creates a new instance of SummerLeague */
    public SummerLeague() {
    }
}
```

```

public SummerLeague(String id, String name,
    String sport) throws IncorrectSportException {
    this.id = id;
    this.name = name;
    if (sport.equalsIgnoreCase("swimming") ||
        sport.equalsIgnoreCase("soccer") ||
        sport.equalsIgnoreCase("basketball") ||
        sport.equalsIgnoreCase("baseball")) {
        this.sport = sport;
    } else {
        throw new IncorrectSportException(
            "Sport is not a summer sport.");
    }
}
}

```

The roster application uses the default mapping strategy of `InheritanceType.SINGLE_TABLE`, so the `@Inheritance` annotation is not required. If you wanted to use a different mapping strategy, decorate `League` with `@Inheritance` and specify the mapping strategy in the `strategy` element:

```

@Entity
@Inheritance(strategy=JOINED)
@Table(name="EJB_ROSTER_LEAGUE")
public abstract class League implements java.io.Serializable {
    ...
}

```

roster uses the default discriminator column name, so the `@DiscriminatorColumn` annotation is not required. Because you are using automatic table generation in roster the Persistence provider will create a discriminator column in the `EJB_ROSTER_LEAGUE` table called `DTYPE`, which will store the name of the inherited entity used to create the league. If you want to use a different name for the discriminator column, decorate `League` with `@DiscriminatorColumn` and set the `name` element:

```

@Entity
@DiscriminatorColumn(name="DISCRIMINATOR")
@Table(name="EJB_ROSTER_LEAGUE")
public abstract class League implements java.io.Serializable {
    ...
}

```

## Criteria Queries in the roster Application

The roster application uses Criteria API queries, as opposed to the JPQL queries used in order. Criteria queries are Java programming language, type-safe queries defined in the business tier of roster, in the RequestBean stateless session bean.

### Metamodel Classes in the roster Application

Metamodel classes are classes that model an entities attributes, and are used by Criteria queries to navigate to an entities attributes. Each entity class in roster has a corresponding Metamodel class, generated at compile-time, with the same package name as the entity, and appended with an underscore character (\_). For example, the roster.entity.Person entity has a corresponding Metamodel class, roster.entity.Person\_.

Each persistent field or property in the entity class has a corresponding attribute in the entity's Metamodel class. For the Person entity, the corresponding Metamodel class is:

```
@StaticMetamodel(Person.class)
public class Person_ {
    public static volatile SingularAttribute<Player, String> id;
    public static volatile SingularAttribute<Player, String> name;
    public static volatile SingularAttribute<Player, String> position;
    public static volatile SingularAttribute<Player, Double> salary;
    public static volatile CollectionAttribute<Player, Team> teams;
}
```

### Obtaining a CriteriaBuilder Instance in RequestBean

The CrtriteriaBuilder interface defines methods to create criteria query objects, create expressions for modifying those query objects. RequestBean creates an instance of CriteriaBuilder by using a @PostConstruct method, init.

```
@PersistenceContext
private EntityManager em;
private CriteriaBuilder cb;

@PostConstruct
private void init() {
    cb = em.getCriteriaBuilder();
}
```

The EntityManager instance is injected at runtime, and then that EntityManager object is used to create the CriteriaBuilder instance by calling getCriteriaBuilder. The CriteriaBuilder instance is created in a @PostConstruct method to ensure that the EntityManager instance has been injected by the enterprise bean container.

## Creating Criteria Queries in RequestBean's Business Methods

Many of the business methods in RequestBean define Criteria queries. One business method, `getPlayersByPosition`, returns a list of players that play a particular position on a team.

```
public List<PlayerDetails> getPlayersByPosition(String position) {
    logger.info("getPlayersByPosition");
    List<Player> players = null;

    try {
        CriteriaQuery<Player> cq = cb.createQuery(Player.class);
        if (cq != null) {
            Root<Player> player = cq.from(Player.class);

            // set the where clause
            cq.where(cb.equal(player.get(Player_.position), position));
            cq.select(player);
            TypedQuery<Player> q = em.createQuery(cq);
            players = q.getResultList();
        }

        return copyPlayersToDetails(players);
    } catch (Exception ex) {
        throw new EJBException(ex);
    }
}
```

A query object is created by calling the `CriteriaBuilder` object's `createQuery` method, with the type set to `Player` because the query will return a list of players.

The query root is the base entity from which the query will navigate to find the entity's attributes and related entities, and is created by calling the `from` method of the query object. This sets the from clause of the query.

The where clause, set by calling the `where` method on the query object, restricts the results of the query according to the conditions of an expression. The `CriteriaBuilder.equal` method compares the two expressions. In `getPlayersByPosition`, the `position` attribute of the `Player_ Metamodel` class, accessed by calling the `get` method of the query root, is compared to the `position` parameter passed to `getPlayersByPosition`.

The select clause of the query is set by calling the `select` method of the query object. The query will return `Player` entities, so the query root object is passed as a parameter to `select`.

The query object is prepared for execution by calling `EntityManager.createQuery`, which returns a `TypedQuery<T>` object with the type of the query, in this case `Player`. This typed query object is used to execute the query, which occurs when the `getResultList` method is called, and a `List<Player>` collection is returned.

## Automatic Table Generation in the roster Application

At deployment time the Enterprise Server will automatically drop and create the database tables used by roster. This is done by setting the `toplink.ddl-generation` property to `drop-and-create-tables` in `persistence.xml`.

```
<?xml version="1.0" encoding="UTF-8"?>
<persistence xmlns="http://java.sun.com/xml/ns/persistence"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://java.sun.com/xml/ns/persistence
http://java.sun.com/xml/ns/persistence/persistence_1_0.xsd"
version="1.0">
  <persistence-unit name="em" transaction-type="JTA">
    <jta-data-source>jdbc/___default</jta-data-source>
    <properties>
      <property name="toplink.ddl-generation"
                value="drop-and-create-tables"/>
    </properties>
  </persistence-unit>
</persistence>
```

This feature is specific to the Java Persistence API provider used by the Enterprise Server, and is non-portable across Java EE servers. Automatic table creation is useful for development purposes, however, and the `toplink.ddl-generation` property may be removed from `persistence.xml` when preparing the application for production use, or when deploying to other Java EE servers.

## Building and Running the roster Application

This section describes how to build, package, deploy, and run the roster application. You can do this using either NetBeans IDE or Ant.

### Building, Packaging, Deploying, and Running roster in NetBeans IDE

Follow these instructions to build, package, deploy, and run the roster example to your Enterprise Server instance using NetBeans IDE.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to `tut-install/examples/persistence/`.
3. Select the roster folder.
4. Select the Open as Main Project and Open Required Projects check boxes.
5. Click Open Project.
6. In the Projects tab, right-click the roster project and select Run.

You will see the following partial output from the application client in the Output tab:

```
List all players in team T2:  
P6 Ian Carlyle goalkeeper 555.0  
P7 Rebecca Struthers midfielder 777.0  
P8 Anne Anderson forward 65.0  
P9 Jan Wesley defender 100.0  
P10 Terry Smithson midfielder 100.0
```

```
List all teams in league L1:  
T1 Honey Bees Visalia  
T2 Gophers Manteca  
T5 Crows Orland
```

```
List all defenders:  
P2 Alice Smith defender 505.0  
P5 Barney Bold defender 100.0  
P9 Jan Wesley defender 100.0  
P22 Janice Walker defender 857.0  
P25 Frank Fletcher defender 399.0  
...
```

## Building, Packaging, Deploying, and Running roster Using Ant

To build the application components of `roster`, enter the following command:

```
ant
```

This runs the default task, which compiles the source files and packages the application into an EAR file located at `tut-install/examples/persistence/roster/dist/roster.ear`.

To deploy the EAR, make sure the Enterprise Server is started, then enter the following command:

```
ant deploy
```

The build system will check to see if the JavaDB database server is running and start it if it is not running, then deploy `roster.ear`. The Enterprise Server will then drop and create the database tables during deployment, as specified in `persistence.xml`.

After `roster.ear` is deployed, a client JAR, `rosterClient.jar`, is retrieved. This contains the application client.

To run the application client, enter the following command:

```
ant run
```

You will see the output, which begins:

```
[echo] running application client container.  
[exec] List all players in team T2:  
[exec] P6 Ian Carlyle goalkeeper 555.0  
[exec] P7 Rebecca Struthers midfielder 777.0  
[exec] P8 Anne Anderson forward 65.0  
[exec] P9 Jan Wesley defender 100.0  
[exec] P10 Terry Smithson midfielder 100.0  
  
[exec] List all teams in league L1:  
[exec] T1 Honey Bees Visalia  
[exec] T2 Gophers Manteca  
[exec] T5 Crows Orland  
  
[exec] List all defenders:  
[exec] P2 Alice Smith defender 505.0  
[exec] P5 Barney Bold defender 100.0  
[exec] P9 Jan Wesley defender 100.0  
[exec] P22 Janice Walker defender 857.0  
[exec] P25 Frank Fletcher defender 399.0  
...
```

## The all Task

As a convenience, the `all` task will build, package, deploy, and run the application. To do this, enter the following command:

```
ant all
```

## Undeploying order

To undeploy `roster.ear`, enter the following command:

```
ant undeploy
```

# The Java Persistence Query Language

---

The Java Persistence query language defines queries for entities and their persistent state. The query language allows you to write portable queries that work regardless of the underlying data store.

The query language uses the abstract persistence schemas of entities, including their relationships, for its data model, and it defines operators and expressions based on this data model. The scope of a query spans the abstract schemas of related entities that are packaged in the same persistence unit. The query language uses a SQL-like syntax to select objects or values based on entity abstract schema types and relationships among them.

This chapter relies on the material presented in earlier chapters. For conceptual information, see [Chapter 22, “Introduction to the Java Persistence API.”](#) For code examples, see [Chapter 23, “Running the Persistence Examples.”](#)

## Query Language Terminology

The following list defines some of the terms referred to in this chapter.

- **Abstract schema:** The persistent schema abstraction (persistent entities, their state, and their relationships) over which queries operate. The query language translates queries over this persistent schema abstraction into queries that are executed over the database schema to which entities are mapped.
- **Abstract schema type:** All expressions evaluate to a type. The abstract schema type of an entity is derived from the entity class and the metadata information provided by Java language annotations.
- **Backus-Naur Form (BNF):** A notation that describes the syntax of high-level languages. The syntax diagrams in this chapter are in BNF notation.
- **Navigation:** The traversal of relationships in a query language expression. The navigation operator is a period.
- **Path expression:** An expression that navigates to an entity’s state or relationship field.

- **State field:** A persistent field of an entity.
- **Relationship field:** A persistent relationship field of an entity whose type is the abstract schema type of the related entity.

## Creating Queries Using the Java Persistence Query Language

The `EntityManager.createQuery` and `EntityManager.createNamedQuery` methods are used to query the datastore using Java Persistence query language queries.

The `createQuery` method is used to create *dynamic queries*, which are queries that are defined directly within an application's business logic.

```
public List findWithName(String name) {
    return em.createQuery(
        "SELECT c FROM Customer c WHERE c.name LIKE :custName")
        .setParameter("custName", name)
        .setMaxResults(10)
        .getResultList();
}
```

The `createNamedQuery` method is used to create *static queries*, or queries that are defined in metadata using the `javax.persistence.NamedQuery` annotation. The `name` element of `@NamedQuery` specifies the name of the query that will be used with the `createNamedQuery` method. The `query` element of `@NamedQuery` is the query.

```
@NamedQuery(
    name="findAllCustomersWithName",
    query="SELECT c FROM Customer c WHERE c.name LIKE :custName"
)
```

Here's an example of `createNamedQuery`, which uses the `@NamedQuery` defined above.

```
@PersistenceContext
public EntityManager em;
...
customers = em.createNamedQuery("findAllCustomersWithName")
    .setParameter("custName", "Smith")
    .getResultList();
```

## Named Parameters in Queries

Named parameters are parameters in a query that are prefixed with a colon (:). Named parameters in a query are bound to an argument by the `javax.persistence.Query.setParameter(String name, Object value)` method. In the

following example, the name argument to the `findWithName` business method is bound to the `:custName` named parameter in the query by calling `Query.setParameter`.

```
public List findWithName(String name) {
    return em.createQuery(
        "SELECT c FROM Customer c WHERE c.name LIKE :custName")
        .setParameter("custName", name)
        .getResultList();
}
```

Named parameters are case-sensitive, and may be used by both dynamic and static queries.

## Positional Parameters in Queries

You may alternately use positional parameters in queries, instead of named parameters. Positional parameters are prefixed with a question mark (?) followed the numeric position of the parameter in the query. The `Query.setParameter(integer position, Object value)` method is used to set the parameter values.

In the following example, the `findWithName` business method is rewritten to use input parameters:

```
public List findWithName(String name) {
    return em.createQuery(
        "SELECT c FROM Customer c WHERE c.name LIKE ?1")
        .setParameter(1, name)
        .getResultList();
}
```

Input parameters are numbered starting from 1. Input parameters are case-sensitive, and may be used by both dynamic and static queries.

## Simplified Query Language Syntax

This section briefly describes the syntax of the query language so that you can quickly move on to the next section, [“Example Queries” on page 426](#). When you are ready to learn about the syntax in more detail, see the section [“Full Query Language Syntax” on page 432](#).

### Select Statements

A select query has six clauses: `SELECT`, `FROM`, `WHERE`, `GROUP BY`, `HAVING`, and `ORDER BY`. The `SELECT` and `FROM` clauses are required, but the `WHERE`, `GROUP BY`, `HAVING`, and `ORDER BY` clauses are optional. Here is the high-level BNF syntax of a query language query:

```
QL_statement ::= select_clause from_clause  
              [where_clause][groupby_clause][having_clause][orderby_clause]
```

The **SELECT** clause defines the types of the objects or values returned by the query.

The **FROM** clause defines the scope of the query by declaring one or more identification variables, which can be referenced in the **SELECT** and **WHERE** clauses. An identification variable represents one of the following elements:

- The abstract schema name of an entity
- An element of a collection relationship
- An element of a single-valued relationship
- A member of a collection that is the multiple side of a one-to-many relationship

The **WHERE** clause is a conditional expression that restricts the objects or values retrieved by the query. Although it is optional, most queries have a **WHERE** clause.

The **GROUP BY** clause groups query results according to a set of properties.

The **HAVING** clause is used with the **GROUP BY** clause to further restrict the query results according to a conditional expression.

The **ORDER BY** clause sorts the objects or values returned by the query into a specified order.

## Update and Delete Statements

Update and delete statements provide bulk operations over sets of entities. They have the following syntax:

```
update_statement ::= update_clause [where_clause] delete_statement ::=  
                    delete_clause [where_clause]
```

The update and delete clauses determine the type of the entities to be updated or deleted. The **WHERE** clause may be used to restrict the scope of the update or delete operation.

## Example Queries

The following queries are from the `Player` entity of the `roster` application, which is documented in [“The roster Application” on page 415](#).

### Simple Queries

If you are unfamiliar with the query language, these simple queries are a good place to start.

## A Basic Select Query

```
SELECT p
FROM Player p
```

**Data retrieved:** All players.

**Description:** The FROM clause declares an identification variable named `p`, omitting the optional keyword AS. If the AS keyword were included, the clause would be written as follows:

```
FROM Player AS
  p
```

The `Player` element is the abstract schema name of the `Player` entity.

**See also:** [“Identification Variables” on page 438](#)

## Eliminating Duplicate Values

```
SELECT DISTINCT
  p
FROM Player p
WHERE p.position = ?1
```

**Data retrieved:** The players with the position specified by the query’s parameter.

**Description:** The DISTINCT keyword eliminates duplicate values.

The WHERE clause restricts the players retrieved by checking their `position`, a persistent field of the `Player` entity. The `?1` element denotes the input parameter of the query.

**See also:** [“Input Parameters” on page 443](#), [“The DISTINCT Keyword” on page 453](#)

## Using Named Parameters

```
SELECT DISTINCT p
FROM Player p
WHERE p.position = :position AND p.name = :name
```

**Data retrieved:** The players having the specified positions and names.

**Description:** The `position` and `name` elements are persistent fields of the `Player` entity. The WHERE clause compares the values of these fields with the named parameters of the query, set using the `Query.setNamedParameter` method. The query language denotes a named input parameter using colon (`:`) followed by an identifier. The first input parameter is `:position`, the second is `:name`.

## Queries That Navigate to Related Entities

In the query language, an expression can traverse (or navigate) to related entities. These expressions are the primary difference between the Java Persistence query language and SQL. Queries navigates to related entities, whereas SQL joins tables.

### A Simple Query with Relationships

```
SELECT DISTINCT p
FROM Player p, IN(p.teams) t
```

**Data retrieved:** All players who belong to a team.

**Description:** The FROM clause declares two identification variables: `p` and `t`. The `p` variable represents the `Player` entity, and the `t` variable represents the related `Team` entity. The declaration for `t` references the previously declared `p` variable. The `IN` keyword signifies that `teams` is a collection of related entities. The `p.teams` expression navigates from a `Player` to its related `Team`. The period in the `p.teams` expression is the navigation operator.

You may also use the `JOIN` statement to write the same query:

```
SELECT DISTINCT p
FROM Player p JOIN p.teams t
```

This query could also be rewritten as:

```
SELECT DISTINCT p
FROM Player p
WHERE p.team IS NOT EMPTY
```

### Navigating to Single-Valued Relationship Fields

Use the `JOIN` clause statement to navigate to a single-valued relationship field:

```
SELECT t
FROM Team t JOIN t.league l
WHERE l.sport = 'soccer' OR l.sport = 'football'
```

In this example, the query will return all teams that are in either soccer or football leagues.

### Traversing Relationships with an Input Parameter

```
SELECT DISTINCT p
FROM Player p, IN (p.teams) AS t
WHERE t.city = :city
```

**Data retrieved:** The players whose teams belong to the specified city.

**Description:** This query is similar to the previous example, but it adds an input parameter. The `AS` keyword in the `FROM` clause is optional. In the `WHERE` clause, the period preceding the persistent variable `city` is a delimiter, not a navigation operator. Strictly speaking, expressions can navigate to relationship fields (related entities), but not to persistent fields. To access a persistent field, an expression uses the period as a delimiter.

Expressions cannot navigate beyond (or further qualify) relationship fields that are collections. In the syntax of an expression, a collection-valued field is a terminal symbol. Because the `teams` field is a collection, the `WHERE` clause cannot specify `p.teams.city` (an illegal expression).

See also: “[Path Expressions](#)” on page 441

## Traversing Multiple Relationships

```
SELECT DISTINCT p
FROM Player p, IN (p.teams) t
WHERE t.league = :league
```

**Data retrieved:** The players that belong to the specified league.

**Description:** The expressions in this query navigate over two relationships. The `p.teams` expression navigates the `Player-Team` relationship, and the `t.league` expression navigates the `Team-League` relationship.

In the other examples, the input parameters are `String` objects, but in this example the parameter is an object whose type is a `League`. This type matches the `league` relationship field in the comparison expression of the `WHERE` clause.

## Navigating According to Related Fields

```
SELECT DISTINCT p
FROM Player p, IN (p.teams) t
WHERE t.league.sport = :sport
```

**Data retrieved:** The players who participate in the specified sport.

**Description:** The `sport` persistent field belongs to the `League` entity. To reach the `sport` field, the query must first navigate from the `Player` entity to `Team` (`p.teams`) and then from `Team` to the `League` entity (`t.league`). Because the `league` relationship field is not a collection, it can be followed by the `sport` persistent field.

## Queries with Other Conditional Expressions

Every `WHERE` clause must specify a conditional expression, of which there are several kinds. In the previous examples, the conditional expressions are comparison expressions that test for equality. The following examples demonstrate some of the other kinds of conditional expressions. For descriptions of all conditional expressions, see the section “[WHERE Clause](#)” on page 442.

## The LIKE Expression

```
SELECT p
FROM Player p
WHERE p.name LIKE 'Mich%'
```

**Data retrieved:** All players whose names begin with “Mich.”

**Description:** The LIKE expression uses wildcard characters to search for strings that match the wildcard pattern. In this case, the query uses the LIKE expression and the % wildcard to find all players whose names begin with the string “Mich.” For example, “Michael” and “Michelle” both match the wildcard pattern.

**See also:** [“LIKE Expressions” on page 445](#)

## The IS NULL Expression

```
SELECT t
FROM Team t
WHERE t.league IS NULL
```

**Data retrieved:** All teams not associated with a league.

**Description:** The IS NULL expression can be used to check if a relationship has been set between two entities. In this case, the query checks to see if the teams are associated with any leagues, and returns the teams that do not have a league.

**See also:** [“NULL Comparison Expressions” on page 446](#), [“NULL Values” on page 450](#)

## The IS EMPTY Expression

```
SELECT p
FROM Player p
WHERE p.teams IS EMPTY
```

**Data retrieved:** All players who do not belong to a team.

**Description:** The teams relationship field of the Player entity is a collection. If a player does not belong to a team, then the teams collection is empty and the conditional expression is TRUE.

**See also:** [“Empty Collection Comparison Expressions” on page 446](#)

## The BETWEEN Expression

```
SELECT DISTINCT p
FROM Player p
WHERE p.salary BETWEEN :lowerSalary AND :higherSalary
```

**Data retrieved:** The players whose salaries fall within the range of the specified salaries.

**Description:** This BETWEEN expression has three arithmetic expressions: a persistent field (`p.salary`) and the two input parameters (`:lowerSalary` and `:higherSalary`). The following expression is equivalent to the BETWEEN expression:

```
p.salary >= :lowerSalary AND p.salary <= :higherSalary
```

See also: [“BETWEEN Expressions” on page 444](#)

## Comparison Operators

```
SELECT DISTINCT p1
FROM Player p1, Player p2
WHERE p1.salary > p2.salary AND p2.name = :name
```

**Data retrieved:** All players whose salaries are higher than the salary of the player with the specified name.

**Description:** The FROM clause declares two identification variables (`p1` and `p2`) of the same type (Player). Two identification variables are needed because the WHERE clause compares the salary of one player (`p2`) with that of the other players (`p1`).

See also: [“Identification Variables” on page 438](#)

## Bulk Updates and Deletes

The following examples show how to use the UPDATE and DELETE expressions in queries. UPDATE and DELETE operate on multiple entities according to the condition or conditions set in the WHERE clause. The WHERE clause in UPDATE and DELETE queries follows the same rules as SELECT queries.

### Update Queries

```
UPDATE Player p
SET p.status = 'inactive'
WHERE p.lastPlayed < :inactiveThresholdDate
```

**Description:** This query sets the status of a set of players to inactive if the player’s last game was longer than the date specified in `inactiveThresholdDate`.

### Delete Queries

```
DELETE
FROM Player p
WHERE p.status = 'inactive'
AND p.teams IS EMPTY
```

**Description:** This query deletes all inactive players who are not on a team.

# Full Query Language Syntax

This section discusses the query language syntax, as defined in the Java Persistence specification. Much of the following material paraphrases or directly quotes the specification.

## BNF Symbols

Table 24–1 describes the BNF symbols used in this chapter.

TABLE 24–1 BNF Symbol Summary

Symbol	Description
::=	The element to the left of the symbol is defined by the constructs on the right.
*	The preceding construct may occur zero or more times.
{...}	The constructs within the curly braces are grouped together.
[...]	The constructs within the square brackets are optional.
	An exclusive OR.
<b>BOLDFACE</b>	A keyword (although capitalized in the BNF diagram, keywords are not case-sensitive).
White space	A white space character can be a space, a horizontal tab, or a line feed.

## BNF Grammar of the Java Persistence Query Language

Here is the entire BNF diagram for the query language:

```

QL_statement ::= select_statement | update_statement | delete_statement
select_statement ::= select_clause from_clause [where_clause] [groupby_clause]
                  [having_clause] [orderby_clause]
update_statement ::= update_clause [where_clause]
delete_statement ::= delete_clause [where_clause]
from_clause ::=
    FROM identification_variable_declaration
        {, {identification_variable_declaration |
            collection_member_declaration}}*
identification_variable_declaration ::=
    range_variable_declaration { join | fetch_join }*
range_variable_declaration ::= abstract_schema_name [AS]
    identification_variable
join ::= join_spec join_association_path_expression [AS]
    identification_variable
fetch_join ::= join_specFETCH join_association_path_expression

```

```

association_path_expression ::=
    collection_valued_path_expression |
    single_valued_association_path_expression
join_spec ::= [LEFT [OUTER] |INNER] JOIN
join_association_path_expression ::=
    join_collection_valued_path_expression |
    join_single_valued_association_path_expression
join_collection_valued_path_expression ::=
    identification_variable.collection_valued_association_field
join_single_valued_association_path_expression ::=
    identification_variable.single_valued_association_field
collection_member_declaration ::=
    IN (collection_valued_path_expression) [AS]
    identification_variable
single_valued_path_expression ::=
    state_field_path_expression |
    single_valued_association_path_expression
state_field_path_expression ::=
    {identification_variable |
    single_valued_association_path_expression}.state_field
single_valued_association_path_expression ::=
    identification_variable.{single_valued_association_field}*
    single_valued_association_field
collection_valued_path_expression ::=
    identification_variable.{single_valued_association_field}*
    collection_valued_association_field
state_field ::=
    {embedded_class_state_field.*}simple_state_field
update_clause ::= UPDATE abstract_schema_name [[AS]
    identification_variable] SET update_item {, update_item}*
update_item ::= [identification_variable.]{state_field |
    single_valued_association_field} = new_value
new_value ::=
    simple_arithmetic_expression |
    string_primary |
    datetime_primary |
    boolean_primary |
    enum_primary simple_entity_expression |
    NULL
delete_clause ::= DELETE FROM abstract_schema_name [[AS]
    identification_variable]
select_clause ::= SELECT [DISTINCT] select_expression {,
    select_expression}*
select_expression ::=
    single_valued_path_expression |
    aggregate_expression |
    identification_variable |
    OBJECT(identification_variable) |

```

```
    constructor_expression
constructor_expression ::=
    NEW constructor_name(constructor_item {,
    constructor_item}*)
constructor_item ::= single_valued_path_expression |
    aggregate_expression
aggregate_expression ::=
    {AVG |MAX |MIN |SUM} ([DISTINCT]
    state_field_path_expression) |
    COUNT ([DISTINCT] identification_variable |
    state_field_path_expression |
    single_valued_association_path_expression)
where_clause ::= WHERE conditional_expression
groupby_clause ::= GROUP BY groupby_item {, groupby_item}*
groupby_item ::= single_valued_path_expression
having_clause ::= HAVING conditional_expression
orderby_clause ::= ORDER BY orderby_item {, orderby_item}*
orderby_item ::= state_field_path_expression [ASC |DESC]
subquery ::= simple_select_clause subquery_from_clause
    [where_clause] [groupby_clause] [having_clause]
subquery_from_clause ::=
    FROM subselect_identification_variable_declaration
    {, subselect_identification_variable_declaration}*
subselect_identification_variable_declaration ::=
    identification_variable_declaration |
    association_path_expression [AS] identification_variable |
    collection_member_declaration
simple_select_clause ::= SELECT [DISTINCT]
    simple_select_expression
simple_select_expression ::=
    single_valued_path_expression |
    aggregate_expression |
    identification_variable
conditional_expression ::= conditional_term |
    conditional_expression OR conditional_term
conditional_term ::= conditional_factor | conditional_term AND
    conditional_factor
conditional_factor ::= [NOT] conditional_primary
conditional_primary ::= simple_cond_expression |(
    conditional_expression)
simple_cond_expression ::=
    comparison_expression |
    between_expression |
    like_expression |
    in_expression |
    null_comparison_expression |
    empty_collection_comparison_expression |
    collection_member_expression |
```

```

exists_expression
between_expression ::=
    arithmetic_expression [NOT] BETWEEN
        arithmetic_expression AND arithmetic_expression |
    string_expression [NOT] BETWEEN string_expression AND
        string_expression |
    datetime_expression [NOT] BETWEEN
        datetime_expression AND datetime_expression
in_expression ::=
    state_field_path_expression [NOT] IN (in_item {, in_item}*
        | subquery)
in_item ::= literal | input_parameter
like_expression ::=
    string_expression [NOT] LIKE pattern_value [ESCAPE
        escape_character]
null_comparison_expression ::=
    {single_valued_path_expression | input_parameter} IS [NOT]
        NULL
empty_collection_comparison_expression ::=
    collection_valued_path_expression IS [NOT] EMPTY
collection_member_expression ::= entity_expression
    [NOT] MEMBER [OF] collection_valued_path_expression
exists_expression ::= [NOT] EXISTS (subquery)
all_or_any_expression ::= {ALL |ANY |SOME} (subquery)
comparison_expression ::=
    string_expression comparison_operator {string_expression |
    all_or_any_expression} |
    boolean_expression {= |<> } {boolean_expression |
    all_or_any_expression} |
    enum_expression {= |<> } {enum_expression |
    all_or_any_expression} |
    datetime_expression comparison_operator
        {datetime_expression | all_or_any_expression} |
    entity_expression {= |<> } {entity_expression |
    all_or_any_expression} |
    arithmetic_expression comparison_operator
        {arithmetic_expression | all_or_any_expression}
comparison_operator ::= = |> |>= |< |<= |<>
arithmetic_expression ::= simple_arithmetic_expression |
    (subquery)
simple_arithmetic_expression ::=
    arithmetic_term | simple_arithmetic_expression {+ |- }
        arithmetic_term
arithmetic_term ::= arithmetic_factor | arithmetic_term {* |/ }
    arithmetic_factor
arithmetic_factor ::= [{+ |- }] arithmetic_primary
arithmetic_primary ::=
    state_field_path_expression |

```

```
    numeric_literal |
    (simple_arithmetic_expression) |
    input_parameter |
    functions_returning_numerics |
    aggregate_expression
string_expression ::= string_primary | (subquery)
string_primary ::=
    state_field_path_expression |
    string_literal |
    input_parameter |
    functions_returning_strings |
    aggregate_expression
datetime_expression ::= datetime_primary | (subquery)
datetime_primary ::=
    state_field_path_expression |
    input_parameter |
    functions_returning_datetime |
    aggregate_expression
boolean_expression ::= boolean_primary | (subquery)
boolean_primary ::=
    state_field_path_expression |
    boolean_literal |
    input_parameter
enum_expression ::= enum_primary | (subquery)
enum_primary ::=
    state_field_path_expression |
    enum_literal |
    input_parameter
entity_expression ::=
    single_valued_association_path_expression |
    simple_entity_expression
simple_entity_expression ::=
    identification_variable |
    input_parameter
functions_returning_numerics ::=
    LENGTH(string_primary) |
    LOCATE(string_primary, string_primary[,
        simple_arithmetic_expression]) |
    ABS(simple_arithmetic_expression) |
    SQRT(simple_arithmetic_expression) |
    MOD(simple_arithmetic_expression,
        simple_arithmetic_expression) |
    SIZE(collection_valued_path_expression)
functions_returning_datetime ::=
    CURRENT_DATE |
    CURRENT_TIME |
    CURRENT_TIMESTAMP
functions_returning_strings ::=
```

```

CONCAT(string_primary, string_primary) |
SUBSTRING(string_primary,
           simple_arithmetic_expression,
           simple_arithmetic_expression) |
TRIM([[trim_specification] [trim_character] FROM]
      string_primary) |
LOWER(string_primary) |
UPPER(string_primary)
trim_specification ::= LEADING | TRAILING | BOTH

```

## FROM Clause

The FROM clause defines the domain of the query by declaring identification variables.

### Identifiers

An identifier is a sequence of one or more characters. The first character must be a valid first character (letter, \$, \_) in an identifier of the Java programming language (hereafter in this chapter called simply “Java”). Each subsequent character in the sequence must be a valid non-first character (letter, digit, \$, \_) in a Java identifier. (For details, see the Java SE API documentation of the `isJavaIdentifierStart` and `isJavaIdentifierPart` methods of the `Character` class.) The question mark (?) is a reserved character in the query language and cannot be used in an identifier.

A query language identifier is case-sensitive with two exceptions:

- Keywords
- Identification variables

An identifier cannot be the same as a query language keyword. Here is a list of query language keywords:

ABS	ALL	AND	ANY	AS
ASC	AVG	BETWEEN	BIT_LENGTH	BOTH
BY	CASE	CHAR_LENGTH	CHARACTER_LENGTH	CLASS
COALESCE	CONCAT	COUNT	CURRENT_DATE	CURRENT_TIMESTAMP
DELETE	DESC	DISTINCT	ELSE	EMPTY
END	ENTRY	ESCAPE	EXISTS	FALSE
FETCH	FROM	GROUP	HAVING	IN
INDEX	INNER	IS	JOIN	KEY

LEADING	LEFT	LENGTH	LIKE	LOCATE
LOWER	MAX	MEMBER	MIN	MOD
NEW	NOT	NULL	NULLIF	OBJECT
OF	OR	ORDER	OUTER	POSITION
SELECT	SET	SIZE	SOME	SQRT
SUBSTRING	SUM	THEN	TRAILING	TRIM
TRUE	TYPE	UNKNOWN	UPDATE	UPPER
VALUE	WHEN	WHERE		

It is not recommended that you use a SQL keyword as an identifier, because the list of keywords may expand to include other reserved SQL words in the future.

## Identification Variables

An *identification variable* is an identifier declared in the FROM clause. Although the SELECT and WHERE clauses can reference identification variables, they cannot declare them. All identification variables must be declared in the FROM clause.

Because an identification variable is an identifier, it has the same naming conventions and restrictions as an identifier with the exception that an identification variables is case-insensitive. For example, an identification variable cannot be the same as a query language keyword. (See the preceding section for more naming rules.) Also, within a given persistence unit, an identification variable name must not match the name of any entity or abstract schema.

The FROM clause can contain multiple declarations, separated by commas. A declaration can reference another identification variable that has been previously declared (to the left). In the following FROM clause, the variable `t` references the previously declared variable `p`:

```
FROM Player p, IN (p.teams) AS t
```

Even if an identification variable is not used in the WHERE clause, its declaration can affect the results of the query. For an example, compare the next two queries. The following query returns all players, whether or not they belong to a team:

```
SELECT p  
FROM Player p
```

In contrast, because the next query declares the `t` identification variable, it fetches all players that belong to a team:

```
SELECT p  
FROM Player p, IN (p.teams) AS t
```

The following query returns the same results as the preceding query, but the `WHERE` clause makes it easier to read:

```
SELECT p
FROM Player p
WHERE p.teams IS NOT EMPTY
```

An identification variable always designates a reference to a single value whose type is that of the expression used in the declaration. There are two kinds of declarations: range variable and collection member.

## Range Variable Declarations

To declare an identification variable as an abstract schema type, you specify a range variable declaration. In other words, an identification variable can range over the abstract schema type of an entity. In the following example, an identification variable named `p` represents the abstract schema named `Player`:

```
FROM Player p
```

A range variable declaration can include the optional `AS` operator:

```
FROM Player AS p
```

In most cases, to obtain objects a query uses path expressions to navigate through the relationships. But for those objects that cannot be obtained by navigation, you can use a range variable declaration to designate a starting point (or *root*).

If the query compares multiple values of the same abstract schema type, then the `FROM` clause must declare multiple identification variables for the abstract schema:

```
FROM Player p1, Player p2
```

For a sample of such a query, see [“Comparison Operators” on page 431](#).

## Collection Member Declarations

In a one-to-many relationship, the multiple side consists of a collection of entities. An identification variable can represent a member of this collection. To access a collection member, the path expression in the variable’s declaration navigates through the relationships in the abstract schema. (For more information on path expressions, see the following section.) Because a path expression can be based on another path expression, the navigation can traverse several relationships. See [“Traversing Multiple Relationships” on page 429](#).

A collection member declaration must include the `IN` operator, but it can omit the optional `AS` operator.

In the following example, the entity represented by the abstract schema named `Player` has a relationship field called `teams`. The identification variable called `t` represents a single member of the `teams` collection.

```
FROM Player p, IN (p.teams) t
```

## Joins

The `JOIN` operator is used to traverse over relationships between entities, and is functionally similar to the `IN` operator.

In the following example, the query joins over the relationship between customers and orders:

```
SELECT c
  FROM Customer c JOIN c.orders o
 WHERE c.status = 1 AND o.totalPrice > 10000
```

The `INNER` keyword is optional:

```
SELECT c
  FROM Customer c INNER JOIN c.orders o
 WHERE c.status = 1 AND o.totalPrice > 10000
```

These examples are equivalent to the following query, which uses the `IN` operator:

```
SELECT c
  FROM Customer c, IN(c.orders) o
 WHERE c.status = 1 AND o.totalPrice > 10000
```

You can also join a single-valued relationship.

```
SELECT t
  FROM Team t JOIN t.league l
 WHERE l.sport = :sport
```

A `LEFT JOIN` or `LEFT OUTER JOIN` retrieves a set of entities where matching values in the join condition may be absent. The `OUTER` keyword is optional.

```
SELECT c.name, o.totalPrice
  FROM Order o LEFT JOIN o.customer c
```

A `FETCH JOIN` is a join operation that returns associated entities as a side-effect of running the query. In the following example, the query returns a set of departments, and as a side-effect, the associated employees of the departments, even though the employees were not explicitly retrieved by the `SELECT` clause.

```
SELECT d
FROM Department d LEFT JOIN FETCH d.employees
WHERE d.deptno = 1
```

## Path Expressions

Path expressions are important constructs in the syntax of the query language, for several reasons. First, they define navigation paths through the relationships in the abstract schema. These path definitions affect both the scope and the results of a query. Second, they can appear in any of the main clauses of a query (SELECT, DELETE, HAVING, UPDATE, WHERE, FROM, GROUP BY, ORDER BY). Finally, although much of the query language is a subset of SQL, path expressions are extensions not found in SQL.

### Examples of Path Expressions

Here, the WHERE clause contains a `single_valued_path_expression`. The `p` is an identification variable, and `salary` is a persistent field of `Player`.

```
SELECT DISTINCT p
FROM Player p
WHERE p.salary BETWEEN :lowerSalary AND :higherSalary
```

Here, the WHERE clause also contains a `single_valued_path_expression`. The `t` is an identification variable, `league` is a single-valued relationship field, and `sport` is a persistent field of `league`.

```
SELECT DISTINCT p
FROM Player p, IN (p.teams) t
WHERE t.league.sport = :sport
```

Here, the WHERE clause contains a `collection_valued_path_expression`. The `p` is an identification variable, and `teams` designates a collection-valued relationship field.

```
SELECT DISTINCT p
FROM Player p
WHERE p.teams IS EMPTY
```

### Expression Types

The type of a path expression is the type of the object represented by the ending element, which can be one of the following:

- Persistent field
- Single-valued relationship field
- Collection-valued relationship field

For example, the type of the expression `p.salary` is `double` because the terminating persistent field (`salary`) is a `double`.

In the expression `p.teams`, the terminating element is a collection-valued relationship field (`teams`). This expression's type is a collection of the abstract schema type named `Team`. Because `Team` is the abstract schema name for the `Team` entity, this type maps to the entity. For more information on the type mapping of abstract schemas, see the section [“Return Types” on page 452](#).

## Navigation

A path expression enables the query to navigate to related entities. The terminating elements of an expression determine whether navigation is allowed. If an expression contains a single-valued relationship field, the navigation can continue to an object that is related to the field. However, an expression cannot navigate beyond a persistent field or a collection-valued relationship field. For example, the expression `p.teams.league.sport` is illegal, because `teams` is a collection-valued relationship field. To reach the `sport` field, the `FROM` clause could define an identification variable named `t` for the `teams` field:

```
FROM Player AS p, IN (p.teams) t
WHERE t.league.sport = 'soccer'
```

## WHERE Clause

The `WHERE` clause specifies a conditional expression that limits the values returned by the query. The query returns all corresponding values in the data store for which the conditional expression is `TRUE`. Although usually specified, the `WHERE` clause is optional. If the `WHERE` clause is omitted, then the query returns all values. The high-level syntax for the `WHERE` clause follows:

```
where_clause ::= WHERE conditional_expression
```

## Literals

There are four kinds of literals: string, numeric, Boolean, and enum.

### String Literals

A string literal is enclosed in single quotes:

```
'Duke'
```

If a string literal contains a single quote, you indicate the quote by using two single quotes:

```
'Duke''s'
```

Like a Java `String`, a string literal in the query language uses the Unicode character encoding.

## Numeric Literals

There are two types of numeric literals: exact and approximate.

An exact numeric literal is a numeric value without a decimal point, such as 65, -233, and +12. Using the Java integer syntax, exact numeric literals support numbers in the range of a Java `long`.

An approximate numeric literal is a numeric value in scientific notation, such as 57., -85.7, and +2.1. Using the syntax of the Java floating-point literal, approximate numeric literals support numbers in the range of a Java `double`.

## Boolean Literals

A Boolean literal is either `TRUE` or `FALSE`. These keywords are not case-sensitive.

## Enum Literals

The Java Persistence Query Language supports the use of enum literals using the Java enum literal syntax. The enum class name must be specified as fully qualified class name.

```
SELECT e
FROM Employee e
WHERE e.status = com.xyz.EmployeeStatus.FULL_TIME
```

## Input Parameters

An input parameter can be either a named parameter or a positional parameter.

A named input parameter is designated by a colon (`:`) followed by a string. For example, `:name`.

A positional input parameter is designated by a question mark (`?`) followed by an integer. For example, the first input parameter is `?1`, the second is `?2`, and so forth.

The following rules apply to input parameters:

- They can be used only in a `WHERE` or `HAVING` clause.
- Positional parameters must be numbered, starting with the integer 1.
- Named parameters and positional parameters may not be mixed in a single query.
- Named parameters are case-sensitive.

## Conditional Expressions

A `WHERE` clause consists of a conditional expression, which is evaluated from left to right within a precedence level. You can change the order of evaluation by using parentheses.

## Operators and Their Precedence

Table 24–2 lists the query language operators in order of decreasing precedence.

TABLE 24-2 Query Language Order Precedence

Type	Precedence Order
Navigation	. (a period)
Arithmetic	+ - (unary) * / (multiplication and division) + - (addition and subtraction)
Comparison	= > >= < <= <> (not equal) [NOT] BETWEEN [NOT] LIKE [NOT] IN IS [NOT] NULL IS [NOT] EMPTY [NOT] MEMBER OF
Logical	NOT AND OR

---

## BETWEEN Expressions

A **BETWEEN** expression determines whether an arithmetic expression falls within a range of values.

These two expressions are equivalent:

```
p.age BETWEEN 15 AND 19  
p.age >= 15 AND p.age <= 19
```

The following two expressions are also equivalent:

```
p.age NOT BETWEEN 15 AND 19  
p.age < 15 OR p.age > 19
```

If an arithmetic expression has a NULL value, then the value of the BETWEEN expression is unknown.

## IN Expressions

An IN expression determines whether or not a string belongs to a set of string literals, or whether a number belongs to a set of number values.

The path expression must have a string or numeric value. If the path expression has a NULL value, then the value of the IN expression is unknown.

In the following example, if the country is UK the expression is TRUE. If the country is Peru it is FALSE.

```
o.country IN ('UK', 'US', 'France')
```

You may also use input parameters:

```
o.country IN ('UK', 'US', 'France', :country)
```

## LIKE Expressions

A LIKE expression determines whether a wildcard pattern matches a string.

The path expression must have a string or numeric value. If this value is NULL, then the value of the LIKE expression is unknown. The pattern value is a string literal that can contain wildcard characters. The underscore ( `_` ) wildcard character represents any single character. The percent ( `%` ) wildcard character represents zero or more characters. The ESCAPE clause specifies an escape character for the wildcard characters in the pattern value. [Table 24–3](#) shows some sample LIKE expressions.

TABLE 24–3 LIKE Expression Examples

Expression	TRUE	FALSE
<code>address.phone LIKE '12%3'</code>	<code>'123'</code> <code>'12993'</code>	<code>'1234'</code>
<code>asentence.word LIKE 'l_se'</code>	<code>'lose'</code>	<code>'loose'</code>
<code>aword.underscored LIKE '\_%' ESCAPE '\'</code>	<code>'_foo'</code>	<code>'bar'</code>
<code>address.phone NOT LIKE '12%3'</code>	<code>'1234'</code>	<code>'123'</code> <code>'12993'</code>

## NULL Comparison Expressions

A NULL comparison expression tests whether a single-valued path expression or an input parameter has a NULL value. Usually, the NULL comparison expression is used to test whether or not a single-valued relationship has been set.

```
SELECT t
FROM Team t
WHERE t.league IS NULL
```

This query selects all teams where the league relationship is not set. Please note, the following query is *not* equivalent:

```
SELECT t
FROM Team t
WHERE t.league = NULL
```

The comparison with NULL using the equals operator (=) always returns an unknown value, even if the relationship is not set. The second query will always return an empty result.

## Empty Collection Comparison Expressions

The IS [NOT] EMPTY comparison expression tests whether a collection-valued path expression has no elements. In other words, it tests whether or not a collection-valued relationship has been set.

If the collection-valued path expression is NULL, then the empty collection comparison expression has a NULL value.

Here is an example that finds all orders that do not have any line items:

```
SELECT o
FROM Order o
WHERE o.lineItems IS EMPTY
```

## Collection Member Expressions

The [NOT] MEMBER [OF] collection member expression determines whether a value is a member of a collection. The value and the collection members must have the same type.

If either the collection-valued or single-valued path expression is unknown, then the collection member expression is unknown. If the collection-valued path expression designates an empty collection, then the collection member expression is FALSE.

The OF keyword is optional.

The following example tests whether a line item is part of an order:

```
SELECT o
FROM Order o
WHERE :lineItem MEMBER OF o.lineItems
```

## Subqueries

Subqueries may be used in the `WHERE` or `HAVING` clause of a query. Subqueries must be surrounded by parentheses.

The following example find all customers who have placed more than 10 orders:

```
SELECT c
FROM Customer c
WHERE (SELECT COUNT(o) FROM c.orders o) > 10
```

## EXISTS Expressions

The `[NOT] EXISTS` expression is used with a subquery, and is true only if the result of the subquery consists of one or more values and is false otherwise.

The following example finds all employees whose spouse is also an employee:

```
SELECT DISTINCT emp
FROM Employee emp
WHERE EXISTS (
    SELECT spouseEmp
    FROM Employee spouseEmp
    WHERE spouseEmp = emp.spouse)
```

## ALL and ANY Expressions

The `ALL` expression is used with a subquery, and is true if all the values returned by the subquery are true, or if the subquery is empty.

The `ANY` expression is used with a subquery, and is true if some of the values returned by the subquery are true. An `ANY` expression is false if the subquery result is empty, or if all the values returned are false. The `SOME` keyword is synonymous with `ANY`.

The `ALL` and `ANY` expressions are used with the `=`, `<`, `<=`, `>`, `>=`, `<>` comparison operators.

The following example finds all employees whose salary is higher than the salary of the managers in the employee's department:

```
SELECT emp
FROM Employee emp
WHERE emp.salary > ALL (
    SELECT m.salary
    FROM Manager m
    WHERE m.department = emp.department)
```

## Functional Expressions

The query language includes several string, arithmetic, and date/time functions which may be used in the `SELECT`, `WHERE`, or `HAVING` clause of a query. The functions are listed in the following tables. In [Table 24–4](#), the `start` and `length` arguments are of type `int`. They designate positions in the `String` argument. The first position in a string is designated by 1. In [Table 24–5](#), the `number` argument can be either an `int`, a `float`, or a `double`.

TABLE 24–4 String Expressions

Function Syntax	Return Type
<code>CONCAT(String, String)</code>	<code>String</code>
<code>LENGTH(String)</code>	<code>int</code>
<code>LOCATE(String, String [, start])</code>	<code>int</code>
<code>SUBSTRING(String, start, length)</code>	<code>String</code>
<code>TRIM([[LEADING TRAILING BOTH] char) FROM] (String)</code>	<code>String</code>
<code>LOWER(String)</code>	<code>String</code>
<code>UPPER(String)</code>	<code>String</code>

The `CONCAT` function concatenates two strings into one string.

The `LENGTH` function returns the length of a string in characters as an integer.

The `LOCATE` function returns the position of a given string within a string. It returns the first position at which the string was found as an integer. The first argument is the string to be located. The second argument is the string to be searched. The optional third argument is an integer that represents the starting string position. By default, `LOCATE` starts at the beginning of the string. The starting position of a string is 1. If the string cannot be located, `LOCATE` returns 0.

The `SUBSTRING` function returns a string that is a substring of the first argument based on the starting position and length.

The `TRIM` function trims the specified character from the beginning and/or end of a string. If no character is specified, `TRIM` removes spaces or blanks from the string. If the optional `LEADING` specification is used, `TRIM` removes only the leading characters from the string. If the optional `TRAILING` specification is used, `TRIM` removes only the trailing characters from the string. The default is `BOTH`, which removes the leading and trailing characters from the string.

The `LOWER` and `UPPER` functions convert a string to lower or upper case, respectively.

TABLE 24-5 Arithmetic Expressions

Function Syntax	Return Type
ABS(number)	int, float, or double
MOD(int, int)	int
SQRT(double)	double
SIZE(Collection)	int

The ABS function takes a numeric expression and returns a number of the same type as the argument.

The MOD function returns the remainder of the first argument divided by the second.

The SQRT function returns the square root of a number.

The SIZE function returns an integer of the number of elements in the given collection.

Date/time functions return the date, time, or timestamp on the database server.

TABLE 24-6 Date/Time Expressions

Function Syntax	Return Type
CURRENT_DATE	java.sql.Date
CURRENT_TIME	java.sql.Time
CURRENT_TIMESTAMP	java.sql.Timestamp

## Case Expressions

Case expressions are expressions that change based on a condition, similar to the case keyword of the Java programming language. The CASE keyword indicates the start of a case expression, and the expression is terminated by the END keyword. The WHEN and THEN keyword define individual conditions, and the ELSE keyword defines the default condition should none of the other conditions be satisfied.

### EXAMPLE 24-1 Using Case Expressions in the Query Language

The following query selects the name of a person and a conditional string, depending on the subtype of the Person entity. If the subtype is Student, it returns the string kid. If the subtype is Guardian or Staff, it returns adult. If the entity is some other subtype of Person, the string unknown is returned.

```
SELECT p.name
CASE TYPE(p)
  WHEN Student THEN 'kid'
```

**EXAMPLE 24-1** Using Case Expressions in the Query Language *(Continued)*

```

    WHEN Guardian THEN 'adult'
    WHEN Staff THEN 'adult'
    ELSE 'unknown'
END
FROM Person p

```

The following query sets a discount for different types of customers. Gold-level customers get a 20% discount, silver-level customers get a 15% discount, bronze-level customers get a 10% discount, and everyone else gets a 5% discount.

```

UPDATE Customer c
SET c.discount =
    CASE c.level
        WHEN 'Gold' THEN 20
        WHEN 'SILVER' THEN 15
        WHEN 'Bronze' THEN 10
        ELSE 5
    END

```

## NULL Values

If the target of a reference is not in the persistent store, then the target is NULL. For conditional expressions containing NULL, the query language uses the semantics defined by SQL92. Briefly, these semantics are as follows:

- If a comparison or arithmetic operation has an unknown value, it yields a NULL value.
- Two NULL values are not equal. Comparing two NULL values yields an unknown value.
- The IS NULL test converts a NULL persistent field or a single-valued relationship field to TRUE. The IS NOT NULL test converts them to FALSE.
- Boolean operators and conditional tests use the three-valued logic defined by [Table 24-7](#) and [Table 24-8](#). (In these tables, T stands for TRUE, F for FALSE, and U for unknown.)

**TABLE 24-7** AND Operator Logic

AND	T	F	U
T	T	F	U
F	F	F	F
U	U	F	U

TABLE 24-8 OR Operator Logic

OR	T	F	U
T	T	T	T
F	T	F	U
U	T	U	U

## Equality Semantics

In the query language, only values of the same type can be compared. However, this rule has one exception: Exact and approximate numeric values can be compared. In such a comparison, the required type conversion adheres to the rules of Java numeric promotion.

The query language treats compared values as if they were Java types and not as if they represented types in the underlying data store. For example, if a persistent field could be either an integer or a NULL, then it must be designated as an Integer object and not as an int primitive. This designation is required because a Java object can be NULL but a primitive cannot.

Two strings are equal only if they contain the same sequence of characters. Trailing blanks are significant; for example, the strings 'abc' and 'abc ' are not equal.

Two entities of the same abstract schema type are equal only if their primary keys have the same value. Table 24-9 shows the operator logic of a negation, and Table 24-10 shows the truth values of conditional tests.

TABLE 24-9 NOT Operator Logic

NOT Value	Value
T	F
F	T
U	U

TABLE 24-10 Conditional Test

Conditional Test	T	F	U
Expression IS TRUE	T	F	F
Expression IS FALSE	F	T	F
Expression is unknown	F	F	T

## SELECT Clause

The SELECT clause defines the types of the objects or values returned by the query.

### Return Types

The return type of the SELECT clause is defined by the result types of the select expressions contained within it. If multiple expressions are used, the result of the query is an `Object[]`, and the elements in the array correspond to the order of the expressions in the SELECT clause, and in type to the result types of each expression.

A SELECT clause cannot specify a collection-valued expression. For example, the SELECT clause `p.teams` is invalid because `teams` is a collection. However, the clause in the following query is valid because the `t` is a single element of the `teams` collection:

```
SELECT t
FROM Player p, IN (p.teams) t
```

The following query is an example of a query with multiple expressions in the select clause:

```
SELECT c.name, c.country.name
FROM customer c
WHERE c.lastname = 'Coss' AND c.firstname = 'Roxane'
```

It returns a list of `Object[]` elements where the first array element is a string denoting the customer name and the second array element is a string denoting the name of the customer's country.

### Aggregate Functions in the SELECT Clause

The result of a query may be the result of an aggregate function, listed in [Table 24–11](#).

TABLE 24–11 Aggregate Functions in Select Statements

Name	Return Type	Description
AVG	Double	Returns the mean average of the fields.
COUNT	Long	Returns the total number of results.
MAX	the type of the field	Returns the highest value in the result set.
MIN	the type of the field	Returns the lowest value in the result set.

TABLE 24-11 Aggregate Functions in Select Statements (Continued)

Name	Return Type	Description
SUM	Long (for integral fields)Double (for floating point fields)BigInteger (for BigInteger fields)BigDecimal (for BigDecimal fields)	Returns the sum of all the values in the result set.

For select method queries with an aggregate function (AVG, COUNT, MAX, MIN, or SUM) in the SELECT clause, the following rules apply:

- For the AVG, MAX, MIN, and SUM functions, the functions return null if there are no values to which the function can be applied.
- For the COUNT function, if there are no values to which the function can be applied, COUNT returns 0.

The following example returns the average order quantity:

```
SELECT AVG(o.quantity)
FROM Order o
```

The following example returns the total cost of the items ordered by Roxane Coss:

```
SELECT SUM(l.price)
FROM Order o JOIN o.lineItems l JOIN o.customer c
WHERE c.lastname = 'Coss' AND c.firstname = 'Roxane'
```

The following example returns the total number of orders:

```
SELECT COUNT(o)
FROM Order o
```

The following example returns the total number of items in Hal Incandenza's order that have prices:

```
SELECT COUNT(l.price)
FROM Order o JOIN o.lineItems l JOIN o.customer c
WHERE c.lastname = 'Incandenza' AND c.firstname = 'Hal'
```

## The DISTINCT Keyword

The DISTINCT keyword eliminates duplicate return values. If a query returns a `java.util.Collection`, which allows duplicates, then you must specify the DISTINCT keyword to eliminate duplicates.

## Constructor Expressions

Constructor expressions allow you to return Java instances that store a query result element instead of an `Object[]`.

The following query creates a `CustomerDetail` instance per `Customer` matching the `WHERE` clause. A `CustomerDetail` stores the customer name and customer's country name. So the query returns a `List` of `CustomerDetail` instances:

```
SELECT NEW com.xyz.CustomerDetail(c.name, c.country.name)
FROM customer c
WHERE c.lastname = 'Coss' AND c.firstname = 'Roxane'
```

## ORDER BY Clause

As its name suggests, the `ORDER BY` clause orders the values or objects returned by the query.

If the `ORDER BY` clause contains multiple elements, the left-to-right sequence of the elements determines the high-to-low precedence.

The `ASC` keyword specifies ascending order (the default), and the `DESC` keyword indicates descending order.

When using the `ORDER BY` clause, the `SELECT` clause must return an orderable set of objects or values. You cannot order the values or objects for values or objects not returned by the `SELECT` clause. For example, the following query is valid because the `ORDER BY` clause uses the objects returned by the `SELECT` clause:

```
SELECT o
FROM Customer c JOIN c.orders o JOIN c.address a
WHERE a.state = 'CA'
ORDER BY o.quantity, o.totalcost
```

The following example is *not* valid because the `ORDER BY` clause uses a value not returned by the `SELECT` clause:

```
SELECT p.product_name
FROM Order o, IN(o.lineItems) l JOIN o.customer c
WHERE c.lastname = 'Faehmel' AND c.firstname = 'Robert'
ORDER BY o.quantity
```

## The GROUP BY Clause

The `GROUP BY` clause allows you to group values according to a set of properties.

The following query groups the customers by their country and returns the number of customers per country:

```
SELECT c.country, COUNT(c)
FROM Customer c GROUP BY c.country
```

## The HAVING Clause

The HAVING clause is used with the GROUP BY clause to further restrict the returned result of a query.

The following query groups orders by the status of their customer and returns the customer status plus the average totalPrice for all orders where the corresponding customers has the same status. In addition, it considers only customers with status 1, 2, or 3, so orders of other customers are not taken into account:

```
SELECT c.status, AVG(o.totalPrice)
  FROM Order o JOIN o.customer c
 GROUP BY c.status HAVING c.status IN (1, 2, 3)
```



## Creating Queries Using the Criteria API

---

The Criteria API is used to define queries for entities and their persistent state by creating query-defining objects. Criteria queries are written using Java programming language APIs, are type-safe, and are portable queries that work regardless of the underlying data store.

### Overview of the Criteria and Metamodel APIs

Similar to JPQL, the Criteria API is based on the abstract schema of persistent entities, their relationships, and embedded objects. The Criteria API operates on this abstract schema to allow developers to find, modify, and delete persistent entities by invoking Java Persistence API entity operations.

The Metamodel API works in concert with the Criteria API to model persistent entity classes for Criteria queries.

The Criteria API and JPQL are closely related, and designed to allow similar operations in their queries. Developers familiar with JPQL syntax will find equivalent object-level operations in the Criteria API.

#### EXAMPLE 25-1 A Simple Criteria Query

The following simple Criteria query returns all instances of the `Pet` entity in the data source.

```
EntityManager em = ...;
CriteriaBuilder cb = em.getCriteriaBuilder();
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.select(pet);
TypeQuery<Pet> q = em.createQuery(cq);
List<Pet> allPets = q.getResultList();
```

The equivalent JPQL query is:

**EXAMPLE 25-1** A Simple Criteria Query *(Continued)*

```
SELECT p
FROM Pet p
```

This query demonstrates the basic steps to create a Criteria query:

- Use an `EntityManager` instance to create a `CriteriaBuilder` object.
- Create a query object by creating an instance of the `CriteriaQuery` interface. This query object's attributes will be modified with the details of the query.
- Set the query root by calling the `from` method on the `CriteriaQuery` object.
- Specify what the type of the query result will be by calling the `select` method of the `CriteriaQuery` object.
- Prepare the query for execution by creating a `TypedQuery<T>` instance, specifying the type of the query result.
- Execute the query by calling the `getResultList` method on the `TypedQuery<T>` object. Because this query returns a collection of entities, the result is stored in a `List`.

The tasks associated with each step are discussed in detail in this chapter.

To create a `CriteriaBuilder` instance call the `getCriteriaBuilder` method on the `EntityManager` instance:

```
CriteriaBuilder cb = em.getCriteriaBuilder();
```

The actual query object is created using the `CriteriaBuilder` instance:

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
```

The query will return instances of the `Pet` entity, so the type of the query is specified when the `CriteriaQuery` object is created to create a type-safe query.

The `from` clause of the query is set, and the root of the query specified, by calling the `from` method of the query object:

```
Root<Pet> pet = cq.from(Pet.class);
```

The `select` clause of the query is set by calling the `select` method of the query object, and passing in the query root:

```
cq.select(pet);
```

The query object is now used to create a `TypedQuery<T>` object that can be executed against the data source. The modifications to the query object are captured to create a ready-to-execute query.

```
TypeQuery<Pet> q = em.createQuery(cq);
```

**EXAMPLE 25-1** A Simple Criteria Query (Continued)

This typed query object is executed by calling its `getResultList` method, because this query will return multiple entity instances. The results are stored in a `List<Pet>` collection-valued object.

```
List<Pet> allPets = q.getResultList();
```

## Modeling Entity Classes with the Metamodel API

The Metamodel API is used to create a metamodel of the managed entities in a particular persistence unit. For each entity class in a particular package, a metamodel class is created with a trailing underscore, and with attributes that correspond to the persistent fields or properties of the entity class.

**EXAMPLE 25-2** Example Entity Class and Corresponding Metamodel Class

The following entity class `com.example.Pet` has four persistent fields, `id`, `name`, `color`, and `owners`.

```
package com.example;

...

@Entity
public class Pet {
    @Id
    protected Long id;
    protected String name;
    protected String color;
    @ManyToOne
    protected Set<Person> owners;
    ...
}
```

The corresponding Metamodel class is:

```
package com.example;

...

@Static Metamodel(Pet.class)
public class Pet_ {

    public static volatile SingularAttribute<Pet, Long> id;
    public static volatile SingularAttribute<Pet, String> name;
```

**EXAMPLE 25-2** Example Entity Class and Corresponding Metamodel Class *(Continued)*

```
public static volatile SingularAttribute<Pet, String> color;
public static volatile SetAttribute<Pet, Person> owners;
}
```

The metamodel class and its attributes are used in Criteria queries to refer to the managed entity classes and their persistent state and relationships.

## Using Metamodel Classes

Metamodel classes that correspond to entity classes are of type `javax.persistence.metamodel.EntityType<T>`, and are typically generated by annotation processors either at development time or at runtime. Developers of applications that use Criteria queries may generate static metamodel classes using the persistence provider's annotation processor, or may obtain the metamodel class either by calling the `getModel` method on the query root object or first obtaining an instance of the `Metamodel` interface and then passing the entity type to the instance's `entity` method.

**EXAMPLE 25-3** Obtaining a Metamodel Class Dynamically Using the `Root<T>.getModel` Method

The following code snippet shows how to obtain the `Pet` entity's metamodel class by calling `Root<T>.getModel`.

```
EntityManager em = ...;
CriteriaBuilder cb = em.getCriteriaBuilder();
CriteriaQuery cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
EntityType<Pet> Pet_ = pet.getModel();
```

**EXAMPLE 25-4** Obtaining a Metamodel Class Dynamically Using the `Metamodel.getModel` Method

The following code snippet shows how to obtain the `Pet` entity's metamodel class by first obtaining a metamodel instance using `EntityManager.getModel`, and then calling `entity` on the metamodel instance.

```
EntityManager em = ...;
Metamodel m = em.getModel();
EntityType<Pet> Pet_ = m.entity(Pet.class);
```

# Basic Type-Safe Queries Using the Criteria API and Metamodel API

The basic semantics of a Criteria query consists of a select clause, a from clause, and an optional where clause, similar to a JPQL query. Criteria queries set these clauses using Java programming language objects, so the query can be created in a type-safe manner.

## Creating a Criteria Query

The `javax.persistence.criteria.CriteriaBuilder` interface is used to construct:

- Criteria queries
- selections
- expressions
- predicates
- ordering

To obtain an instance of the `CriteriaBuilder` interface, call the `getCriteriaBuilder` method on either an `EntityManager` or `EntityManagerFactory` instance.

**EXAMPLE 25-5** Obtaining a `CriteriaBuilder` Instance Using the `EntityManager.getCriteriaManager` Method

The following code shows how to obtain a `CriteriaBuilder` instance using the `EntityManager.getCriteriaBuilder` method.

```
EntityManager em = ...;  
CriteriaBuilder cb = em.getCriteriaBuilder();
```

Criteria queries are constructed by obtaining an instance of the `javax.persistence.criteria.CriteriaQuery` interface. `CriteriaQuery` objects define a particular query that will navigate over one or more entities. Obtain `CriteriaQuery` instances by calling one of the `CriteriaBuilder.createQuery` methods. For creating type-safe queries, call the `CriteriaBuilder.createQuery` method as follows:

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
```

The `CriteriaQuery` object's type should be set to the expected result type of the query. In the code above, the object's type is set to `CriteriaQuery<Pet>` for a query which will find instances of the `Pet` entity.

**EXAMPLE 25-6** Creating a CriteriaQuery Instance for a Query that Returns a String

In the following code snippet, a `CriteriaQuery` object is created for a query that returns a string.

```
CriteriaQuery<String> cq = cb.createQuery(String.class);
```

## Query Roots

For a particular `CriteriaQuery` object, the root entity of the query, from which all navigation originates, is called the *query root*. It is similar to the `FROM` clause in a JPQL query.

Create the query root by calling the `from` method on the `CriteriaQuery` instance. The argument to the `from` method is either the entity class, or an `EntityType<T>` instance for the entity.

**EXAMPLE 25-7** Setting the Query Root Using the Entity Class

The following code sets the query root to the `Pet` entity.

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
Root<Pet> pet = cq.from(Pet.class);
```

**EXAMPLE 25-8** Setting the Query Root Using a Metamodel `EntityType<T>` Instance

The following code sets the query root to the `Pet` class using an `EntityType<T>` instance.

```
EntityManager em = ...;  
Metamodel m = em.getMetamodel();  
EntityType<Pet> Pet_ = m.entity(Pet.class);  
Root<Pet> pet = cq.from(Pet_);
```

Criteria queries may have more than one query root. This usually occurs when the query navigates from several entities.

**EXAMPLE 25-9** Creating a Query with Multiple Query Roots

The following code has two `Root` instances.

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
Root<Pet> pet1 = cq.from(Pet.class);  
Root<Pet> pet2 = cq.from(Pet.class);
```

## Querying Relationships Using Joins

For queries that navigate to related entity classes, the query must define a join to the related entity by calling one of the `From.join` methods on the query root object, or another join object. The `join` methods are similar to the `JOIN` keyword in JPQL.

The target of the join uses the Metamodel class of type `EntityType<T>` to specify the persistent field or property of the joined entity.

The `join` methods return an object of type `Join<X, Y>`, where `X` is the source entity and `Y` is the target of the join.

### EXAMPLE 25-10 Joining a Query

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);

Root<Pet> pet = cq.from(Pet.class);
Join<Pet, Owner> owner = pet.join(Pet_.owners);
```

Joins can be chained together to navigate to related entities of the target entity without having to create a `Join<X, Y>` instance for each join.

### EXAMPLE 25-11 Chaining Joins Together in a Query

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);
EntityType<Owner> Owner_ = m.entity(Owner.class);

Root<Pet> pet = cq.from(Pet.class);
Join<Owner, Address> address = cq.join(Pet_.owners).join(Owner_.addresses);
```

## Path Navigation in Criteria Queries

Path objects are used in the `select` and `where` clauses of a Criteria query, and can be query root entities, join entities, or other Path objects. The `Path.get` method is used to navigate to attributes of the entities of a query.

The argument to the `get` method is the corresponding attribute of the entity's Metamodel class. The attribute can either be a single-valued attribute (specified by `@SingularAttribute` in the Metamodel class) or a collection-valued attribute (specified by one of `@CollectionAttribute`, `@SetAttribute`, `@ListAttribute`, or `@MapAttribute`).

**EXAMPLE 25-12** Using Path Objects in the Select Clause of a Query

```
CriteriaQuery<String> cq = cb.createQuery(String.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);

Root<Pet> pet = cq.from(Pet.class);
cq.select(pet.get(Pet_.name));
```

This query returns the names of all the pets in the data store. The `get` method is called on the query root, `pet`, with the name attribute of the `Pet` entity's Metamodel class, `Pet_` as the argument.

## Restricting Criteria Query Results

The results of a query can be restricted on the `CriteriaQuery` object according to conditions set by calling the `CriteriaQuery.where` method. Calling the `where` method is analogous to setting the `WHERE` clause in a JPQL query.

The `where` method evaluates instances of the `Expression` interface to restrict the results according to the conditions of the expressions. `Expression` instances are created using methods defined in the `Expression` and `CriteriaBuilder` interfaces.

### The Expression Interface Methods

An `Expression` object is used to in a query's `select`, `where`, or `having` clause

**TABLE 25-1** Conditional Methods in the Expression Interface

Method	Description
<code>isNull</code>	Tests whether an expression is null.
<code>isNotNull</code>	Tests whether an expression is not null.
<code>in</code>	Tests whether an expression is within a list of values.

**EXAMPLE 25-13** Using the `Expression.isNull` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.where(pet.get(Pet_.color).isNull());
```

This query finds all pets where the `color` attribute is null.

**EXAMPLE 25-14** Using the `Expression.in` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.where(pet.get(Pet_.color).in("brown", "black");
```

This query finds all brown and black pets.

The `in` method also can check whether an attribute is a member of a collection.

## Expression Methods in the CriteriaBuilder Interface

The `CriteriaBuilder` interface defines additional methods for creating expressions. These methods correspond to the arithmetic, string, date, time, and case operators and functions of JPQL.

**TABLE 25-2** Conditional Methods in the `CriteriaBuilder` Interface

Conditional Method	Description
<code>equal</code>	Tests whether two expressions are equal.
<code>notEqual</code>	Tests whether two expressions are not equal.
<code>gt</code>	Tests whether the first numeric expression is greater than the second numeric expression.
<code>ge</code>	Tests whether the first numeric expression is greater than or equal to the second numeric expression.
<code>lt</code>	Tests whether the first numeric expression is less than the second numeric expression.
<code>le</code>	Tests whether the first numeric expression is less than or equal to the second numeric expression.
<code>between</code>	Tests whether the first expression is between the second and third expression in value.
<code>like</code>	Tests whether the expression matches a given pattern.

**EXAMPLE 25-15** Using the `CriteriaBuilder.equal` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
```

**EXAMPLE 25-15** Using the `CriteriaBuilder.equal` Method *(Continued)*

```
cq.where(cb.equal(pet.get(Pet_.name)), "Fido");  
...
```

**EXAMPLE 25-16** Using the `CriteriaBuilder.gt` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
Metamodel m = em.getMetamodel();  
EntityType<Pet> Pet_ = m.entity(Pet.class);  
Root<Pet> pet = cq.from(Pet.class);  
Date someDate = new Date(...);  
cq.where(cb.gt(pet.get(Pet_.birthday)), date);
```

**EXAMPLE 25-17** Using the `CriteriaBuilder.between` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
Metamodel m = em.getMetamodel();  
EntityType<Pet> Pet_ = m.entity(Pet.class);  
Root<Pet> pet = cq.from(Pet.class);  
Date firstDate = new Date(...);  
Date secondDate = new Date(...);  
cq.where(cb.between(pet.get(Pet_.birthday)), firstDate, secondDate);
```

**EXAMPLE 25-18** Using the `CriteriaBuilder.like` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
Metamodel m = em.getMetamodel();  
EntityType<Pet> Pet_ = m.entity(Pet.class);  
Root<Pet> pet = cq.from(Pet.class);  
cq.where(cb.like(pet.get(Pet_.name)), "*do");
```

Multiple conditional predicates can be specified by using the compound predicate methods of the `CriteriaBuilder` interface.

**TABLE 25-3** Compound Predicate Methods in the `CriteriaBuilder` Interface

Method	Description
<code>and</code>	A logical conjunction of two boolean expressions.
<code>or</code>	A logical disjunction of two boolean expressions.
<code>not</code>	A logical negation of the given boolean expression.

**EXAMPLE 25-19** Using Compound Predicates in Queries

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Metamodel m = em.getMetamodel();
EntityType<Pet> Pet_ = m.entity(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.where(cb.equal(pet.get(Pet_.name), "Fido")
        .and(cb.equal(pet.get(Pet_.color), "brown"));
```

## Managing Criteria Query Results

For queries that return more than one result, it's often helpful to organize those results. The `CriteriaQuery` interface defines the `orderBy` method to order query results according to attributes of an entity. The `CriteriaQuery` interface also defines the `groupBy` method to group the results of a query together according to attributes of an entity, and the `and having` method to restrict those groups according to a condition.

### Ordering Results

The order of the results of a query can be set by calling the `CriteriaQuery.orderBy` method and passing in an `Order` object. `Order` objects are created by calling either the `CriteriaBuilder.asc` or `CriteriaBuilder.desc` methods. The `asc` method is used to order the results by ascending value of the passed expression parameter. The `desc` method is used to order the results by descending value of the passed expression parameter.

**EXAMPLE 25-20** Ordering Results in Descending Order

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.select(pet);
cq.orderBy(cb.desc(pet.get(Pet_.birthday)));
```

In this query, the results will be ordered by the pet's birthday from highest to lowest. That is, pets born in December will appear before pets born in May.

**EXAMPLE 25-21** Ordering Results in Ascending Order

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
Join<Owner, Address> address = cq.join(Pet_.owners).join(Owner_.address);
cq.select(pet);
cq.orderBy(cb.asc(address.get(Address_.postalCode)));
```

In this query, the results will be ordered by the pet owner's postal code from lowest to highest. That is, pets whose owner lives in the 10001 zip code will appear before pets whose owner lives in the 91000 zip code.

If more than one `Order` object is passed to `orderBy`, the precedence is determined by the order in which they appear in the argument list of `orderBy`. The first `Order` object has precedence.

**EXAMPLE 25-22** Ordering Results By More than One Criteria

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
Join<Pet, Owner> owner = cq.join(Pet_.owners);
cq.select(pet);
cq.orderBy(cb.asc(owner.get(Owner_.lastName), owner.get(Owner_.firstName)));
```

The results of this query will be ordered alphabetically by the pet owner's last name, then first name.

## Grouping Results

The `CriteriaQuery.groupBy` method partitions the query results into groups. These groups are set by passing an expression to `groupBy`.

**EXAMPLE 25-23** Grouping Results Using the `CriteriaQuery.groupBy` Method

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.groupBy(pet.get(Pet_.color));
```

This query returns all `Pet` entities, and groups the results by pet's color.

The `CriteriaQuery.having` method is used in conjunction with `groupBy` to filter over the groups. The `having` method takes a conditional expression as a parameter. By calling the `having` method, the query result is restricted according to the conditional expression.

**EXAMPLE 25-24** Grouping Results Using the `CriteriaQuery.groupBy` and `CriteriaQuery.having` Methods

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);
Root<Pet> pet = cq.from(Pet.class);
cq.groupBy(pet.get(Pet_.color));
cq.having(cb.in(pet.get(Pet_.color)).value("brown").value("blonde"));
```

In this example, the query groups the returned `Pet` entities by color, as in the example above. However, the only returned groups will be `Pet` entities where the `color` attribute is set to `brown` or `blonde`. That is, no grey colored pets will be returned in this query.

## Executing Queries

To prepare a query for execution, create a `TypedQuery<T>` object with the type of the query result by passing the `CriteriaQuery` object to `EntityManager.createQuery`.

Queries are executed by calling either `getSingleResult` or `getResultList` on the `TypedQuery<T>` object.

## Single-Valued Query Results

The `TypedQuery<T>.getSingleResult` method is used for executing queries that return a single result.

**EXAMPLE 25-25** Retrieving Single-Valued Query Result

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
...  
TypedQuery<Pet> q = em.createQuery(cq);  
Pet result = q.getSingleResult();
```

## Collection-Valued Query Results

The `TypedQuery<T>.getResultList` method is used for executing queries that return a collection of objects.

**EXAMPLE 25-26** Retrieving Collection-Valued Query Results

```
CriteriaQuery<Pet> cq = cb.createQuery(Pet.class);  
...  
TypedQuery<Pet> q = em.createQuery(cq);  
List<Pet> results = q.getResultList();
```



( PART VII  
Security

Part Seven introduces basic security concepts and examples.



# Introduction to Security in the Java EE Platform

---

This and subsequent chapters discuss how to address security requirements in Java EE, web, and web services applications. Every enterprise that has sensitive resources that can be accessed by many users, or resources that traverse unprotected, open, networks, such as the Internet, needs to be protected.

This chapter introduces basic security concepts and security implementation mechanisms. More information on these concepts and mechanisms can be found in the *Security* chapter of the Java EE 6 specification. This document is available for download online at <http://www.jcp.org/en/jsr/detail?id=316>.

Other chapters in this tutorial that address security requirements include the following:

- Chapter 27, “Getting Started Securing Enterprise Applications,” discusses adding security to Java EE components such as enterprise beans and application clients.
- Chapter 28, “Getting Started Securing Web Applications,” discusses and provides examples for adding security to web components such as servlets, and to web services such as JAX-WS.

Some of the material in this chapter assumes that you understand basic security concepts. To learn more about these concepts, you should explore the Java SE security web site before you begin this chapter. The URL for this site is <http://java.sun.com/javase/6/docs/technotes/guides/security/>.

This tutorial assumes deployment onto the Sun GlassFishEnterprise Server v3 and provides some information regarding configuration of the Enterprise Server. The best source for information regarding configuration of the Enterprise Server, however, is the *Sun GlassFish Enterprise Server v3 Administration Guide*. The best source for development tips specific to the Enterprise Server is the *Sun GlassFish Enterprise Server v3 Application Development Guide*. The best source for tips on deploying applications to the Enterprise Server is the *Sun GlassFish Enterprise Server v3 Application Deployment Guide*.

# Overview of Java EE Security

Java EE, web, and web services applications are made up of components that can be deployed into different containers. These components are used to build a multi-tier enterprise application. Security for components is provided by their containers. A container provides two kinds of security: declarative and programmatic security.

- *Declarative security* expresses an application component's security requirements using *deployment descriptors*. Deployment descriptors are external to an application, and include information that specifies how security roles and access requirements are mapped into environment-specific security roles, users, and policies. For this volume of the tutorial, deployment descriptors are used to secure web applications. For more information about deployment descriptors, read [“Using Deployment Descriptors for Declarative Security” on page 482](#).

*Annotations* (also called *metadata*) are used to specify information about security within a class file. When the application is deployed, this information can either be used by or overridden by the application deployment descriptor. Annotations save you from having to write declarative information inside XML descriptors. Instead, you just put annotations on the code and the required information gets generated. For this tutorial, annotations are used for securing enterprise applications. For more information about annotations, read [“Using Annotations” on page 483](#).

- *Programmatic security* is embedded in an application and is used to make security decisions. Programmatic security is useful when declarative security alone is not sufficient to express the security model of an application. For more information about programmatic security, read [“Using Programmatic Security” on page 484](#).

## A Simple Security Example

The security behavior of a Java EE environment may be better understood by examining what happens in a simple application with a web client, a JSP user interface, and enterprise bean business logic.

In the following example, which is taken from JSR-316, the [Java EE 6 Specification](#), the web client relies on the web server to act as its authentication proxy by collecting user authentication data from the client and using it to establish an authenticated session.

### Step 1: Initial Request

In the first step of this example, the web client requests the main application URL. This action is shown in [Figure 26–1](#).

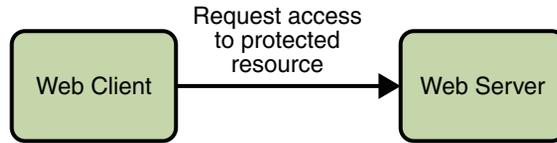


FIGURE 26-1 Initial Request

Since the client has not yet authenticated itself to the application environment, the server responsible for delivering the web portion of the application (hereafter referred to as *web server*) detects this and invokes the appropriate authentication mechanism for this resource. For more information on these mechanisms, read [“Security Implementation Mechanisms”](#) on page 479.

## Step 2: Initial Authentication

The web server returns a form that the web client uses to collect authentication data (for example, user name and password) from the user. The web client forwards the authentication data to the web server, where it is validated by the web server, as shown in [Figure 26-2](#).

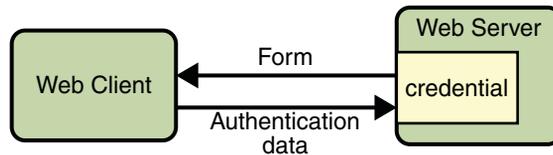


FIGURE 26-2 Initial Authentication

The validation mechanism may be local to a server, or it may leverage the underlying security services. On the basis of the validation, the web server sets a credential for the user.

## Step 3: URL Authorization

The credential is used for future determinations of whether the user is authorized to access restricted resources it may request. The web server consults the security policy (derived from the deployment descriptor) associated with the web resource to determine the security roles that are permitted access to the resource. The web container then tests the user’s credential against each role to determine if it can map the user to the role. [Figure 26-3](#) shows this process.

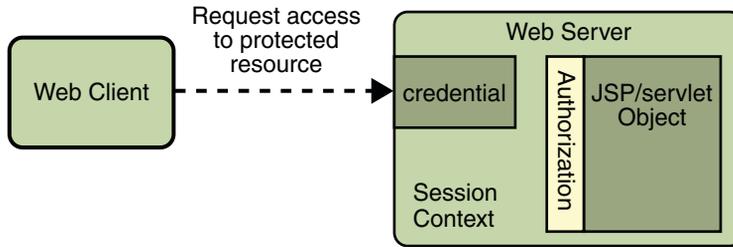


FIGURE 26-3 URL Authorization

The web server’s evaluation stops with an “is authorized” outcome when the web server is able to map the user to a role. A “not authorized” outcome is reached if the web server is unable to map the user to any of the permitted roles.

### Step 4: Fulfilling the Original Request

If the user is authorized, the web server returns the result of the original URL request, as shown in Figure 26-4.

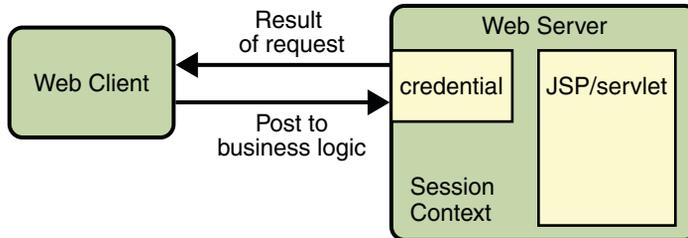


FIGURE 26-4 Fulfilling the Original Request

In our example, the response URL of a JSP page is returned, enabling the user to post form data that needs to be handled by the business logic component of the application. Read [Chapter 28, “Getting Started Securing Web Applications,”](#) for more information on protecting web applications.

### Step 5: Invoking Enterprise Bean Business Methods

The JSP page performs the remote method call to the enterprise bean, using the user’s credential to establish a secure association between the JSP page and the enterprise bean (as shown in Figure 26-5). The association is implemented as two related security contexts, one in the web server and one in the EJB container.

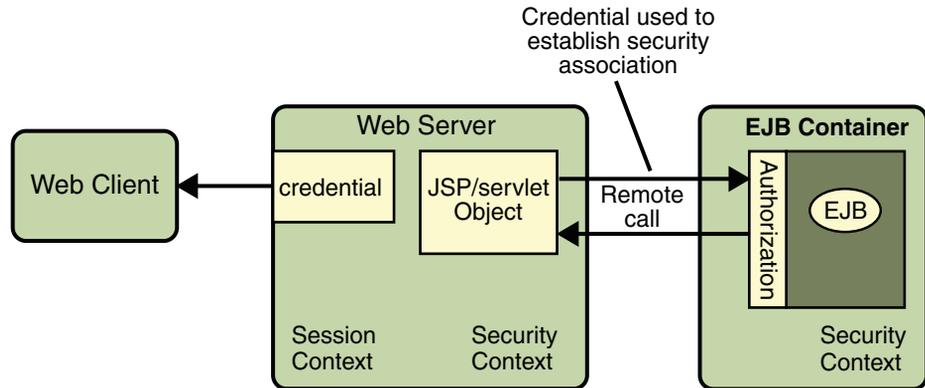


FIGURE 26-5 Invoking an Enterprise Bean Business Method

The EJB container is responsible for enforcing access control on the enterprise bean method. It consults the security policy (derived from the deployment descriptor) associated with the enterprise bean to determine the security roles that are permitted access to the method. For each role, the EJB container uses the security context associated with the call to determine if it can map the caller to the role.

The container's evaluation stops with an "is authorized" outcome when the container is able to map the caller's credential to a role. A "not authorized" outcome is reached if the container is unable to map the caller to any of the permitted roles. A "not authorized" result causes an exception to be thrown by the container, and propagated back to the calling JSP page.

If the call is authorized, the container dispatches control to the enterprise bean method. The result of the bean's execution of the call is returned to the JSP, and ultimately to the user by the web server and the web client.

## Security Functions

A properly implemented security mechanism will provide the following functionality:

- Prevent unauthorized access to application functions and business or personal data (authentication)
- Hold system users accountable for operations they perform (non-repudiation)
- Protect a system from service interruptions and other breaches that affect quality of service (QoS)

Ideally, properly implemented security mechanisms will also provide the following functionality:

- Easy to administer
- Transparent to system users

- Interoperable across application and enterprise boundaries

## Characteristics of Application Security

Java EE applications consist of components that can contain both protected and unprotected resources. Often, you need to protect resources to ensure that only authorized users have access. *Authorization* provides controlled access to protected resources. Authorization is based on identification and authentication. *Identification* is a process that enables recognition of an entity by a system, and *authentication* is a process that verifies the identity of a user, device, or other entity in a computer system, usually as a prerequisite to allowing access to resources in a system.

Authorization and authentication are not required for an entity to access unprotected resources. Accessing a resource without authentication is referred to as unauthenticated or anonymous access.

The characteristics of application security that, when properly addressed, help to minimize the security threats faced by an enterprise, include the following:

- **Authentication:** The means by which communicating entities (for example, client and server) prove to one another that they are acting on behalf of specific identities that are authorized for access. This ensures that users are who they say they are.
- **Authorization, or Access Control:** The means by which interactions with resources are limited to collections of users or programs for the purpose of enforcing integrity, confidentiality, or availability constraints. This ensures that users have permission to perform operations or access data.
- **Data integrity:** The means used to prove that information has not been modified by a third party (some entity other than the source of the information). For example, a recipient of data sent over an open network must be able to detect and discard messages that were modified after they were sent. This ensures that only authorized users can modify data.
- **Confidentiality or Data Privacy:** The means used to ensure that information is made available only to users who are authorized to access it. This ensures that only authorized users can view sensitive data.
- **Non-repudiation:** The means used to prove that a user performed some action such that the user cannot reasonably deny having done so. This ensures that transactions can be proven to have happened.
- **Quality of Service (QoS):** The means used to provide better service to selected network traffic over various technologies.
- **Auditing:** The means used to capture a tamper-resistant record of security-related events for the purpose of being able to evaluate the effectiveness of security policies and mechanisms. To enable this, the system maintains a record of transactions and security information.

# Security Implementation Mechanisms

The characteristics of an application should be considered when deciding the layer and type of security to be provided for applications. The following sections discuss the characteristics of the common mechanisms that can be used to secure Java EE applications. Each of these mechanisms can be used individually or with others to provide protection layers based on the specific needs of your implementation.

## Java SE Security Implementation Mechanisms

Java SE provides support for a variety of security features and mechanisms, including:

- **Java Authentication and Authorization Service (JAAS):** JAAS is a set of APIs that enable services to authenticate and enforce access controls upon users. JAAS provides a pluggable and extensible framework for programmatic user authentication and authorization. JAAS is a core Java SE API and is an underlying technology for Java EE security mechanisms.
- **Java Generic Security Services (Java GSS-API):** Java GSS-API is a token-based API used to securely exchange messages between communicating applications. The GSS-API offers application programmers uniform access to security services atop a variety of underlying security mechanisms, including Kerberos.
- **Java Cryptography Extension (JCE):** JCE provides a framework and implementations for encryption, key generation and key agreement, and Message Authentication Code (MAC) algorithms. Support for encryption includes symmetric, asymmetric, block, and stream ciphers. Block ciphers operate on groups of bytes while stream ciphers operate on one byte at a time. The software also supports secure streams and sealed objects.
- **Java Secure Sockets Extension (JSSE):** JSSE provides a framework and an implementation for a Java version of the SSL and TLS protocols and includes functionality for data encryption, server authentication, message integrity, and optional client authentication to enable secure Internet communications.
- **Simple Authentication and Security Layer (SASL):** SASL is an Internet standard (RFC 2222) that specifies a protocol for authentication and optional establishment of a security layer between client and server applications. SASL defines how authentication data is to be exchanged but does not itself specify the contents of that data. It is a framework into which specific authentication mechanisms that specify the contents and semantics of the authentication data can fit.

Java SE also provides a set of tools for managing keystores, certificates, and policy files; generating and verifying JAR signatures; and obtaining, listing, and managing Kerberos tickets.

For more information on Java SE security, visit its web page at <http://java.sun.com/javase/6/docs/technotes/guides/security/>.

## Java EE Security Implementation Mechanisms

Java EE security services are provided by the component container and can be implemented using declarative or programmatic techniques (container security is discussed more in [“Securing Containers” on page 482](#)). Java EE security services provide a robust and easily configured security mechanism for authenticating users and authorizing access to application functions and associated data at many different layers. Java EE security services are separate from the security mechanisms of the operating system.

### Application-Layer Security

In Java EE, component containers are responsible for providing application-layer security. Application-layer security provides security services for a specific application type tailored to the needs of the application. At the application layer, application firewalls can be employed to enhance application protection by protecting the communication stream and all associated application resources from attacks.

Java EE security is easy to implement and configure, and can offer fine-grained access control to application functions and data. However, as is inherent to security applied at the application layer, security properties are not transferable to applications running in other environments and only protect data while it is residing in the application environment. In the context of a traditional application, this is not necessarily a problem, but when applied to a web services application, where data often travels across several intermediaries, you would need to use the Java EE security mechanisms along with transport-layer security and message-layer security for a complete security solution.

The advantages of using application-layer security include the following:

- Security is uniquely suited to the needs of the application.
- Security is fine-grained, with application-specific settings.

The disadvantages of using application-layer security include the following:

- The application is dependent on security attributes that are not transferable between application types.
- Support for multiple protocols makes this type of security vulnerable.
- Data is close to or contained within the point of vulnerability.

For more information on providing security at the application layer, read [“Securing Containers” on page 482](#).

### Transport-Layer Security

Transport-layer security is provided by the transport mechanisms used to transmit information over the wire between clients and providers, thus transport-layer security relies on secure HTTP transport (HTTPS) using Secure Sockets Layer (SSL). Transport security is a

point-to-point security mechanism that can be used for authentication, message integrity, and confidentiality. When running over an SSL-protected session, the server and client can authenticate one another and negotiate an encryption algorithm and cryptographic keys before the application protocol transmits or receives its first byte of data. Security is “live” from the time it leaves the consumer until it arrives at the provider, or vice versa, even across intermediaries. The problem is that it is not protected once it gets to its destination. One solution is to encrypt the message before sending.

Transport-layer security is performed in a series of phases, which are listed here:

- The client and server agree on an appropriate algorithm.
- A key is exchanged using public-key encryption and certificate-based authentication.
- A symmetric cipher is used during the information exchange.

Digital certificates are necessary when running secure HTTP transport (HTTPS) using Secure Sockets Layer (SSL). The HTTPS service of most web servers will not run unless a digital certificate has been installed. Digital certificates have already been created for the Enterprise Server. If you are using a different server, check to see if it has digital certificates, and, if not, use the procedure outlined in [“Working with Digital Certificates” on page 497](#) to set up a digital certificate that can be used by your web or application server to enable SSL.

The advantages of using transport-layer security include the following:

- Relatively simple, well understood, standard technology.
- Applies to message body and attachments.

The disadvantages of using transport-layer security include the following:

- Tightly-coupled with transport-layer protocol.
- All or nothing approach to security. This implies that the security mechanism is unaware of message contents, and as such, you cannot selectively apply security to portions of the message as you can with message-layer security.
- Protection is transient. The message is only protected while in transit. Protection is removed automatically by the endpoint when it receives the message.
- Not an end-to-end solution, simply point-to-point.

For more information on transport-layer security, read [“Establishing a Secure Connection Using SSL” on page 494](#).

## Message-Layer Security

In message-layer security, security information is contained within the SOAP message and/or SOAP message attachment, which allows security information to travel along with the message or attachment. For example, a portion of the message may be signed by a sender and encrypted for a particular receiver. When the message is sent from the initial sender, it may pass through intermediate nodes before reaching its intended receiver. In this scenario, the encrypted

portions continue to be opaque to any intermediate nodes and can only be decrypted by the intended receiver. For this reason, message-layer security is also sometimes referred to as *end-to-end security*.

The advantages of message-layer security include the following:

- Security stays with the message over all hops and after the message arrives at its destination.
- Security can be selectively applied to different portions of a message (and to attachments if using XWSS).
- Message security can be used with intermediaries over multiple hops.
- Message security is independent of the application environment or transport protocol.

The disadvantage of using message-layer security is that it is relatively complex and adds some overhead to processing.

The Enterprise Server supports message security using Metro, a web services stack that uses Web Services Security (WSS) to secure messages. Because this message security is specific to Metro and is not a part of the Java EE platform, this tutorial does not discuss using WSS to secure messages. See the *Metro User's Guide* at <https://metro.dev.java.net/guide/>. Other sources of information for message security include *Sun GlassFish Enterprise Server v3 Administration Guide* and *Sun GlassFish Enterprise Server v3 Application Development Guide*.

## Securing Containers

In Java EE, the component containers are responsible for providing application security. A container provides two types of security: declarative and programmatic. The following sections discuss these concepts in more detail.

- “Using Deployment Descriptors for Declarative Security” on page 482
- “Using Annotations” on page 483
- “Using Programmatic Security” on page 484

## Using Deployment Descriptors for Declarative Security

Declarative security expresses an application component's security requirements using *deployment descriptors*. A deployment descriptor is an XML document with an `.xml` extension that describes the deployment settings of an application, a module, or a component. Because deployment descriptor information is declarative, it can be changed without the need to modify the source code. At runtime, the Java EE server reads the deployment descriptor and acts upon the application, module, or component accordingly.

This tutorial does not document how to write the deployment descriptors from scratch, only what configurations each example requires its deployment descriptors to define. For help with writing deployment descriptors, you can view the provided deployment descriptors in a text editor. Another way to learn how to write deployment descriptors is to read the specification in which the deployment descriptor elements are defined.

Deployment descriptors must provide certain structural information for each component if this information has not been provided in annotations or is not to be defaulted.

Different types of components use different formats, or *schema*, for their deployment descriptors. The security elements of deployment descriptors which are discussed in this tutorial include the following:

- Enterprise JavaBeans components may use an EJB deployment descriptor named `META-INF/ejb-jar.xml` and would be contained in the EJB JAR file.

The schema for enterprise bean deployment descriptors is provided in the EJB 3.1 Specification (JSR-318), Chapter 9, *Deployment Descriptor*, which can be downloaded from <http://jcp.org/en/jsr/detail?id=318>.

- Deployment descriptor elements for web services components are defined in JSR-109. This deployment descriptor provides deployment-time mapping functionality between Java and WSDL. In conjunction with JSR-181, JAX-WS 2.2 complements this mapping functionality with development-time Java annotations that control mapping between Java and WSDL.

The schema for web services deployment descriptors is provided in *Web Services for Java EE* (JSR-109), section 7.1, *Web Services Deployment Descriptor XML Schema*, which can be downloaded from <http://jcp.org/en/jsr/detail?id=109>.

Schema elements for web application deployment descriptors are discussed in this tutorial in the section “[Introduction to Web Application Deployment Descriptors](#)” on page 530.

- Web components use a web application deployment descriptor named `web.xml`.

The schema for web component deployment descriptors is provided in the Java Servlet 3.0 Specification (JSR-315), chapter 14, *Deployment Descriptor*, which can be downloaded from <http://jcp.org/en/jsr/detail?id=315>.

Security elements for web application deployment descriptors are discussed in this tutorial in the section “[Introduction to Web Application Deployment Descriptors](#)” on page 530.

## Using Annotations

*Annotations* enable a declarative style of programming, and so encompass both the declarative and programmatic security concepts. Users can specify information about security within a class file using annotations. When the application is deployed, this information is used by the Enterprise Server. Not all security information can be specified using annotations, however. Some information must be specified in the application deployment descriptors.

Annotations let you avoid writing boilerplate code under many circumstances by enabling tools to generate it from annotations in the source code. This leads to a declarative programming style, where the programmer says what should be done and tools emit the code to do it. It also eliminates the need for maintaining side files that must be kept up to date with changes in source files. Instead the information can be maintained in the source file.

In this tutorial, specific annotations that can be used to specify security information within a class file are described in [“Securing an Enterprise Bean Using Declarative Security and Annotations” on page 507](#). In this version of the tutorial, annotations are not shown for securing web applications. This is because the use of annotations vary between web components, and deployment descriptors are the common method of securing web applications.

The following are sources for more information on annotations:

- *JSR 175: A Metadata Facility for the Java Programming Language*
- *JSR 181: Web Services Metadata for the Java Platform*
- *JSR 250: Common Annotations for the Java Platform*
- The Java SE discussion of annotations

Links to this information are provided in [“Further Information about Security” on page 500](#).

## Using Programmatic Security

Programmatic security is embedded in an application and is used to make security decisions. Programmatic security is useful when declarative security alone is not sufficient to express the security model of an application. The API for programmatic security consists of two methods of the `EJBContext` interface and six methods of the `HttpServletRequest` interface. These methods allow components to make business logic decisions based on the security role of the caller or remote user.

Programmatic security is discussed in more detail in the following sections:

- [“Securing an Enterprise Bean Programmatically” on page 515](#)
- [“Using Programmatic Security with Web Applications” on page 548](#)

## Securing the Enterprise Server

This tutorial describes deployment to the Sun GlassFish Enterprise Server v3, which provides highly secure, interoperable, and distributed component computing based on the Java EE security model. The Enterprise Server supports the Java EE 6 security model. You can configure the Enterprise Server for the following purposes:

- Adding, deleting, or modifying authorized users. For more information on this topic, read [“Working with Realms, Users, Groups, and Roles” on page 486](#).
- Configuring secure HTTP and IIOP listeners.

- Configuring secure Java Management Extensions (JMX) connectors.
- Adding, deleting, or modifying existing or custom realms.  
A discussion of this topic is available in [“Realm Configuration” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).
- Defining an interface for pluggable authorization providers using Java Authorization Contract for Containers (JACC).  
Java Authorization Contract for Containers (JACC) defines security contracts between the Enterprise Server and authorization policy modules. These contracts specify how the authorization providers are installed, configured, and used in access decisions. JACC is discussed in [“JACC Support” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).
- Using pluggable audit modules.  
A discussion of using pluggable audit modules can be found in [“Pluggable Audit Module Support” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#)
- Customizing authentication mechanisms  
All implementations of Java EE 6-compatible Servlet containers are required to support the Servlet Profile of JSR-196, which offers an avenue for customizing the authentication mechanism applied by the web container on behalf of one or more applications. Some discussion of using the JSR-196 contract to configure a custom container authentication mechanism is discussed in [“Adding Authentication Mechanisms to the Servlet Container” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).
- Setting and changing policy permissions for an application.  
You can read more about changing policy permissions in [“The server.policy File” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).

The following features are specific to the Enterprise Server:

- Message security  
Message security is discussed in [“Configuring Message Security for Web Services” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).
- Single sign-on across all Enterprise Server applications within a single security domain.  
You can learn more about single sign-on in [“User Authentication for Single Sign-on” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).
- Programmatic login  
You can learn more about programmatic login in [“Programmatic Login” in \*Sun GlassFish Enterprise Server v3 Application Development Guide\*](#).

For more information about configuring the Enterprise Server, read the [Sun GlassFish Enterprise Server v3 Application Development Guide](#) and [Sun GlassFish Enterprise Server v3 Administration Guide](#).

## Working with Realms, Users, Groups, and Roles

You often need to protect resources to ensure that only *authorized users* have access. *Authorization* provides controlled access to protected resources. Authorization is based on identification and authentication. *Identification* is a process that enables recognition of an entity by a system, and *authentication* is a process that verifies the identity of a user, device, or other entity in a computer system, usually as a prerequisite to allowing access to resources in a system. These concepts are discussed in more detail in [“Characteristics of Application Security” on page 478](#).

This section discusses setting up users so that they can be correctly identified and either given access to protected resources, or denied access if the user is not authorized to access the protected resources. To authenticate a user, you need to follow these basic steps:

1. The Application Developer writes code to prompt the user for their user name and password. The different methods of authentication are discussed in [“Specifying an Authentication Mechanism” on page 537](#).
2. The Application Developer communicates how to set up security for the deployed application by use of a deployment descriptor or metadata annotation. This step is discussed in [“Setting Up Security Roles” on page 491](#).
3. The Server Administrator sets up authorized users and groups on the Enterprise Server. This is discussed in [“Managing Users and Groups on the Enterprise Server” on page 489](#).
4. The Application Deployer maps the application’s security roles to users, groups, and principals defined on the Enterprise Server. This topic is discussed in [“Mapping Roles to Users and Groups” on page 493](#).

## What Are Realms, Users, Groups, and Roles?

A realm is a security policy domain defined for a web or application server. It is also a string, passed as part of an HTTP request during basic authentication, that defines a protection space. The protected resources on a server can be partitioned into a set of protection spaces, each with its own authentication scheme and/or authorization database containing a collection of users, which may or may not be assigned to a group. Managing users on the Enterprise Server is discussed in [“Managing Users and Groups on the Enterprise Server” on page 489](#).

An application will often prompt a user for their user name and password before allowing access to a protected resource. After the user has entered their user name and password, that information is passed to the server, which either authenticates the user and sends the protected resource, or does not authenticate the user, in which case access to the protected resource is denied. This type of user authentication is discussed in [“Specifying an Authentication Mechanism” on page 537](#).

In some applications, authorized users are assigned to roles. In this situation, the role assigned to the user in the application must be mapped to a principal or group defined on the application

server. Figure 26–6 shows this. More information on mapping roles to users and groups can be found in “Setting Up Security Roles” on page 491.

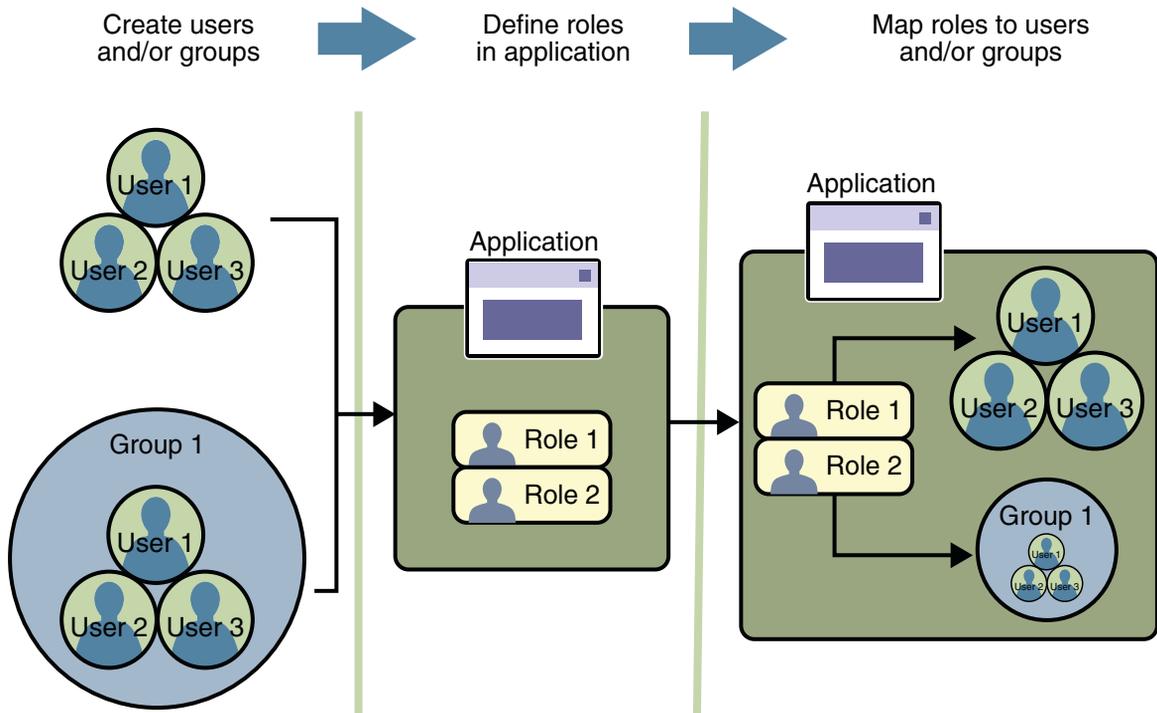


FIGURE 26–6 Mapping Roles to Users and Groups

**Note** – The concepts of users, groups, principals, and realms exist in most application or web servers, but might use different names in different products. If you are using a server other than Enterprise Server, consult your product's documentation for the terminology specific to that server.

The following sections provide more information on realms, users, groups, and roles.

## What Is a Realm?

A realm is a security policy domain defined for a web or application server. It is also a string, passed as part of an HTTP request during basic authentication, that defines a protection space. The protected resources on a server can be partitioned into a set of protection spaces, each with its own authentication scheme and/or authorization database containing a collection of users,

which may or may not be assigned to a group. For a web application, a *realm* is a complete database of *users* and *groups* that identify valid users of a web application (or a set of web applications) and are controlled by the same authentication policy.

The Java EE server authentication service can govern users in multiple realms. In this release of the Enterprise Server, the `file`, `admin-realm`, and `certificate` realms come preconfigured for the Enterprise Server.

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. When using the `file` realm, the server authentication service verifies user identity by checking the `file` realm. This realm is used for the authentication of all clients except for web browser clients that use the HTTPS protocol and certificates.

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. To verify the identity of a user in the `certificate` realm, the authentication service verifies an X.509 certificate. For step-by-step instructions for creating this type of certificate, see [“Working with Digital Certificates” on page 497](#). The common name field of the X.509 certificate is used as the principal name.

The `admin-realm` is also a `FileRealm` and stores administrator user credentials locally in a file named `admin-keyfile`. You can 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 Users and Groups on the Enterprise Server” on page 489](#).

## What Is a User?

A *user* is an individual (or application program) identity that has been defined in the *Enterprise Server*. In a web application, a user can have a set of *roles* associated with that identity, which entitles them to access all resources protected by those roles. Users can be associated with a group.

A Java EE user is similar to an operating system user. Typically, both types of users represent people. However, these two types of users are not the same. The Java EE server authentication service has no knowledge of the user name and password you provide when you log on to the operating system. The Java EE server authentication service is not connected to the security mechanism of the operating system. The two security services manage users that belong to different realms.

## What Is a Group?

A *group* is a set of authenticated *users*, classified by common traits, defined in the *Enterprise Server*.

A Java EE user of the `file` realm can belong to an Enterprise Server group. (A user in the `certificate` realm cannot.) An Enterprise Server *group* is a category of users classified by common traits, such as job title or customer profile. For example, most 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.

An Enterprise Server *group* has a different scope from a *role*. An Enterprise Server group is designated for the entire Enterprise Server, whereas a role is associated only with a specific application in the Enterprise Server.

## What Is a Role?

A *role* is an abstract name for the permission to access a particular set of resources in an *application*. A *role* can be compared to a key that can open a lock. Many people might have a copy of the key. The lock doesn't care who you are, only that you have the right key.

## Some Other Terminology

The following terminology is also used to describe the security requirements of the Java EE platform:

- **Principal:** A principal is an entity that can be authenticated by an authentication protocol in a security service that is deployed in an enterprise. A principal is identified using a principal name and authenticated using authentication data.
- **Security policy domain** (also known as **security domain** or **realm**): A security policy domain is a scope over which a common security policy is defined and enforced by the security administrator of the security service.
- **Security attributes:** A set of security attributes is associated with every principal. The security attributes have many uses, for example, access to protected resources and auditing of users. Security attributes can be associated with a principal by an authentication protocol.
- **Credential:** A credential contains or references information (security attributes) used to authenticate a principal for Java EE product services. A principal acquires a credential upon authentication, or from another principal that allows its credential to be used.

## Managing Users and Groups on the Enterprise Server

Managing users on the Enterprise Server is discussed in more detail in the [Sun GlassFish Enterprise Server v3 Administration Guide](#).

This tutorial provides steps for managing users that will need to be completed to work through the tutorial examples.

## Adding Users to the Enterprise Server

To add users to the Enterprise Server, follow these steps:

1. Start the Enterprise Server if you haven't already done so. Information on starting the Enterprise Server is available in [“Starting and Stopping the Enterprise Server” on page 63](#).
2. Start the Admin Console if you haven't already done so. You can start the Admin Console by starting a web browser and entering the URL `http://localhost:4848/asadmin`. If you changed the default Admin port during installation, enter the correct port number in place of 4848.
3. To log in to the Admin Console, enter the user name and password of a user in the `admin-realm` who belongs to the `asadmin` group. The name and password entered during installation will work, as will any users added to this realm and group subsequent to installation.
4. Expand the Configuration node in the Admin Console tree.
5. Expand the Security node in the Admin Console tree.
6. Expand the Realms node.
  - Select the `file` realm to add users you want to access applications running in this realm. (For the example security applications, select the `file` realm.)
  - Select the `admin-realm` to add users you want to enable as system administrators of the Enterprise Server.
  - You cannot enter users into the `certificate` realm using the Admin Console. You can only add certificates to the `certificate` realm. For information on adding (importing) certificates to the `certificate` realm, read [“Adding Users to the Certificate Realm” on page 491](#).
7. Click the Manage Users button.
8. Click New to add a new user to the realm.
9. Enter the correct information into the User ID, Password, and Group(s) fields.
  - If you are adding a user to the `file` realm, enter a name to identify the user, a password to allow the user access to the realm, and a group to which this user belongs. For more information on these properties, read [“Working with Realms, Users, Groups, and Roles” on page 486](#).  
For the example security applications, enter a user with any name and password you like, but make sure that the user is assigned to the group of `TutorialUser`. The user name and password are case-sensitive. Keep a record of the user name and password for working with examples later in this tutorial.
  - If you are adding a user to the `admin-realm`, enter the name to identify the user, a password to allow the user access to the Enterprise Server, and enter `asadmin` in the Group field. Restart the Enterprise Server and Admin Console after making this change.
10. Click OK to add this user to the list of users in the realm.

11. Click Logout when you have completed this task.

## Adding Users to the Certificate Realm

In the certificate realm, user identity is set up in the Enterprise Server security context and populated with user data obtained from cryptographically-verified client certificates. For step-by-step instructions for creating this type of certificate, see [“Working with Digital Certificates” on page 497](#).

## Setting Up Security Roles

When you design an enterprise bean or web component, you should always think about the kinds of users who will access the component. For example, a web application for a human resources department might have a different request URL for someone who has been assigned the role of `DEPT_ADMIN` than for someone who has been assigned the role of `DIRECTOR`. The `DEPT_ADMIN` role may let you view employee data, but the `DIRECTOR` role enables you to modify employee data, including salary data. Each of these security roles is an abstract logical grouping of users that is defined by the person who assembles the application. When an application is deployed, the deployer will map the roles to security identities in the operational environment, as shown in [Figure 26–6](#).

For enterprise applications, you define security roles using the `@DeclareRoles` and `@RolesAllowed` metadata annotations. For web applications, you define roles in the application deployment descriptor, `web.xml`.

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

The following is an example of an application where the role of `DEPT-ADMIN` is authorized for methods that review employee payroll data and the role of `DIRECTOR` is authorized for methods that change employee payroll data.

The enterprise bean would be annotated as shown in the following code:

```
@DeclareRoles({"DEPT-ADMIN", "DIRECTOR"})
@Stateless public class PayrollBean implements Payroll {
    @Resource SessionContext ctx;

    @RolesAllowed("DEPT-ADMIN")
    public void reviewEmployeeInfo(EmplInfo info) {
```

```
        oldInfo = ... read from database;

        /           ...
    }
    @RolesAllowed("DIRECTOR")
    public void updateEmployeeInfo(EmplInfo info) {

        newInfo = ... update database;

        ...
    }
    ...
}
```

The deployment descriptor would include security constraints, as shown in the following example:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>view dept data</web-resource-name>
    <url-pattern>/hr/employee/*</url-pattern>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>DEPT_ADMIN</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>

<security-constraint>
  <web-resource-collection>
    <web-resource-name>change dept data</web-resource-name>
    <url-pattern>/hr/employee/*</url-pattern>
    <http-method>GET</http-method>
    <http-method>PUT</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>DIRECTOR</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>
```

The `web.xml` application deployment descriptor is described in more detail in [“Introduction to Web Application Deployment Descriptors”](#) on page 530.

These annotations are discussed in more detail in [“Securing an Enterprise Bean Using Declarative Security and Annotations”](#) on page 507.

After users have provided their login information, and the application has declared what roles are authorized to access protected parts of an application, the next step is to map the security role to the name of a user, or principal. This step is discussed in the following section.

## Mapping Roles to Users and Groups

When you are developing a Java EE application, you don’t need to know what categories of users have been defined for the realm in which the application will be run. In the Java EE platform, the security architecture provides a mechanism for mapping the roles defined in the application to the users or groups defined in the runtime realm. To map a role name permitted by the application or module to principals (users) and groups defined on the server, use the `security-role-mapping` element in the runtime deployment descriptor (`sun-application.xml`, `sun-web.xml`, or `sun-ejb-jar.xml`) file. The entry needs to declare a mapping between a security role used in the application and one or more groups or principals defined for the applicable realm of the Enterprise Server. An example for the `sun-web.xml` file is shown below:

```
<sun-web-app>
  <security-role-mapping>
    <role-name>DIRECTOR</role-name>
    <principal-name>schwartz</principal-name>
  </security-role-mapping>
  <security-role-mapping>
    <role-name>DEPT-ADMIN</role-name>
    <group-name>dept-admins</group-name>
  </security-role-mapping>
</sun-web-app>
```

The role name can be mapped to either a specific principal (user), a group, or both. The principal or group names referenced must be valid principals or groups in the current default realm of the Enterprise Server. The `role-name` in this example must exactly match the `role-name` in the `security-role` element of the corresponding `web.xml` file or the role name defined in the `@DeclareRoles` and/or `@RolesAllowed` annotations.

Sometimes the role names used in the application are the same as the group names defined on the Enterprise Server. Under these circumstances, you can enable a default principal-to-role mapping on the Enterprise Server using the Admin Console. From the Admin Console, select Configuration, then Security, then check the enable box beside Default Principal to Role Mapping. If you need more information about using the Admin Console, see [“Adding Users to the Enterprise Server”](#) on page 490 or its online help.

## Establishing a Secure Connection Using SSL

*Secure Socket Layer* (SSL) technology is security that is implemented at the transport layer (see “[Transport-Layer Security](#)” on page 480, for more information about transport layer security). SSL allows web browsers and web servers to communicate over a secure connection. In this secure connection, the data that is being sent is encrypted before being sent and then is decrypted upon receipt and before processing. Both the browser and the server encrypt all traffic before sending any data. SSL addresses the following important security considerations.

- **Authentication:** During your initial attempt to communicate with a web server over a secure connection, that server will present your web browser with a set of credentials in the form of a server certificate. The purpose of the certificate is to verify that the site is who and what it claims to be. In some cases, the server may request a certificate that the client is who and what it claims to be (which is known as client authentication).
- **Confidentiality:** When data is being passed between the client and the server on a network, third parties can view and intercept this data. SSL responses are encrypted so that the data cannot be deciphered by the third party and the data remains confidential.
- **Integrity:** When data is being passed between the client and the server on a network, third parties can view and intercept this data. SSL helps guarantee that the data will not be modified in transit by that third party.

## Installing and Configuring SSL Support

An SSL HTTPS connector is already enabled in the Enterprise Server. For more information on configuring SSL for the Enterprise Server, refer to the [Sun GlassFish Enterprise Server v3 Administration Guide](#).

If you are using a different application server or web server, an SSL HTTPS connector might or might not be enabled. If you are using a server that needs its SSL connector to be configured, consult the documentation for that server.

As a general rule, to enable SSL for a server, you must address the following issues:

- There must be a Connector element for an SSL connector in the server deployment descriptor.
- There must be valid keystore and certificate files.
- The location of the keystore file and its password must be specified in the server deployment descriptor.

You can verify whether or not SSL is enabled by following the steps in “[Verifying SSL Support](#)” on page 496.

## Specifying a Secure Connection in Your Application Deployment Descriptor

To specify a requirement that protected resources be received over a protected transport layer connection (SSL), specify a user data constraint in the application deployment descriptor. The following is an example of a `web.xml` application deployment descriptor that specifies that SSL be used:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>view_dept_data</web-resource-name>
    <url-pattern>/hr/employee/*</url-pattern>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>DEPT_ADMIN</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>
```

A user data constraint (`<user-data-constraint>` in the deployment descriptor) requires that all constrained URL patterns and HTTP methods specified in the security constraint are received over a protected transport layer connection such as HTTPS (HTTP over SSL). A user data constraint specifies a transport guarantee (`<transport-guarantee>` in the deployment descriptor). The choices for transport guarantee include `CONFIDENTIAL`, `INTEGRAL`, or `NONE`. If you specify `CONFIDENTIAL` or `INTEGRAL` as a security constraint, that type of security constraint applies to all requests that match the URL patterns in the web resource collection and not just to the login dialog box.

The strength of the required protection is defined by the value of the transport guarantee.

- Specify `CONFIDENTIAL` when the application requires that data be transmitted so as to prevent other entities from observing the contents of the transmission.
- Specify `INTEGRAL` when the application requires that the data be sent between client and server in such a way that it cannot be changed in transit.
- Specify `NONE` to indicate that the container must accept the constrained requests on any connection, including an unprotected one.

The user data constraint is handy to use with basic and form-based user authentication. When the login authentication method is set to `BASIC` or `FORM`, passwords are not protected, meaning that passwords sent between a client and a server on an unprotected session can be viewed and intercepted by third parties. Using a user data constraint with the user authentication mechanism can alleviate this concern.

## Verifying SSL Support

For testing purposes, and to verify that SSL support has been correctly installed, load the default introduction page with a URL that connects to the port defined in the server deployment descriptor:

```
https://localhost:8181/
```

The `https` in this URL indicates that the browser should be using the SSL protocol. The `localhost` in this example assumes that you are running the example on your local machine as part of the development process. The `8181` in this example is the secure port that was specified where the SSL connector was created. If you are using a different server or port, modify this value accordingly.

The first time that you load this application, the New Site Certificate or Security Alert dialog box displays. Select Next to move through the series of dialog boxes, and select Finish when you reach the last dialog box. The certificates will display only the first time. When you accept the certificates, subsequent hits to this site assume that you still trust the content.

## Tips on Running SSL

The SSL protocol is designed to be as efficient as securely possible. However, encryption and decryption are computationally expensive processes from a performance standpoint. It is not strictly necessary to run an entire web application over SSL, and it is customary for a developer to decide which pages require a secure connection and which do not. Pages that might require a secure connection include login pages, personal information pages, shopping cart checkouts, or any pages where credit card information could possibly be transmitted. Any page within an application can be requested over a secure socket by simply prefixing the address with `https`: instead of `http`:. Any pages that absolutely require a secure connection should check the protocol type associated with the page request and take the appropriate action if `https` is not specified.

Using name-based virtual hosts on a secured connection 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 before 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 should not interfere with normal SSL operations on the server. Be aware, however, that most client browsers will compare the server's domain name against the domain name listed in the certificate, if any (this is applicable primarily to official, CA-signed certificates). If the domain names do not match, these browsers will display a warning to the client. In general, only address-based virtual hosts are commonly used with SSL in a production environment.

## Working with Digital Certificates

Digital certificates for the Enterprise Server have already been generated and can be found in the directory `as-install/domain-dir/config/`. These digital certificates are self-signed and are intended for use in a development environment; they are not intended for production purposes. For production purposes, generate your own certificates and have them signed by a CA.

The instructions in this section apply to the developer and cluster profiles of the Enterprise Server. In the enterprise profile, the `certutil` utility is used to create digital certificates. For more information, see the *Sun GlassFish Enterprise Server v3 Administration Guide*.

To use SSL, an application or web server must have an associated certificate for each external interface, or IP address, that accepts secure connections. The theory behind this design is that a server should provide some kind of reasonable assurance that its owner is who you think it is, particularly before receiving any sensitive information. It may be useful to think of a certificate as a “digital driver’s license” for an Internet address. It states with which company the site is associated, along with some basic contact information about the site owner or administrator.

The digital certificate is cryptographically signed by its owner and is difficult for anyone else to forge. For sites involved in e-commerce or in any other business transaction in which authentication of identity is important, a certificate can be purchased from a well-known certificate authority (CA) such as VeriSign or Thawte. If your server certificate is self-signed, you must install it in the Enterprise Server keystore file (`keystore.jks`). If your client certificate is self-signed, you should install it in the Enterprise Server truststore file (`cacerts.jks`).

Sometimes authentication is not really a concern. For example, an administrator might simply want to ensure that data being transmitted and received by the server is private and cannot be snooped by anyone eavesdropping on the connection. In such cases, you can save the time and expense involved in obtaining a CA certificate and simply use a self-signed certificate.

SSL uses *public key cryptography*, which is based on *key pairs*. Key pairs contain one public key and one private key. If data is encrypted with one key, it can be decrypted only with the other key of the pair. This property is fundamental to establishing trust and privacy in transactions. For example, using SSL, the server computes a value and encrypts the value using its private key. The encrypted value is called a *digital signature*. The client decrypts the encrypted value using the server’s public key and compares the value to its own computed value. If the two values match, the client can trust that the signature is authentic, because only the private key could have been used to produce such a signature.

Digital certificates are used with the HTTPS protocol to authenticate web clients. The HTTPS service of most web servers will not run unless a digital certificate has been installed. Use the procedure outlined in the next section, “[Creating a Server Certificate](#)” on page 498, to set up a digital certificate that can be used by your application or web server to enable SSL.

One tool that can be used to set up a digital certificate is `keytool`, a key and certificate management utility that ships with the Java SE SDK. It enables users to administer their own public/private key pairs and associated certificates for use in self-authentication (where the user

authenticates himself or herself to other users or services) or data integrity and authentication services, using digital signatures. It also allows users to cache the public keys (in the form of certificates) of their communicating peers. For a better understanding of `keytool` and public key cryptography, read the `keytool` documentation at <http://java.sun.com/javase/6/docs/technotes/tools/solaris/keytool.html>.

## Creating a Server Certificate

A server certificate has already been created for the Enterprise Server. The certificate can be found in the `domain-dir/config/` directory. The server certificate is in `keystore.jks`. The `cacerts.jks` file contains all the trusted certificates, including client certificates.

If necessary, you can use `keytool` to generate certificates. The `keytool` utility stores the keys and certificates in a file termed a *keystore*, a repository of certificates used for identifying a client or a server. Typically, a keystore is a file that contains one client or one server's identity. It protects private keys by using a password.

If you don't specify a directory when specifying the keystore file name, the keystores are created in the directory from which the `keytool` command is run. This can be the directory where the application resides, or it can be a directory common to many applications.

To create a server certificate, follow these steps:

1. Create the keystore.
2. Export the certificate from the keystore.
3. Sign the certificate.
4. Import the certificate into a *truststore*: a repository of certificates used for verifying the certificates. A truststore typically contains more than one certificate.

Run `keytool` to generate a new key pair in the default development keystore file, `keystore.jks`. This example uses the alias `server-alias` to generate a new public/private key pair and wrap the public key into a self-signed certificate inside `keystore.jks`. The key pair is generated using an algorithm of type `RSA`, with a default password of `changeit`. For more information and other examples of creating and managing keystore files, read the `keytool` online help at <http://java.sun.com/javase/6/docs/technotes/tools/solaris/keytool.html>.

---

**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.

---

From the directory in which you want to create the key pair, run `keytool` with the following parameters.

1. Generate the server certificate. (Type the `keytool` command all on one line.)

```
java-home/bin/keytool -genkey -alias server-alias -keyalg RSA -keypass changeit
-storepass changeit -keystore keystore.jks
```

When you press Enter, `keytool` prompts you to enter the server name, organizational unit, organization, locality, state, and country code.

You must enter the server name in response to `keytool`'s first prompt, in which it asks for first and last names. For testing purposes, this can be `localhost`.

When you run the example applications, the host (server name) specified in the keystore must match the host identified in the `javae.server.name` property specified in the file `tut-install/examples/bp-project/build.properties`.

- Export the generated server certificate in `keystore.jks` into the file `server.cer`. (Type the `keytool` command all on one line.)

```
java-home/bin/keytool -export -alias server-alias -storepass changeit
-file server.cer -keystore keystore.jks
```

- If you want to have the certificate signed by a CA, read the example at <http://java.sun.com/javase/6/docs/technotes/tools/solaris/keytool.html>.
- To add the server certificate to the truststore file, `cacerts.jks`, run `keytool` from the directory where you created the keystore and server certificate. Use the following parameters:

```
java-home/bin/keytool -import -v -trustcacerts -alias server-alias -file server.cer
-keystore cacerts.jks -keypass changeit -storepass changeit
```

Information on the certificate, such as that shown next, will display.

```
% keytool -import -v -trustcacerts -alias server-alias -file server.cer
-keystore cacerts.jks -keypass changeit -storepass changeit
Owner: CN=localhost, OU=Sun Micro, O=Docs, L=Santa Clara, ST=CA,
C=USIssuer: CN=localhost, OU=Sun Micro, O=Docs, L=Santa Clara, ST=CA,
C=USSerial number: 3e932169Valid from: Tue Apr 08Certificate
fingerprints:MD5: 52:9F:49:68:ED:78:6F:39:87:F3:98:B3:6A:6B:0F:90 SHA1:
EE:2E:2A:A6:9E:03:9A:3A:1C:17:4A:28:5E:97:20:78:3F:
Trust this certificate? [no]:
```

- Enter `yes`, and then press the Enter or Return key. The following information displays:

```
Certificate was added to keystore[Saving cacerts.jks]
```

## Miscellaneous Commands for Certificates

To check the contents of a keystore that contains a certificate with an alias `server-alias`, use this command:

```
keytool -list -keystore keystore.jks -alias server-alias -v
```

To check the contents of the `cacerts` file, use this command:

```
keytool -list -keystore cacerts.jks
```

## Further Information about Security

For more information about security in Java EE applications, see:

- *Java EE 6 Specification*:  
<http://jcp.org/en/jsr/detail?id=316>
- The *Sun GlassFish Enterprise Server v3 Application Development Guide* includes security information for application developers.
- The *Sun GlassFish Enterprise Server v3 Administration Guide* includes information on setting security settings for the Enterprise Server.
- The *Sun GlassFish Enterprise Server v3 Application Deployment Guide* includes information on security settings in the deployment descriptors specific to the Enterprise Server.
- *EJB 3.1 Specification (JSR-318)*:  
<http://jcp.org/en/jsr/detail?id=318>
- *Implementing Enterprise Web Services 1.3 (JSR-109)*:  
<http://jcp.org/en/jsr/detail?id=109>
- Java Platform, Standard Edition 6 security information:  
<http://java.sun.com/javase/6/docs/technotes/guides/security/>
- *Java Servlet Specification, Version 3.0*:  
<http://jcp.org/en/jsr/detail?id=315>
- *JSR 181: Web Services Metadata for the Java Platform*:  
<http://jcp.org/en/jsr/detail?id=181>
- *JSR 250: Common Annotations for the Java Platform*:  
<http://jcp.org/en/jsr/detail?id=250>
- The Java SE discussion of annotations:  
<http://java.sun.com/javase/6/docs/technotes/guides/language/annotations.html>
- *JSR 115: Java Authorization Contract for Containers 1.3*:  
<http://jcp.org/en/jsr/detail?id=115>
- Chapter 24 of the CORBA/IIOP specification, *Secure Interoperability*:  
<http://www.omg.org/cgi-bin/doc?formal/02-06-60>
- *Java Authentication and Authorization Service (JAAS) in Java Platform, Standard Edition*:  
<http://java.sun.com/developer/technicalArticles/Security/jaasv2/index.html>
- *Java Authentication and Authorization Service (JAAS) Reference Guide*:  
<http://java.sun.com/javase/6/docs/technotes/guides/security/jaas/JAASRefGuide.html>

- *Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide:*  
[http://java.sun.com/  
javase/6/docs/technotes/guides/security/jaas/JAASLMDevGuide.html](http://java.sun.com/javase/6/docs/technotes/guides/security/jaas/JAASLMDevGuide.html)



# Getting Started Securing Enterprise Applications

---

*Enterprise applications* provide business logic support functionality for an enterprise, typically in commercial organizations, to improve the enterprise's productivity and efficiency. Services provided by enterprise applications are typically business-oriented applications such as online shopping and online payment processing, interactive product catalogue, automated billing systems, security, content management, CRM, ERP, Business Intelligence, HR management, manufacturing, EAI, Enterprise Forms Automation, etc. Enterprise applications typically have interfaces to other enterprise software (for example, from a database to Enterprise JavaBeans™) and are centrally managed.<sup>1</sup>

## Responsibility for Administering Security

The following parties are responsible for administering security for enterprise applications:

- **System Administrator**

The system administrator is responsible for setting up a database of users and assigning those users to the proper group. The system administrator is also responsible for setting properties on the Enterprise Server that enable the applications to run properly. In terms of security, some examples include setting up a default principal to role mapping, anonymous users, default users, and propagated identities. More information on system administrator responsibilities is found in *Sun GlassFish Enterprise Server v3 Administration Guide*. When needed for this tutorial, the steps for performing specific steps will also be provided herein.
- **Application Developer/Bean Provider**

The application developer/bean provider is responsible for annotating the classes and methods of the enterprise application in order to provide information to the deployer about which methods need to have restricted access. This tutorial describes the steps necessary to complete this task.

---

<sup>1</sup> Enterprise software. (2009, October 20). In Wikipedia, The Free Encyclopedia. Retrieved 19:41, October 29, 2009, from [http://en.wikipedia.org/w/index.php?title=Enterprise\\_software](http://en.wikipedia.org/w/index.php?title=Enterprise_software)

- **Deployer**

The deployer is responsible for taking the security view provided by the application developer and implementing that security upon deployment. This document provides the information needed to accomplish this task for the tutorial example applications. For more information on deployment, the best source for the Enterprise Server is the [Sun GlassFish Enterprise Server v3 Application Deployment Guide](#).

## Securing Enterprise Beans

Enterprise beans are Java EE components that implement *Enterprise JavaBeans* (EJB) technology. Enterprise beans run in the EJB container, a runtime environment within the Enterprise Server. Although transparent to the application developer, the EJB container provides system-level services such as transactions and *security* to its enterprise beans, which form the core of transactional Java EE applications.

Enterprise bean methods can be secured using one of the following methods:

- *Declarative Security* (preferred)

Declarative security enables the application developer to specify which users are authorized to access which methods of the enterprise beans, what type of authentication will be used, and whether or not the protected data will use a secure connection.

In older versions of Java EE, declarative security was specified in the application *deployment descriptor*. A deployment descriptor is an XML file that is external to the application and expresses, among other things, an application's security structure. With the introduction of Java EE 5, *metadata annotations* (or simply, *annotations*,) were introduced to specify which users were authorized to access protected methods of the enterprise applications. Beginning with Java EE 6, the presence of an annotation in the business method of an enterprise application that specifies method permissions is all that is needed for method protection and authentication in some situations. This section discusses this simple and efficient method of securing enterprise beans

There are some limitations to the simplified method of securing enterprise beans, so there are some instances where you would want to continue to use the deployment descriptor to specify security information. An authentication mechanism must be configured on the server for the simple solution to work. The Enterprise Server assumes a default authentication method of basic authentication, but not all servers are configured for a default authentication mechanism, in which case you would need to use a vendor-specific deployment descriptor to specify an authentication method. Using deployment descriptors to specify security information for enterprise applications is discussed in Part VII, "Security," in *The Java EE 6 Tutorial, Volume II*.

For this tutorial, we will describe how to invoke username-password authentication of authorized users by decorating the enterprise application's business methods with *annotations* that specify method permissions. To make the deployer's task easier, the application developer can define security roles. A security role is a grouping of permissions

that a given type of users of the application must have in order to successfully use the application. For example, in a payroll application, there will be users who want to view their own payroll information (*employee*), user who need to view others payroll information (*manager*), and users who need to be able to change others payroll information (*payrollDept*). The application developer would determine who the potential users of the application would be, and which methods would be accessible to which users. The application developer would then decorate classes or methods of the enterprise bean with annotations that specify the types of users authorized to access those methods. Using annotations to specify authorized users is described in [“Specifying Authorized Users by Declaring Security Roles” on page 508](#).

When one of the annotations is used to define method permissions, the deployment system will automatically require username-password authentication. Username-password authentication is used often in applications. In this type of authentication, a user will be prompted to enter their username and password. These will be compared against a database of known users. If the user is found and the password matches, the roles that the user is assigned will be compared against the roles that are authorized to access the method. If the user is authenticated and found to have a role that is authorized to access that method, the data will be returned to the user.

Using declarative security is discussed in [“Securing an Enterprise Bean Using Declarative Security and Annotations” on page 507](#).

- *Programmatic Security*

Programmatic security is code that is embedded in a business method, is used to access a caller's identity programmatically, and uses this information to make security decisions within the method itself. Programmatic security is useful when declarative security alone is not sufficient to express the security model of an application.

In general, security management should be enforced by the container in a manner that is transparent to the enterprise beans' business methods. The programmatic security API's described in this chapter should be used only in the less frequent situations in which the enterprise bean business methods need to access the security context information, such as when you want to grant access based on the time of day (or other nontrivial condition checks) for a particular role.

Programmatic security is discussed in [“Securing an Enterprise Bean Programmatically” on page 515](#).

Some of the material in this chapter assumes that you have already read [Chapter 26, “Introduction to Security in the Java EE Platform,”](#) [Chapter 16, “Enterprise Beans,”](#) and [Chapter 17, “Getting Started with Enterprise Beans.”](#)

As mentioned earlier, enterprise beans run in the EJB container, a runtime environment within the Enterprise Server, as shown in [Figure 27–1](#).

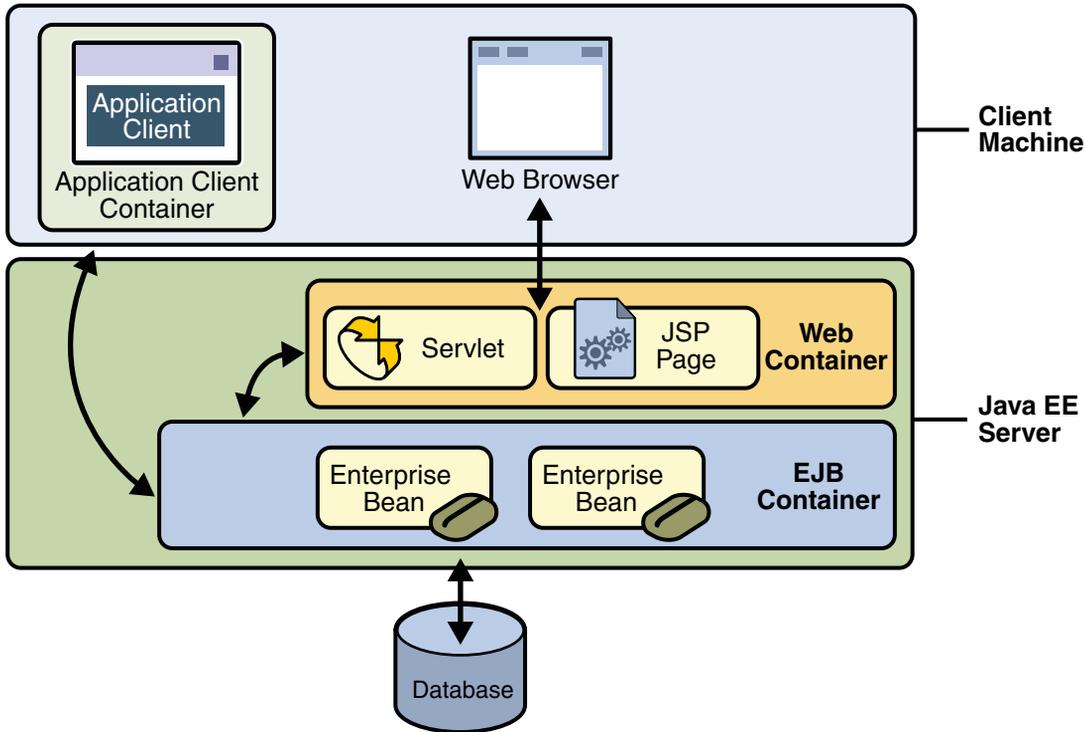


FIGURE 27-1 Java EE Server and Containers

This section discusses how to secure a Java EE application where one or more modules (such as EJB JAR files) are packaged into an EAR file, the archive file that holds the application. Security annotations will be used in the Java programming class files to specify authorized users and basic, or username-password, authentication.

Enterprise beans often provide the business logic of a *web* application. In these cases, packaging the enterprise bean within the web application's WAR module simplifies deployment and application organization. Enterprise beans may be packaged within a WAR module as Java programming language class files or within a JAR file that is bundled within the WAR module. When a servlet or JSP page handles the web front-end, and the application is packaged into a WAR module as a Java programming class file, then security for the *application* can be handled in the application's `web.xml` file. The EJB in the WAR file can have its own deployment descriptor, `ejb-jar.xml`, if required. Securing web applications using `web.xml` is discussed in [Chapter 28, "Getting Started Securing Web Applications,"](#) and Part VII, "Security," in *The Java EE 6 Tutorial, Volume II*.

The following sections describe declarative and programmatic security mechanisms that can be used to protect enterprise bean resources. The protected resources include methods of enterprise beans that are called from application clients, web components, or other enterprise beans.

You can protect enterprise beans by doing the following:

- “Securing an Enterprise Bean Using Declarative Security and Annotations” on page 507
- “Securing an Enterprise Bean Programmatically” on page 515
- “Propagating a Security Identity (Run-As)” on page 518
- “Deploying Secure Enterprise Beans” on page 520

You should also read *JSR-318: Enterprise JavaBeans 3.1* for more information on this topic. This document can be downloaded from <http://jcp.org/en/jsr/detail?id=318>. Chapter 17 of this specification, *Security Management*, discusses security management for enterprise beans.

## Securing an Enterprise Bean Using Declarative Security and Annotations

Declarative security enables the application developer to specify which users are authorized to access which methods of the enterprise beans, and to authenticate these users with basic, or username-password, authentication.

Frequently, the person who is developing an enterprise application is not the same person who is responsible for deploying the application. When an application developer uses declarative security to define method permissions and authentications mechanisms, they are passing along to the deployer a *security view* of the enterprise beans contained in the EJB JAR. When a security view is passed on to the deployer, the deployer uses this information to define method permissions for security roles. If you don't define a security view, the deployer will have to determine what each business method does to determine which users are authorized to call each method.

A security view consists of a set of *security roles*, a semantic grouping of permissions that a given type of users of an application must have to successfully access the application. Security roles are meant to be logical roles, representing a type of user. You can define method permissions for each security role. A method permission is a permission to invoke a specified group of methods of the enterprise beans' business interface, home interface, component interface, and/or web service endpoints. After method permissions are defined, username-password authentication will be used to verify the identity of the user.

It is important to keep in mind that security roles are used to define the logical security view of an application. They should not be confused with the user groups, users, principals, and other concepts that exist in the Enterprise Server. An additional step is required to map the roles defined in the application to users, groups, and principals that are the components of the user database in the file realm of the Enterprise Server. These steps are outlined in “[Mapping Security Roles to Enterprise Server Groups](#)” on page 511.

The following sections show how an application developer uses declarative security to either secure an application or to create a security view to pass along to the deployer.

- “Specifying Authorized Users by Declaring Security Roles” on page 508
- “Specifying an Authentication Mechanism and Secure Connection” on page 512
- “Example: Securing an Enterprise Bean” on page 512

## Specifying Authorized Users by Declaring Security Roles

This section discusses how to use annotations to specify the method permissions for the methods of a bean class. If you'd like more information on these annotations, refer to [JSR-250 Common Annotations for the Java Platform](#).

Method permissions can be specified on the class, the business methods of the class, or both. Method permissions can be specified on a method of the bean class to override the method permissions value specified on the entire bean class. The following annotations are used to specify method permissions:

- `@DeclareRoles`

This annotation is used to specify all of the roles that will be used by the application, including roles not specifically named in a `RolesAllowed` annotation. The set of security roles used by the application is the total of the security roles defined in the `@DeclareRoles` and `@RolesAllowed` annotations.

The `@DeclareRoles` annotation is specified on a bean class, where it serves to declare roles that can be tested (for example, by calling `isCallerInRole`) from within the methods of the annotated class. When declaring the name of a role used as a parameter to the `isCallerInRole(String roleName)` method, the declared name must be the same as the parameter value.

Here is some example code that demonstrates the use of the `DeclareRoles` annotation.

```
@DeclareRoles("BusinessAdmin")
public class Calculator {
    public void convertCurrency() {
        if(x.isUserInRole("BusinessAdmin")) {
            //....
        }
    }
    //...
}
```

The syntax for declaring more than one role is as shown in the following example:

```
@DeclareRoles({"Administrator", "Manager", "Employee"})
```

- `@RolesAllowed("list-of-roles")`

The `@RolesAllowed` annotation specifies the security roles permitted to access method(s) in an application. This annotation can be specified on a class or on method(s). When specified at the class level, it applies to all methods in the class. When specified on a method, it applies to that method only, and overrides any values specified at the class level.

To specify that no roles are authorized to access method(s) in an application, use the `@DenyAll` annotation. To specify that a user in any role is authorized to access the application, use the `@PermitAll` annotation.

When used in conjunction with the `@DeclareRoles` annotation, the combined set of security roles are used by the application.

Here is some example code that demonstrates the use of the `RolesAllowed` annotation.

```
@RolesAllowed("AllUsers")
public class Calculator {
    @RolesAllowed("Administrator")
    public void setNewRate(int rate) {
        //..
    }
}
```

- `@PermitAll`

The `@PermitAll` annotation specifies that *all* security roles are permitted to execute the specified method(s). The user is not checked against a database to ensure that this user is authorized to access this application.

This annotation can be specified on a class or on method(s). Specifying this annotation on the class means that it applies to all methods of the class. Specifying it at the method level means that it applies to only that method.

Here is some example code that demonstrates the use of the `PermitAll` annotation.

```
import javax.annotation.security.*;
@RolesAllowed("RestrictedUsers")
public class Calculator {
    @RolesAllowed("Administrator")
    public void setNewRate(int rate) {
        //...
    }
    @PermitAll
    public long convertCurrency(long amount) {
        //...
    }
}
```

- `@DenyAll`

The `@DenyAll` annotation specifies that *no* security roles are permitted to execute the specified method(s). This means that these methods are excluded from execution in the Java EE container.

Here is some example code that demonstrates the use of the `DenyAll` annotation.

```
import javax.annotation.security.*;
@RolesAllowed("Users")
public class Calculator {
    @RolesAllowed("Administrator")
    public void setNewRate(int rate) {
        //...
    }
    @DenyAll
    public long convertCurrency(long amount) {
        //...
    }
}
```

**EXAMPLE 27-1** Declaring Roles using `@DeclareRoles`

The following code snippet demonstrates the use of the `@DeclareRoles` annotation with the `isCallerInRole` method. In this example, the `@DeclareRoles` annotation declares a role that the enterprise bean `PayrollBean` uses to make the security check using `isCallerInRole("payroll")` to verify that the caller is authorized to change salary data.

```
@DeclareRoles("payroll")
@Stateless public class PayrollBean implements Payroll {
    @Resource SessionContext ctx;

    public void updateEmployeeInfo(EmplInfo info) {

        oldInfo = ... read from database;

        // The salary field can be changed only by callers
        // who have the security role "payroll"
        if (info.salary != oldInfo.salary &&
            !ctx.isCallerInRole("payroll")) {
            throw new SecurityException(...);
        }
        ...
    }
    ...
}
```

**EXAMPLE 27-2** Declaring Roles using `@RolesAllowed`

The following example code illustrates the use of the `RolesAllowed` annotation:

```
@RolesAllowed("admin")
public class SomeClass {
    public void aMethod () {...}
    public void bMethod () {...}
}
```

EXAMPLE 27-2 Declaring Roles using `@RolesAllowed` (Continued)

```

    ...
}

@Stateless public class MyBean extends SomeClass implements A {

    @RolesAllowed("HR")
    public void aMethod () {...}

    public void cMethod () {...}
    ...
}

```

In this example, assuming `aMethod`, `bMethod`, and `cMethod` are methods of business interface `A`, the method permissions values of methods `aMethod` and `bMethod` are `@RolesAllowed("HR")` and `@RolesAllowed("admin")` respectively. The method permissions for method `cMethod` have not been specified.

To clarify, the annotations are not inherited by the subclass per se, they apply to methods of the superclass which are inherited by the subclass.

## Mapping Security Roles to Enterprise Server Groups

The Enterprise Server assigns users to *principals* or *groups*, rather than to security roles. When you are developing a Java EE application, you don't need to know what categories of users have been defined for the realm in which the application will be run. In the Java EE platform, the security architecture provides a mechanism for mapping the security roles defined in the application to the users, principals, or groups defined in the runtime realm. The deployer will work with the security view provided by the application developer to implement this mapping.

One way to declare a mapping between a security role used in the application and one or more groups and/or principals defined for the applicable realm of the Enterprise Server is to use the `security-role-mapping` element in the runtime deployment descriptor (`sun-application.xml`, `sun-web.xml`, or `sun-ejb-jar.xml`.) This is the method to use when the role name defined in the application does not match the group or principal name defined for the Enterprise Server. An example of this role mapping can be found in Part VII, "Security," in *The Java EE 6 Tutorial, Volume II*.

In the tutorial, the role names used in the application are the same as the group names defined on the Enterprise Server. Under these circumstances, you can enable a default principal-to-role mapping on the Enterprise Server using the Admin Console. To enable the default principal-to-role-mapping, follow these steps:

1. Start the Enterprise Server, then the Admin Console.
2. Expand the Configuration node.
3. Select the Security node.

4. On the Security page, check the Enabled box beside Default Principal to Role Mapping.

## Specifying an Authentication Mechanism and Secure Connection

When method permissions are specified, basic, username-password, authentication will be invoked by the Enterprise Server.

If you would like to specify a different type of authentication, or to require a secure connection using SSL, you would specify this information in an application deployment descriptor. Using application deployment descriptors is discussed in Part VII, “Security,” in *The Java EE 6 Tutorial, Volume II*.

## Example: Securing an Enterprise Bean

This section discusses how to configure an enterprise bean for username-password authentication. When a bean that is constrained in this way is requested, the server requests a user name and password from the client and verifies that the user name and password are valid by comparing them against a database of authorized users on the Enterprise Server.

If the topic of authentication is new to you, please refer to the section titled “[Specifying an Authentication Mechanism](#)” on page 537.

This example demonstrates security by starting with the unsecured enterprise bean application, `cart`, which is found in the directory `tut-install/examples/ejb/cart/` and is discussed in “[The cart Example](#)” on page 307.

In general, the following steps are necessary to add username-password authentication to an existing application that contains an enterprise bean.

1. Create an application like the one in “[The cart Example](#)” on page 307.
2. If you have not already done so, complete the steps in “[Setting Up Your System for Running the Security Examples](#)” on page 556 to configure your system for running the tutorial applications.
3. Modify the source code for the enterprise bean, `CartBean.java`, to specify which roles are authorized to access which protected methods. This step is discussed in “[Annotating the Bean](#)” on page 512.
4. Build, package, and deploy the enterprise bean, then build and run the client application by following the steps in “[Building, Deploying, and Running the Secure Cart Example Using NetBeans IDE](#)” on page 514 or “[Building, Deploying, and Running the Secure Cart Example Using Ant](#)” on page 514.

## Annotating the Bean

The source code for the original `cart` application was modified as shown in the following code snippet (modifications in **bold**, method details are removed to save space). The resulting file can be found in the following location:

```
tut-install/examples/ejb/cart-secure/cart-secure-ejb/src/java/cart/
ejb/CartBean.java
```

The code snippet is as follows:

```
package cart.ejb;

import cart.util.BookException;
import cart.util.IdVerifier;
import java.util.ArrayList;
import java.util.List;
import javax.ejb.Remove;
import javax.ejb.Stateful;
import javax.annotation.security.RolesAllowed;

@Stateful
public class CartBean implements Cart {

    String customerName;
    String customerId;
    List<String> contents;

    public void initialize(String person) throws BookException {
        if (person == null) {
            throw new BookException("Null person not allowed.");
        } else {
            customerName = person;
        }

        customerId = "0";
        contents = new ArrayList<String>();
    }

    public void initialize(String person, String id) throws BookException {
        if (person == null) {
            throw new BookException("Null person not allowed.");
        } else {
            customerName = person;
        }

        IdVerifier idChecker = new IdVerifier();

        if (idChecker.validate(id)) {
            customerId = id;
        } else {
            throw new BookException("Invalid id: " + id);
        }
    }
}
```

```
        contents = new ArrayList<String>();
    }

    @RolesAllowed("TutorialUser")
    public void addBook(String title) {
        contents.add(title);
    }

    @RolesAllowed("TutorialUser")
    public void removeBook(String title) throws BookException {
        boolean result = contents.remove(title);
        if (result == false) {
            throw new BookException("\\" + title + "\" not in cart.");
        }
    }

    @RolesAllowed("TutorialUser")
    public List<String> getContents() {
        return contents;
    }

    @Remove()
    @RolesAllowed("TutorialUser")
    public void remove() {
        contents = null;
    }
}
```

The `@RolesAllowed` annotation is specified on methods for which you want to restrict access. In this example, only users in the role of `TutorialUser` will be allowed to add and remove books from the cart, and to list the contents of the cart. An `@RolesAllowed` annotation implicitly declares a role that will be referenced in the application; therefore, no `@DeclareRoles` annotation is required. The presence of the `@RolesAllowed` annotation also implicitly declares that authentication will be required for a user to access these methods. If no authentication method is specified in the deployment descriptor, the type of authentication will be `username-password` authentication.

## Building, Deploying, and Running the Secure Cart Example Using NetBeans IDE

Follow the instructions for building, deploying, and running the secure cart example by following the instructions in [“Building, Packaging, Deploying, and Running the cart Example” on page 312](#),

## Building, Deploying, and Running the Secure Cart Example Using Ant

Follow the instructions for building, deploying, and running the secure cart example by following the instructions in [“Building, Packaging, Deploying, and Running the cart Example” on page 312](#). Enter your user name and password when prompted to do so.

## Securing an Enterprise Bean Programmatically

Programmatic security is code that is embedded in a business method, is used to access a caller's identity programmatically, and uses this information to make security decisions within the method itself.

### Accessing an Enterprise Bean Caller's Security Context

In general, security management should be enforced by the container in a manner that is transparent to the enterprise beans' business methods. The security API described in this section should be used only in the less frequent situations in which the enterprise bean business methods need to access the security context information, such as when you want to restrict access to a particular time of day.

The `javax.ejb.EJBContext` interface provides two methods that allow the bean provider to access security information about the enterprise bean's caller.

- `java.security.Principal getCallerPrincipal();`

The purpose of the `getCallerPrincipal` method is to allow the enterprise bean methods to obtain the current caller principal's name. The methods might, for example, use the name as a key to information in a database.

The following code sample illustrates the use of the `getCallerPrincipal()` method:

```
@Stateless public class EmployeeServiceBean
    implements EmployeeService{
    @Resource SessionContext ctx;
    @PersistenceContext EntityManager em;

    public void changePhoneNumber(...) {
        ...
        // obtain the caller principal.
        callerPrincipal = ctx.getCallerPrincipal();

        // obtain the caller principal's name.
        callerKey = callerPrincipal.getName();

        // use callerKey as primary key to find EmployeeRecord
        EmployeeRecord myEmployeeRecord =
            em.find(EmployeeRecord.class, callerKey);

        // update phone number
        myEmployeeRecord.setPhoneNumber(...);

        ...
    }
}
```

In this example, the enterprise bean obtains the principal name of the current caller and uses it as the primary key to locate an `EmployeeRecord` entity. This example assumes that application has been deployed such that the current caller principal contains the primary key used for the identification of employees (for example, employee number).

- `boolean isCallerInRole(String roleName);`

The enterprise bean code can use the `isCallerInRole(String roleName)` method to allow the bean provider/application developer to code the security checks that cannot be easily defined using method permissions. Such a check might impose a role-based limit on a request, or it might depend on information stored in the database.

The enterprise bean code can use the `isCallerInRole(String roleName)` method to test whether the current caller has been assigned to a given security role. Security roles are defined by the bean provider or the application assembler, and are assigned to principals or principal groups that exist in the operational environment by the deployer.

The following code sample illustrates the use of the `isCallerInRole(String roleName)` method:

```
@Stateless public class PayrollBean implements Payroll {
    @Resource SessionContext ctx;

    public void updateEmployeeInfo(EmplInfo info) {

        oldInfo = ... read from database;

        // The salary field can be changed only by callers
        // who have the security role "payroll"
        if (info.salary != oldInfo.salary &&
            !ctx.isCallerInRole("payroll")) {
            throw new SecurityException(...);
        }
        ...
    }
    ...
}
```

You would use programmatic security in this way to dynamically control access to a method, for example, when you want to deny access except during a particular time of day. An example application that uses the `getCallerPrincipal` and `isCallerInRole` methods is described in [“Example: Using the `isCallerInRole` and `getCallerPrincipal` Methods” on page 516](#).

## Example: Using the `isCallerInRole` and `getCallerPrincipal` Methods

This example demonstrates how to use the `getCallerPrincipal()` and `isCallerInRole(String role)` methods with an enterprise bean. This example starts with a very simple EJB application, `converter`, and modifies the methods of the `ConverterBean` so that currency conversion will only occur when the requester is in the role of `TutorialUser`.

## Modifying ConverterBean

The source code for the original converter application was modified as shown in the following code snippet (modifications in **bold**) to add the `if . . else` clause that tests if the caller is in the role of `TutorialUser`. If the user is in the correct role, the currency conversion is computed and displayed. If the user is not in the correct role, the computation is not performed, and the application displays the result as `0`.

The code snippet is as follows:

```
package converter.secure.ejb;

import java.math.BigDecimal;
import javax.ejb.*;
import java.security.Principal;
import javax.annotation.Resource;
import javax.ejb.SessionContext;
import javax.annotation.security.DeclareRoles;
import javax.annotation.security.RolesAllowed;
@Stateless()
@DeclareRoles("TutorialUser")
public class ConverterBean{
    @Resource SessionContext ctx;
    private BigDecimal yenRate = new BigDecimal("96.0650");
    private BigDecimal euroRate = new BigDecimal("0.0078");

    @RolesAllowed("TutorialUser")
    public BigDecimal dollarToYen(BigDecimal dollars) {
        BigDecimal result = new BigDecimal("0.0");
        Principal callerPrincipal = ctx.getCallerPrincipal();
        if (ctx.isCallerInRole("TutorialUser")) {
            result = dollars.multiply(yenRate);
            return result.setScale(2, BigDecimal.ROUND_UP);
        }else{
            return result.setScale(2, BigDecimal.ROUND_UP);
        }
    }

    @RolesAllowed("TutorialUser")
    public BigDecimal yenToEuro(BigDecimal yen) {
        BigDecimal result = new BigDecimal("0.0");
        Principal callerPrincipal = ctx.getCallerPrincipal();
        if (ctx.isCallerInRole("TutorialUser")) {
            result = yen.multiply(euroRate);
            return result.setScale(2, BigDecimal.ROUND_UP);
        }else{
            return result.setScale(2, BigDecimal.ROUND_UP);
        }
    }
}
```

## Building, Deploying, and Running the Secure Converter Example Using NetBeans IDE

After you've made the changes to the enterprise bean, follow the instructions in [“Compiling, Packaging, and Running the converter Example”](#) on page 303.

## Building, Deploying, and Running the Secure Converter Example Using Ant

After you've made the changes to the enterprise bean, follow the instructions in [“Compiling, Packaging, and Running the converter Example”](#) on page 303.

## Troubleshooting the Secure Converter Application

**Problem:** The application displays zero values after authentication, as shown here:

```
apclient-command-common:
  [exec] $100.00 is 0.00 Yen.
  [exec] 0.00 Yen is 0.00 Euro.
```

**Solution:** Verify that the user name and password that you entered for authentication match a user name and password in the Enterprise Server, and that this user is assigned to the group named `TutorialUser`. User names and passwords are case-sensitive. Read [“Adding Users to the Enterprise Server”](#) on page 490 for more information on adding users to the `file` realm of the Enterprise Server.

## Propagating a Security Identity (Run-As)

You can specify whether a caller's security identity should be used for the execution of specified methods of an enterprise bean, or whether a specific run-as identity should be used.

[Figure 27-2](#) illustrates this concept.

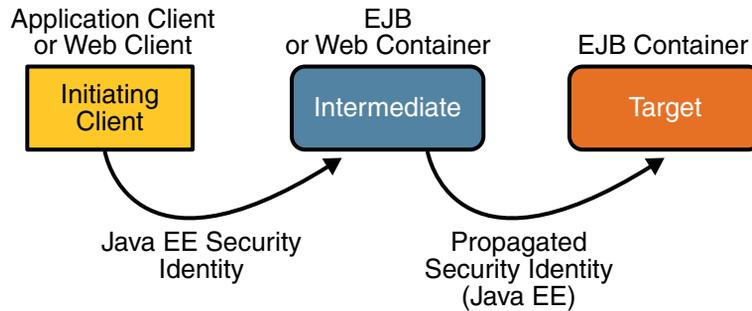


FIGURE 27-2 Security Identity Propagation

In this illustration, an application client is making a call to an enterprise bean method in one EJB container. This enterprise bean method, in turn, makes a call to an enterprise bean method in another container. The security identity during the first call is the identity of the caller. The security identity during the second call can be any of the following options:

- By default, the identity of the caller of the intermediate component is propagated to the target enterprise bean. This technique is used when the target container trusts the intermediate container.
- A *specific identity* is propagated to the target enterprise bean. This technique is used when the target container expects access using a specific identity.

To propagate an identity to the target enterprise bean, configure a run-as identity for the bean as discussed in [“Configuring a Component’s Propagated Security Identity” on page 519](#).

Establishing a run-as identity for an enterprise bean does not affect the identities of its callers, which are the identities tested for permission to access the methods of the enterprise bean. The run-as identity establishes the identity that the enterprise bean will use when it makes calls.

The run-as identity applies to the enterprise bean as a whole, including all the methods of the enterprise bean’s business interface, home interface, component interface, and web service endpoint interfaces, the message listener methods of a message-driven bean, the time-out callback method of an enterprise bean, and all internal methods of the bean that might be called in turn.

## Configuring a Component’s Propagated Security Identity

You can configure an enterprise bean’s run-as, or propagated, security identity using the `@RunAs` annotation. The `RunAs` annotation defines the role of the application during execution in a Java EE container. It can be specified on a class, allowing developers to execute an application under a particular role. The role *must* map to the user/group information in the container’s security realm. The `RunAs` annotation specifies the name of a security role as its parameter.

Here is some example code that demonstrates the use of the `RunAs` annotation.

```
@RunAs("Admin")
public class Calculator {
    //....
}
```

You will have to map the run-as role name to a given principal defined on the Enterprise Server if the given roles associate to more than one user principal. Mapping roles to principals is described in Part VII, “Security,” in *The Java EE 6 Tutorial, Volume II*.

## Trust between Containers

When an enterprise bean is designed so that either the original caller identity or a designated identity is used to call a target bean, the target bean will receive the propagated identity only; it will not receive any authentication data.

There is no way for the target container to authenticate the propagated security identity. However, because the security identity is used in authorization checks (for example, method permissions or with the `isCallerInRole()` method), it is vitally important that the security identity be authentic. Because there is no authentication data available to authenticate the propagated identity, the target must trust that the calling container has propagated an authenticated security identity.

By default, the Enterprise Server is configured to trust identities that are propagated from different containers. Therefore, there are no special steps that you need to take to set up a trust relationship.

## Deploying Secure Enterprise Beans

The deployer is responsible for ensuring that an assembled application is secure after it has been deployed in the target operational environment. If a security view has been provided to the deployer through the use of security annotations and/or a deployment descriptor, the security view is mapped to the mechanisms and policies used by the security domain in the target operational environment, which in this case is the Enterprise Server. If no security view is provided, the deployer must set up the appropriate security policy for the enterprise bean application.

Deployment information is specific to a web or application server. Please read the [Sun GlassFish Enterprise Server v3 Application Deployment Guide](#) for more information on deploying enterprise beans.

## Accepting Unauthenticated Users

Web applications may accept unauthenticated web clients and allow these clients to make calls to the EJB container. The EJB specification requires a security credential for accessing EJB methods. Typically, the credential will be that of a generic unauthenticated user. The way you specify this credential is implementation-specific.

In the Enterprise Server, you must specify the name and password that an unauthenticated user will use to log in by modifying the Enterprise Server using the Admin Console:

1. Start the Enterprise Server, then the Admin Console.
2. Expand the Configuration node.
3. Select the Security node.
4. On the Security page, set the Default Principal and Default Principal Password values.

## Securing Application Clients

The Java EE authentication requirements for application clients are the same as for other Java EE components, and the same authentication techniques can be used as for other Java EE application components.

No authentication is necessary when accessing unprotected web resources. When accessing protected web resources, the usual varieties of authentication can be used, namely HTTP basic authentication, SSL client authentication, or HTTP login form authentication. These authentication methods are discussed in [“Specifying an Authentication Mechanism” on page 537](#).

Authentication is required when accessing protected enterprise beans. The authentication mechanisms for enterprise beans are discussed in [“Securing Enterprise Beans” on page 504](#). Lazy authentication can be used.

An application client makes use of an authentication service provided by the application client container for authenticating its users. The container’s service can be integrated with the native platform’s authentication system, so that a single sign-on capability is employed. The container can authenticate the user when the application is started, or it can use lazy authentication, authenticating the user when a protected resource is accessed.

An application client can provide a class to gather authentication data. If so, the `javax.security.auth.callback.CallbackHandler` interface must be implemented, and the class name must be specified in its deployment descriptor. The application’s callback handler must fully support `Callback` objects specified in the `javax.security.auth.callback` package. Gathering authentication data in this way is discussed in the next section, [“Using Login Modules” on page 521](#).

## Using Login Modules

An application client can use the Java Authentication and Authorization Service (JAAS) to create *login modules* for authentication. A JAAS-based application implements the `javax.security.auth.callback.CallbackHandler` interface so that it can interact with users to enter specific authentication data, such as user names or passwords, or to display error and warning messages.

Applications implement the `CallbackHandler` interface and pass it to the login context, which forwards it directly to the underlying login modules. A login module uses the callback handler both to gather input (such as a password or smart card PIN) from users and to supply information (such as status information) to users. Because the application specifies the callback handler, an underlying login module can remain independent of the various ways that applications interact with users.

For example, the implementation of a callback handler for a GUI application might display a window to solicit user input. Or the implementation of a callback handler for a command-line tool might simply prompt the user for input directly from the command line.

The login module passes an array of appropriate callbacks to the callback handler's `handle` method (for example, a `NameCallback` for the user name and a `PasswordCallback` for the password); the callback handler performs the requested user interaction and sets appropriate values in the callbacks. For example, to process a `NameCallback`, the `CallbackHandler` might prompt for a name, retrieve the value from the user, and call the `setName` method of the `NameCallback` to store the name.

For more information on using JAAS for login modules for authentication, refer to the following sources:

- *Java Authentication and Authorization Service (JAAS) in Java Platform, Standard Edition*
- *Java Authentication and Authorization Service (JAAS) Reference Guide*
- *Java Authentication and Authorization Service (JAAS): LoginModule Developer's Guide*

Links to this information are provided in [“Further Information about Security”](#) on page 500.

## Using Programmatic Login

Programmatic login enables the client code to supply user credentials. If you are using an EJB client, you can use the `com.sun.appserv.security.ProgrammaticLogin` class with their convenient `login` and `logout` methods.

Programmatic login is specific to a server. Information on programmatic login for the Enterprise Server is included in the [Sun GlassFish Enterprise Server v3 Application Development Guide](#) or the documentation for the server you are using.

# Securing Enterprise Information Systems (EIS) Applications

In EIS applications, components request a connection to an EIS resource. As part of this connection, the EIS can require a sign-on for the requester to access the resource. The application component provider has two choices for the design of the EIS sign-on:

- Container-managed sign-on

In the container-managed sign-on approach, the application component lets the container take the responsibility of configuring and managing the EIS sign-on. The container determines the user name and password for establishing a connection to an EIS instance. For more information, read [“Container-Managed Sign-On” on page 523](#).

- Component-managed sign-on

In the component-managed sign-on approach, the application component code manages EIS sign-on by including code that performs the sign-on process to an EIS. For more information, read [“Component-Managed Sign-On” on page 523](#).

You can also configure security for resource adapters. Read [“Configuring Resource Adapter Security” on page 524](#) for more information.

## Container-Managed Sign-On

In container-managed sign-on, an application component does not have to pass any sign-on security information to the `getConnection()` method. The security information is supplied by the container, as shown in the following example.

```
// Business method in an application component
Context initctx = new InitialContext();
// Perform JNDI lookup to obtain a connection factory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory)initctx.lookup(
        "java:comp/env/eis/MainframeCxFactory");
// Invoke factory to obtain a connection. The security
// information is not passed in the getConnection method
javax.resource.cci.Connection cx = cxf.getConnection();
...
```

## Component-Managed Sign-On

In component-managed sign-on, an application component is responsible for passing the needed sign-on security information to the resource to the `getConnection` method. For example, security information might be a user name and password, as shown here:

```
// Method in an application component
Context initctx = new InitialContext();

// Perform JNDI lookup to obtain a connection factory
javax.resource.cci.ConnectionFactory cxf =
    (javax.resource.cci.ConnectionFactory)initctx.lookup(
        "java:comp/env/eis/MainframeCxFactory");

// Get a new ConnectionSpec
com.myeis.ConnectionSpecImpl properties = //..

// Invoke factory to obtain a connection
properties.setUserName("...");
properties.setPassword("...");
javax.resource.cci.Connection cx =
    cxf.getConnection(properties);
...

```

## Configuring Resource Adapter Security

A resource adapter is a system-level software component that typically implements network connectivity to an external resource manager. A resource adapter can extend the functionality of the Java EE platform either by implementing one of the Java EE standard service APIs (such as a JDBC driver), or by defining and implementing a resource adapter for a connector to an external application system. Resource adapters can also provide services that are entirely local, perhaps interacting with native resources. Resource adapters interface with the Java EE platform through the Java EE service provider interfaces (Java EE SPI). A resource adapter that uses the Java EE SPIs to attach to the Java EE platform will be able to work with all Java EE products.

To configure the security settings for a resource adapter, you need to edit the resource adapter descriptor file, `ra.xml`. Here is an example of the part of an `ra.xml` file that configures the following security properties for a resource adapter:

```
<authentication-mechanism>
  <authentication-mechanism-type>BasicPassword</authentication-mechanism-type>
  <credential-interface>
    javax.resource.spi.security.PasswordCredential
  </credential-interface>
</authentication-mechanism>
<reauthentication-support>>false</reauthentication-support>

```

You can find out more about the options for configuring resource adapter security by reviewing *as-install/lib/dtds/connector\_1\_0.dtd*. You can configure the following elements in the resource adapter deployment descriptor file:

- Authentication mechanisms

Use the `authentication-mechanism` element to specify an authentication mechanism supported by the resource adapter. This support is for the resource adapter and not for the underlying EIS instance.

There are two supported mechanism types:

- `BasicPassword`: This mechanism supports the interface `javax.resource.spi.security.PasswordCredential`.
- `Kerbv5`: This mechanism supports the interface `javax.resource.spi.security.GenericCredential`. The Enterprise Server does not currently support this mechanism type.

- Reauthentication support

Use the `reauthentication-support` element to specify whether the resource adapter implementation supports re-authentication of existing `Managed-Connection` instances. Options are `true` or `false`.

- Security permissions

Use the `security-permission` element to specify a security permission that is required by the resource adapter code. Support for security permissions is optional and is not supported in the current release of the Enterprise Server. You can, however, manually update the `server.policy` file to add the relevant permissions for the resource adapter, as described in the *Developing and Deploying Applications* section of the *Sun GlassFish Enterprise Server v3 Application Development Guide*.

The security permissions listed in the deployment descriptor are ones that are different from those required by the default permission set as specified in the connector specification.

Refer to the following URL for more information on Sun's implementation of the security permission specification: <http://java.sun.com/javase/6/docs/technotes/guides/security/PolicyFiles.html#FileSyntax>.

In addition to specifying resource adapter security in the `ra.xml` file, you can create a security map for a connector connection pool to map an application principal or a user group to a back end EIS principal. The security map is usually used in situations where one or more EIS back end principals are used to execute operations (on the EIS) initiated by various principals or user groups in the application. You can find out more about security maps in the *Configuring Security* chapter section of the *Sun GlassFish Enterprise Server v3 Administration Guide*.

## Mapping an Application Principal to EIS Principals

When using the Enterprise Server, you can 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 Enterprise 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 Enterprise 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.

To work with security maps, use the Admin Console. From the Admin Console, follow these steps to get to the security maps page:

1. Expand the Resources node.
2. Expand the Connectors node.
3. Select the Work Security Maps node.
4. Click New to create a new work security map for a resource adapter.
5. Enter a name by which you will refer to the security map, as well as the other required information. Click the Help button for more information on the individual options.

## Getting Started Securing Web Applications

---

A *web application* is an application that is accessed using a web browser over a network such as the Internet or a company's intranet. As discussed in [Chapter 27, “Getting Started Securing Enterprise Applications,”](#) the Java™ EE platform uses a distributed multi-tiered application model. As discussed in [“Distributed Multitiered Applications” on page 35,](#) web applications run in the web-tier.

Web applications contain resources that can be accessed by many users. These resources often traverse unprotected, open networks, such as the Internet. In such an environment, a substantial number of web applications will require some type of security. The ways to implement security for Java EE web applications are discussed in a general way in [“Securing Containers” on page 482.](#) This chapter provides more detail and a few examples that explore these security services as they relate to web components.

Both types of web applications can be secured using the same security model:

- **Presentation-oriented:** A presentation-oriented web application generates interactive web pages containing various types of markup language such as HTML, XHTML, and XML, and generates dynamic content in response to requests. The technologies that are discussed in this chapter and that are considered presentation-oriented web applications include Java Servlets and JavaServer™ Faces technology. You can read more about web applications in [Chapter 3, “Getting Started with Web Applications.”](#)
- **Service-oriented:** A service-oriented web application implements the endpoint of a web service. Presentation-oriented applications are often clients of service-oriented web applications. The technologies that are discussed in the chapter and that are considered service-oriented web applications include the Java API for XML-Based Web Services (JAX-WS) and the Java API for RESTful Web Services (JAX-RS). You can read more about web services in [Chapter 12, “Introduction to Web Services”](#)

Securing applications and their clients in the business tier and the EIS tier is discussed in [Chapter 27, “Getting Started Securing Enterprise Applications.”](#)

The following topics are included in this chapter:

- “Overview of Web Application Security” on page 528
- “Using Deployment Descriptors to Secure Web Applications” on page 530
- “Using Programmatic Security with Web Applications” on page 548
- “Using Message Security with Web Applications” on page 555
- “Examples: Securing Web Applications” on page 555

## Overview of Web Application Security

In the Java EE platform, *web components* provide the dynamic extension capabilities for a web server. Web components are either Java servlets, JSF pages, or web service endpoints. The interaction between a web client and a web application is illustrated in Figure 28–1.

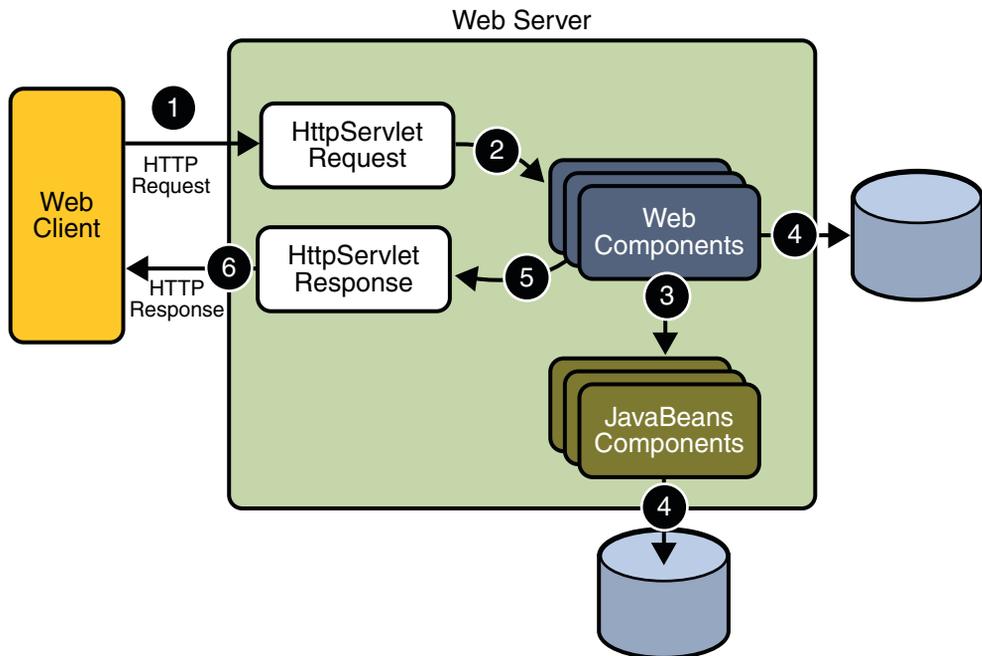


FIGURE 28–1 Java Web Application Request Handling

Web components are supported by the services of a runtime platform called a *web container*. A web container provides services such as request dispatching, security, concurrency, and life-cycle management. Certain aspects of web application security can be configured when the application is installed, or *deployed*, to the web container. Annotations and/or deployment descriptors are used to relay information to the deployer about security and other aspects of the

application. Specifying this information in annotations or in the deployment descriptor helps the deployer set up the appropriate security policy for the web application. Any values explicitly specified in the deployment descriptor override any values specified in annotations.

Security for Java EE web applications can be implemented in the following ways:

- **Declarative security**

Declarative security can be implemented using either an application's deployment descriptor or using metadata annotations. *Metadata annotations* (or simply, *annotations*) are used to specify information about security within a class file. An application *deployment descriptor* is an XML file that is external to the application and which expresses an application's security structure, including security roles, access control, and authentication requirements. When an application is deployed, security information specified using annotations can be overridden by the application deployment descriptor.

Declarative security is described in [“Using Deployment Descriptors to Secure Web Applications” on page 530](#).

- **Programmatic security**

*Programmatic security* is embedded in an application and is used to make security decisions. Programmatic security is useful when declarative security alone is not sufficient to express the security model of an application. Declarative security alone may not be sufficient in cases where conditional login in a particular work flow, instead of for all cases, is required in the middle of an application.

New in Java EE 6 and Servlet specification 3.0 are the `authenticate`, `login`, and `logout`, methods of the `HttpServletRequest` interface. With the addition of the `authenticate`, `login`, and `logout` methods to the Servlet specification, an application deployment descriptor is no longer required for web applications, but may still be used to further specify security requirements beyond the basic default values.

Programmatic security is discussed in [“Using Programmatic Security with Web Applications” on page 548](#)

- **Message Security**

*Message security* works with web services and incorporates security features, such as digital signatures and encryption, into the header of a SOAP message, working in the application layer, ensuring end-to-end security. Message security is not a component of Java EE 6, but is included here for informational purposes only.

Message security is discussed in [“Using Message Security with Web Applications” on page 555](#).

Some of the material in this chapter builds on material presented earlier in this tutorial. In particular, this chapter assumes that you are familiar with the information in the following chapters:

- Chapter 3, “Getting Started with Web Applications”
- Chapter 11, “Java Servlet Technology”

- Chapter 4, “JavaServer Faces Technology”
- Chapter 12, “Introduction to Web Services”
- Chapter 13, “Building Web Services with JAX-WS”
- Chapter 14, “Building RESTful Web Services with JAX-RS and Jersey”
- Chapter 26, “Introduction to Security in the Java EE Platform”
- Chapter 27, “Getting Started Securing Enterprise Applications”

## Using Deployment Descriptors to Secure Web Applications

Web applications are created by application developers who give, sell, or otherwise transfer the application to an application deployer for installation into a runtime environment. Application developers communicate how the security is to be set up for the deployed application by the use of the *deployment descriptor* mechanism or by the use of *annotations*. When this information is passed on to the deployer, the deployer uses this information to define method permissions for security roles, set up user authentication, and set up the appropriate transport mechanism. If you don't define security requirements, the deployer will have to determine the security requirements independently.

Many elements necessary for security in a web application cannot, as yet, be specified as annotations for all types of web applications. Therefore, in this volume of the tutorial, we describe how to secure web applications using deployment descriptors only. For information on using annotations to secure web applications, refer to the following sources:

- The [Java Servlet Specification, Version 3.0](#), for information on the `@ServletSecurity` annotation
- Part VII, “Security,” in *The Java EE 6 Tutorial, Volume II*

## Introduction to Web Application Deployment Descriptors

The web application deployment descriptor file does pretty much what its name says it does: it describes how the web application should be deployed. The web application deployment descriptor describes a lot more about a web application than just its security information, but this chapter only discusses the elements of the application deployment descriptor that relate to security.

For web applications written using the Java programming language, the web application deployment descriptor is written using the EXtensible Markup Language (XML) syntax. The web application deployment descriptor is named `web.xml`, and, when included with a web application, it must reside in a `WEB-INF` subdirectory at the web application root. The contents of this file direct a deployment tool to deploy a module or application with the specified security settings, and describes other specific configuration requirements and/or container options.

The following XML code is an example of the elements in a deployment descriptor that apply specifically to declaring security for web applications or for resources within web applications. This example comes from *An Example of Security*, from the [Java Servlet Specification, Version 3.0](#).

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<web-app xmlns="http://java.sun.com/xml/ns/j2ee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/j2ee
  http://java.sun.com/xml/ns/j2ee/web-app_2_5.xsd"
  version="2.5">

  <display-name>A Secure Application</display-name>
  <servlet>
    <servlet-name>catalog</servlet-name>
    <servlet-class>com.mycorp.CatalogServlet</servlet-class>
    <init-param>
      <param-name>catalog</param-name>
      <param-value>Spring</param-value>
    </init-param>

    <!-- Defining Security Roles -->

    <security-role-ref>
      <role-name>MGR</role-name>
      <!-- role name used in code -->
      <role-link>manager</role-link>
    </security-role-ref>
  </servlet>
  <security-role>
    <role-name>manager</role-name>
  </security-role>
  <servlet-mapping>
    <servlet-name>catalog</servlet-name>
    <url-pattern>/catalog/*</url-pattern>
  </servlet-mapping>

  <!-- Defining A Security Constraint -->

  <security-constraint>
    <!-- Specifying the Resources to be Protected -->

    <web-resource-collection>
      <web-resource-name>SalesInfo</web-resource-name>
      <url-pattern>/salesinfo/*</url-pattern>
      <http-method>GET</http-method>
      <http-method>POST</http-method>
    </web-resource-collection>
```

```
        <!-- Specifying which Users Can Access Protected Resources -->

        <auth-constraint>
            <role-name>manager</role-name>
        </auth-constraint>
        <!-- Specifying Secure Transport using SSL -->

        <user-data-constraint>
            <transport-guarantee>CONFIDENTIAL </transport-guarantee>
        </user-data-constraint>
    </security-constraint>

    <!-- Specifying an Authentication Method -->
    <login-config>
        <auth-method>BASIC</auth-method>
        <realm-name>file</realm-name>
    </login-config>
</web-app>
```

Even if you are simply using the deployment descriptor to specify security, there are some structural elements that must be included in this file in order for it to work properly. For example, the `<security-constraint>` element is a sub-element of the `<web-app>` element, so the `<web-app>` element must always be included, and it must indicate the version of the web application schema (2.4 or 2.5) it is using. The elements that are specified within the deployment descriptor must comply with the rules for processing that version of the deployment descriptor. Version 3.0 of the Java Servlet Specification, which can be downloaded at <http://jcp.org/en/jsr/detail?id=315>, contains more information regarding the structure of deployment descriptors.

XML files are hierarchical. The elements must be specified in a particular order within the deployment descriptor, between elements that are its parent. To visually see an example of how the deployment descriptor elements are nested within their parent elements, refer to the elements within the `<security-constraint>` element above, which is itself nested within `<web-app>` elements. For this example, the lines have been indented to emphasize the nesting aspect of the file, but the file itself ignores the formatting and relies only on the elements and their content for its processing. Information about the application is specified as a value between the opening (`<element-name>`) and closing (`</element-name>`) elements. For example, between the opening `<transport-guarantee>` element and the closing `</transport-guarantee>` element, there is the value `CONFIDENTIAL`, which describes which type of transport guarantee should be used for this application.

The following sections describe each of the security elements of a deployment descriptor in more detail, listing all of the options available for each element:

- “Specifying Security Constraints” on page 533
- “Specifying an Authentication Mechanism” on page 537
- “Working with Security Roles” on page 544

Some of the elements of web application security must be addressed in server configuration files rather than in the deployment descriptor or annotations for the web application. Configuring security on the Enterprise Server is discussed in the following sections and books:

- “Securing the Enterprise Server” on page 484 (in this book)
- “Managing Users and Groups on the Enterprise Server” on page 489 (in this book)
- “Installing and Configuring SSL Support” on page 494 (in this book)
- “Deploying Secure Enterprise Beans” on page 520 (in this book)
- *Sun GlassFish Enterprise Server v3 Administration Guide*
- *Sun GlassFish Enterprise Server v3 Application Development Guide*

## Specifying Security Constraints

A *security constraint* is used to define the access privileges to a collection of resources using their URL mapping. The following elements can be part of a security constraint:

- *Web resource collection* (web-resource-collection)

A web resource collection is a list of URL patterns (the part of a URL *after* the host name and port which you want to constrain) and HTTP operations (the methods within the files that match the URL pattern which you want to constrain (in this example, POST and GET)) that describe a set of resources to be protected. Web resource collections are discussed in “[Specifying a Web Resource Collection](#)” on page 534.
- *Authorization constraint* (auth-constraint)

*Authorization constraints* establish a requirement for authentication and name the roles authorized to perform the constrained requests. For more information about authorization constraints, read “[Specifying an Authentication Mechanism](#)” on page 537.
- *User data constraint* (TutorialUser-data-constraint)

*User data constraints* establish a requirement that the constrained requests be received over a protected transport layer connection. This guarantees how the data will be transported between client and server. The strength of the required protection is defined by the value of the transport guarantee. A transport guarantee of INTEGRAL is used to establish a requirement for content integrity and a transport guarantee of CONFIDENTIAL is used to establish a requirement for confidentiality. The transport guarantee of NONE indicates that the container must accept the constrained requests when received on any connection including an unprotected one. User data constraints are discussed in “[Specifying a Secure Connection](#)” on page 535.

## Specifying a Web Resource Collection

A web resource collection consists of the following sub-elements:

- `web-resource-name` is the name you use for this resource. Its use is optional, but it is useful to describe the resources being protected as applications get more complex.
- `url-pattern` is used to list the request URI to be protected.

Many applications feature unprotected web content, which any caller can access without authentication. In the web tier, you provide unrestricted access simply by not configuring a security constraint for that particular request URI. It is common to have some unprotected resources and some protected resources. In this case, you will define security constraints and a login method, but they will not be used to control access to the unprotected resources. Users won't be asked to log in until the first time they enter a protected request URI.

The Java Servlet specification defines the request URI as the part of a URL *after* the host name and port. For example, let's say you have an e-commerce site with a browseable catalog that you would want anyone to be able to access, and a shopping cart area for customers only. You could set up the paths for your web application so that the pattern `/cart/*` is protected but nothing else is protected. Assuming that the application is installed at context path `/myapp`, the following are true:

- `http://localhost:8080/myapp/index.jsp` is *not* protected.
- `http://localhost:8080/myapp/cart/index.jsp` is protected.

A user will not be prompted to log in the first time that user accesses a resource in the `cart/` subdirectory.

- `http-method` or `http-method-omission` is used to specify which methods should be protected or which methods should be omitted from protection. An HTTP method is protected by a `web-resource-collection` when no HTTP methods are named in the collection (which means all are protected), or the collection specifically names the HTTP method in a contained `http-method` element, or the collection contains one or more `http-method-omission` elements, none of which names the HTTP method.

## Specifying an Authorization Constraint

An authorization constraint (`auth-constraint`) contains the `role-name` element. You can use as many `role-name` elements as needed here.

An authorization constraint establishes a requirement for authentication and names the roles authorized to access the URL patterns and HTTP methods declared by this security constraint. If there is no authorization constraint, the container must accept the request without requiring user authentication. If there is an authorization constraint, but no roles are specified within it, the container will not allow access to constrained requests under any circumstances. The role name(s) specified here must either correspond to the role name of one of the `<security-role>` elements defined for this web application, or be the specially reserved role name `*`, which is a compact syntax for indicating all roles in the web application. Role names are case sensitive.

The roles defined for the application must be mapped to users and groups defined on the server. For more information about security roles, read [“Working with Security Roles” on page 544](#).

## Specifying a Secure Connection

A user data constraint (`<TutorialUser-data-constraint>` in the deployment descriptor) contains the `<transport-guarantee>` element. A user data constraint can be used to require that a protected transport layer connection such as HTTPS (HTTP over SSL) be used for all constrained URL patterns and HTTP methods specified in the security constraint. The choices for transport guarantee include `CONFIDENTIAL`, `INTEGRAL`, or `NONE`. If you specify `CONFIDENTIAL` or `INTEGRAL` as a security constraint, it generally means that the use of SSL is required, and that type of security constraint applies to all requests that match the URL patterns in the web resource collection and not just to the login dialog box.

The strength of the required protection is defined by the value of the transport guarantee. Specify `CONFIDENTIAL` when the application requires that data be transmitted so as to prevent other entities from observing the contents of the transmission. Specify `INTEGRAL` when the application requires that the data be sent between client and server in such a way that it cannot be changed in transit. Specify `NONE` to indicate that the container must accept the constrained requests on any connection, including an unprotected one.

The user data constraint is handy to use in conjunction with basic and form-based user authentication. When the login authentication method is set to `BASIC` or `FORM`, passwords are not protected, meaning that passwords sent between a client and a server on an unprotected session can be viewed and intercepted by third parties. Using a user data constraint with the user authentication mechanism can alleviate this concern. Configuring a user authentication mechanism is described in [“Specifying an Authentication Mechanism” on page 537](#).

To guarantee that data is transported over a secure connection, ensure that SSL support is configured for your server. If your server is the Sun Java System Enterprise Server, SSL support is already configured. If you are using another server, consult the documentation for that server for information on setting up SSL support. More information on configuring SSL support on the Enterprise Server can be found in [“Establishing a Secure Connection Using SSL” on page 494](#) and in the *Sun GlassFish Enterprise Server v3 Administration Guide*.

---

**Note** – Good Security Practice: If you are using sessions, after you switch to SSL you should never accept any further requests for that session that are non-SSL. For example, a shopping site might not use SSL until the checkout page, and then it might switch to using SSL to accept your card number. After switching to SSL, you should stop listening to non-SSL requests for this session. The reason for this practice is that the session ID itself was not encrypted on the earlier communications. This is not so bad when you’re only doing your shopping, but after the credit card information is stored in the session, you don’t want a bad guy trying to fake the purchase transaction against your credit card. This practice could be easily implemented using a filter.

---

## Specifying Separate Security Constraints for Different Resources

You can create a separate security constraint for different resources within your application. For example, you could allow users with the role of PARTNER access to the GET and POST methods of all resources with the URL pattern /acme/wholesale/\*, and allow users with the role of CLIENT access to the GET and POST methods of all resources with the URL pattern /acme/retail/\*. An example of a deployment descriptor that would demonstrate this functionality is the following:

```
<!-- SECURITY CONSTRAINT #1 -->
<security-constraint>
  <web-resource-collection>
    <web-resource-name>wholesale</web-resource-name>
    <url-pattern>/acme/wholesale/*</url-pattern>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>PARTNER</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>

<!-- SECURITY CONSTRAINT #2 -->
<security-constraint>
  <web-resource-collection>
    <web-resource-name>retail</web-resource-name>
    <url-pattern>/acme/retail/*</url-pattern>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>CLIENT</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>
```

When the same `url-pattern` and `http-method` occur in multiple security constraints, the constraints on the pattern and method are defined by combining the individual constraints, which could result in unintentional denial of access. The *Java Servlet 3.0 Specification* (downloadable from <http://jcp.org/en/jsr/detail?id=315>) gives more detail and an example that illustrates the combination of constraints and how the declarations will be interpreted.

## Specifying an Authentication Mechanism

The *login configuration* element is separate from the *security-constraint* element, as there can be multiple security constraints applying to multiple resources, but the same authentication method will apply to all constrained resources in an application. The *login-config* element is used to specify the user authentication method to be used for access to web content, the realm in which the user will be authenticated (in the case of basic authentication), and, in the case of form-based login, additional attributes. When specified, the user must be authenticated before access to any resource that is constrained by a security constraint will be granted.

The sub-element *auth-method* configures the authentication mechanism for the web application. The element content must be either NONE, BASIC, DIGEST, FORM, or CLIENT-CERT. The *realm-name* element indicates the realm name to use when the basic authentication scheme is chosen for the web application. The *form-login-config* element specifies the login and error pages that should be used when FORM based login is specified.

---

**Note** – Another way to specify form-based authentication is to use the *authenticate*, *login*, and *logout* methods of *HttpServletRequest*, as discussed in [“Authenticating Users Programmatically” on page 549](#).

---

When you try to access a web resource that is constrained by a *security-constraint* element, the web container activates the authentication mechanism that has been configured for that resource. The authentication mechanism you choose specifies how the user is prompted to login. If the *<login-config>* element is present, and the *<auth-method>* element contains a value other than NONE, the user must be authenticated before it can access any resource that is constrained by the use of a *security-constraint* element in the same deployment descriptor (read [“Specifying Security Constraints” on page 533](#) for more information on security constraints). If you do not specify an authentication mechanism, authentication of the user is not required.

Before you can authenticate a user, you must have a database of user names, passwords, and roles configured on your web or application server. For information on setting up the user database, refer to [“Managing Users and Groups on the Enterprise Server” on page 489](#) in this tutorial or the *Sun GlassFish Enterprise Server v3 Administration Guide*.

The choices for authentication mechanisms are discussed further in the following sections:

- [“HTTP Basic Authentication” on page 538](#)
- [“Form-Based Authentication” on page 539](#)
- [“HTTPS Client Authentication” on page 541](#)
- [“Digest Authentication” on page 544](#)

## HTTP Basic Authentication

Specifying *HTTP Basic Authentication* requires that the server request a user name and password from the web client and verify that the user name and password are valid by comparing them against a database of authorized users in the specified or default realm.

When basic authentication is declared, the following actions occur:

1. A client requests access to a protected resource.
2. The web server returns a dialog box that requests the user name and password.
3. The client submits the user name and password to the server.
4. The server authenticates the user in the specified realm and, if successful, returns the requested resource.

Figure 28–2 shows what happens when you specify HTTP basic authentication.

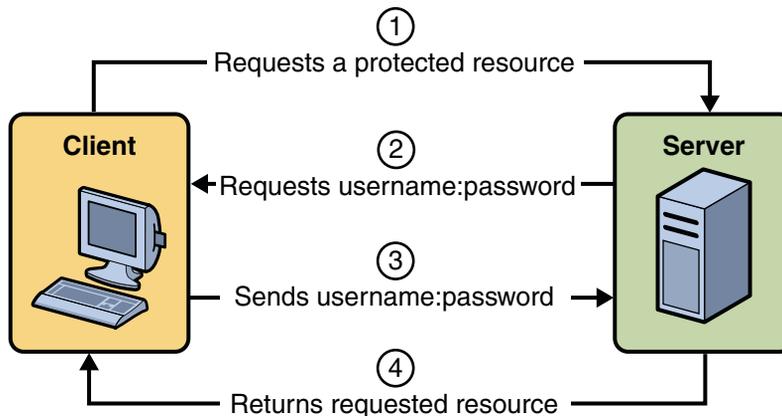


FIGURE 28-2 HTTP Basic Authentication

The following example shows how to specify basic authentication in your deployment descriptor:

```
<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>file</realm-name>
</login-config>
```

HTTP basic authentication is not a secure authentication mechanism. Basic authentication sends user names and passwords over the Internet as text that is Base64 encoded, and the target server is not authenticated. This form of authentication can expose user names and passwords. If someone can intercept the transmission, the user name and password information can easily be decoded. However, when a secure transport mechanism, such as SSL, or security at the

network level, such as the IPSEC protocol or VPN strategies, is used in conjunction with basic authentication, some of these concerns can be alleviated. To specify a secure transport mechanism, use the elements described in [“Specifying a Secure Connection” on page 535](#).

[“Example: Basic Authentication with JAX-WS” on page 562](#) is an example application that uses HTTP basic authentication in a JAX-WS service.

## Form-Based Authentication

Form-based authentication allows the developer to control the look and feel of the login authentication screens by customizing the login screen and error pages that an HTTP browser presents to the end user. When form-based authentication is declared, the following actions occur:

1. A client requests access to a protected resource.
2. If the client is unauthenticated, the server redirects the client to a login page.
3. The client submits the login form to the server.
4. The server attempts to authenticate the user.
  - a. If authentication succeeds, the authenticated user’s principal is checked to ensure it is in a role that is authorized to access the resource. If the user is authorized, the server redirects the client to the resource using the stored URL path.
  - b. If authentication fails, the client is forwarded or redirected to an error page.

[Figure 28–3](#) shows what happens when you specify form-based authentication.

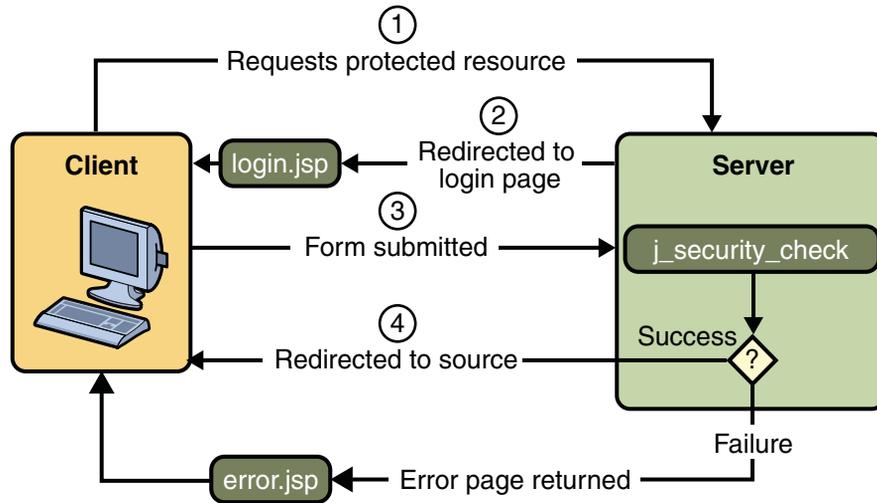


FIGURE 28-3 Form-Based Authentication

The following example shows how to declare form-based authentication in your deployment descriptor:

```

<login-config>
  <auth-method>FORM</auth-method>
  <realm-name>file</realm-name>
  <form-login-config>
    <form-login-page>/logon.jsp</form-login-page>
    <form-error-page>/logonError.jsp</form-error-page>
  </form-login-config>
</login-config>
  
```

The login and error page locations are specified relative to the location of the deployment descriptor. Examples of login and error pages are shown in [“Creating the Login Form and the Error Page”](#) on page 566.

Form-based authentication is not particularly secure. In form-based authentication, the content of the user dialog box is sent as plain text, and the target server is not authenticated. This form of authentication can expose your user names and passwords unless all connections are over SSL. If someone can intercept the transmission, the user name and password information can easily be decoded. However, when a secure transport mechanism, such as SSL, or security at the network level, such as the IPSEC protocol or VPN strategies, is used in conjunction with form-based authentication, some of these concerns can be alleviated. To add a protected transport in your application, use the elements described in [“Specifying a Secure Connection”](#) on page 535.

The section [“Example: Form-Based Authentication with a Servlet”](#) on page 566 is an example application that uses form-based authentication.

## Using Login Forms

When creating a form-based login, be sure to maintain sessions using cookies or SSL session information.

As shown in [“Form-Based Authentication” on page 539](#), for authentication to proceed appropriately, the action of the login form must always be `j_security_check`. This restriction is made so that the login form will work no matter which resource it is for, and to avoid requiring the server to specify the action field of the outbound form. The following code snippet shows how the form should be coded into the HTML page:

```
<form method="POST" action="j_security_check">
<input type="text" name="j_username">
<input type="password" name="j_password">
</form>
```

## HTTPS Client Authentication

HTTPS Client Authentication requires the client to possess a Public Key Certificate (PKC). If you specify *client authentication*, the web server will authenticate the client using the client's public key certificate.

HTTPS Client Authentication is a more secure method of authentication than either basic or form-based authentication. It uses HTTP over SSL (HTTPS), in which the server authenticates the client using the client's Public Key Certificate (PKC). Secure Sockets Layer (SSL) technology provides data encryption, server authentication, message integrity, and optional client authentication for a TCP/IP connection. You can think of a public key certificate as the digital equivalent of a passport. It is issued by a trusted organization, which is called a certificate authority (CA), and provides identification for the bearer.

Before using HTTP Client Authentication, you must make sure that the following actions have been completed:

- Make sure the client has a valid Public Key Certificate. For more information on creating and using public key certificates, read [“Working with Digital Certificates” on page 497](#).
- Make sure that SSL support is configured for your server. If your server is the Sun GlassFishEnterprise Server v3, SSL support is already configured. If you are using another server, consult the documentation for that server for information on setting up SSL support. More information on configuring SSL support on the application server can be found in [“Establishing a Secure Connection Using SSL” on page 494](#) and the *Sun GlassFish Enterprise Server v3 Administration Guide*.

The following example shows how to declare HTTPS client authentication in your deployment descriptor:

```
<login-config>  
  <auth-method>CLIENT-CERT</auth-method>  
</login-config>
```

An example demonstrating HTTPS client authentication may be available in Part VII, “Security,” in *The Java EE 6 Tutorial, Volume II*.

## Mutual Authentication

With *mutual authentication*, the server and the client authenticate one another. There are two types of mutual authentication:

- Certificate-based mutual authentication (see [Figure 28–4](#))
- User name- and password-based mutual authentication (see [Figure 28–5](#))

When using certificate-based mutual authentication, the following actions occur:

1. A client requests access to a protected resource.
2. The web server presents its certificate to the client.
3. The client verifies the server’s certificate.
4. If successful, the client sends its certificate to the server.
5. The server verifies the client’s credentials.
6. If successful, the server grants access to the protected resource requested by the client.

[Figure 28–4](#) shows what occurs during certificate-based mutual authentication.

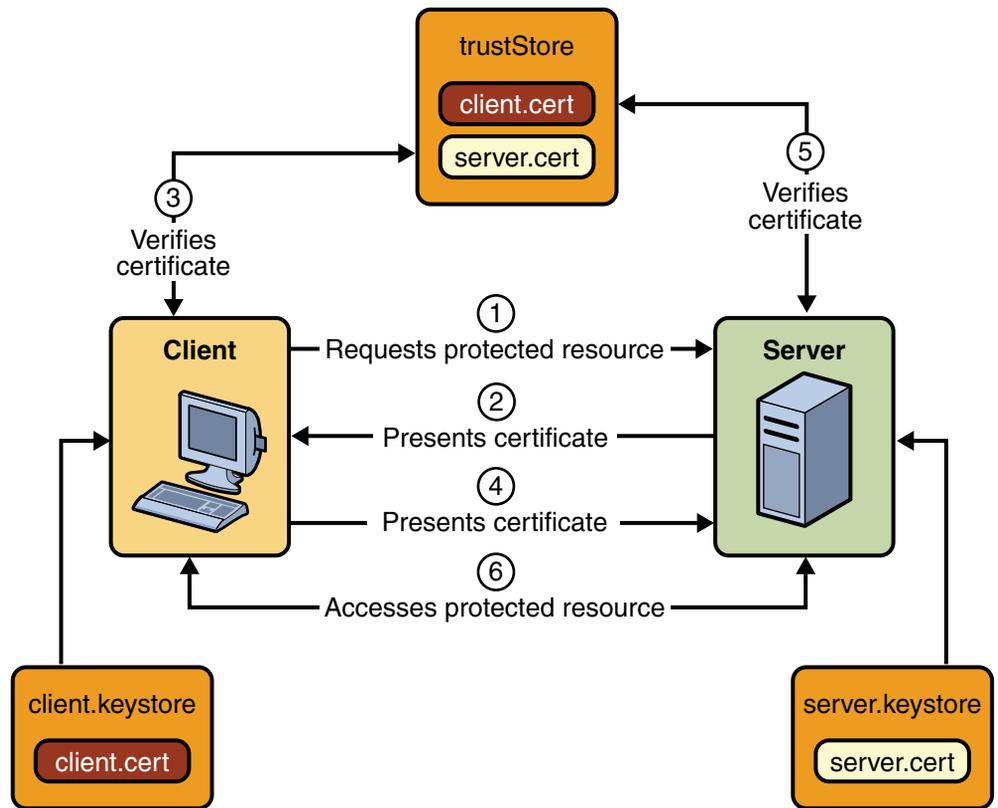


FIGURE 28-4 Certificate-Based Mutual Authentication

In user name- and password-based mutual authentication, the following actions occur:

1. A client requests access to a protected resource.
2. The web server presents its certificate to the client.
3. The client verifies the server's certificate.
4. If successful, the client sends its user name and password to the server, which verifies the client's credentials.
5. If the verification is successful, the server grants access to the protected resource requested by the client.

Figure 28-5 shows what occurs during user name- and password-based mutual authentication.

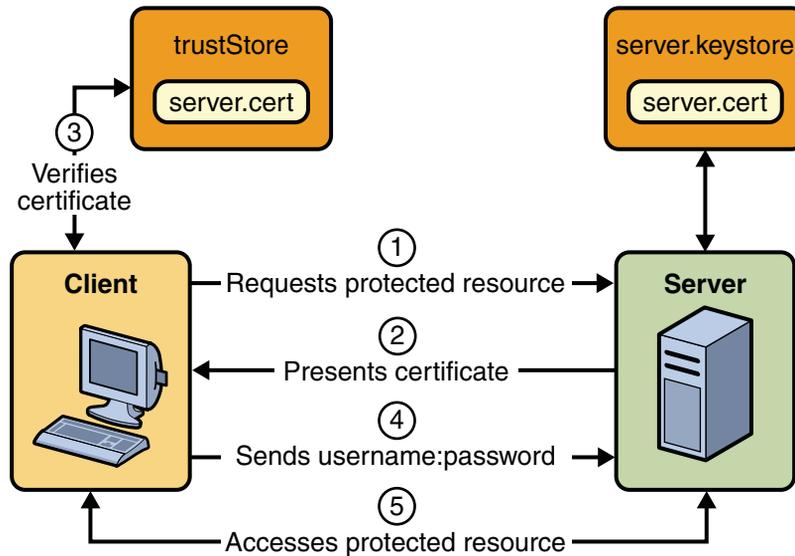


FIGURE 28-5 User Name- and Password-Based Mutual Authentication

## Digest Authentication

Like HTTP Basic Authentication, *HTTP Digest Authentication* authenticates a user based on a username and a password. However, unlike HTTP Basic Authentication, HTTP Digest Authentication does not send user passwords over the network. In HTTP Digest authentication, the client sends a one-way cryptographic hash of the password (and additional data). Although passwords are not sent on the wire, HTTP Digest authentication requires that clear text password equivalents be available to the authenticating container so that it can validate received authenticators by calculating the expected digest.

The following example shows how to declare HTTP Digest authentication in your deployment descriptor:

```
<login-config>
  <auth-method>DIGEST</auth-method>
</login-config>
```

## Working with Security Roles

In an application, a *role* is an abstract name for a group of users. A role can be compared to a key that can open a lock. Many people might have a copy of the key. The lock doesn't care who you are, only that you have the right key.

For example, in a corporation, you might have the roles Director, Manager, HR, and Employee. When an application developer is creating an internal payroll website, the developer would use the same set of data for all of its corporate users, but would allow different access to the data depending on the role the user is in. For example, a person in the role of HR would have permission to create new Employees, and to modify the payroll information for Employees. The Employee would be able to view their own payroll information, but would not be able to change some of the data, such as their pay rate, but could change some other data, such as their address or dependent information. The users' assigned role determines what permissions that user is granted for access to a particular set of resources in an application.

The following elements in a deployment descriptor use security roles in some capacity:

- `<security-role>`

A *security role* element is used to define the security roles that will comprise the entire set of security roles used in the application. The sub-element `role-name` designates the name of the security role. All role names that are used in an application should be specified in its deployment descriptor.
- `<security-role-ref>`

The *security role reference* element is used in conjunction with the `HttpServletRequest.isUserInRole(String role)` programmatic security option. When a call is made to `isUserInRole`, the caller identity is tested for membership in this role. If a `security-role-ref` has been defined, the caller is tested for membership in the role linked, or mapped, to the role name.

The security *role name* specified here is the security role name used in the application code. The value of the `role-name` element must be the `String` used as the parameter to the `HttpServletRequest.isUserInRole(String role)` method. The container uses the mapping of `security-role-ref` to `security-role` when determining whether or not the user is authorized to access the requested information.

The security *role link* specified here contains the value of the name of the security role that the user may be mapped into. The `role-link` element is used to link a security role reference to a defined security role. The `role-link` element must contain the name of one of the security roles defined in the `security-role` elements.

For more information about using `security-role-ref` with the `isUserInRole` method, read [“Declaring and Linking Role References”](#) on page 553.

## Reviewing Security Concepts

If you read [“Working with Realms, Users, Groups, and Roles”](#) on page 486, you will remember the following information:

- On the Enterprise Server, the following concepts need to be understood in order to work with security roles.
  - A *realm* is a complete database of *users* and *groups* that identify valid users of a web application (or a set of web applications) and are controlled by the same authentication policy. For more information, read [“What Is a Realm?”](#) on page 487.
  - A *user* is an individual (or application program) identity that has been defined in the Enterprise Server. On the Enterprise Server, a user generally has a user name, a password, and, optionally, a list of *groups* to which this user has been assigned. For more information, read [“What Is a User?”](#) on page 488.
  - A *group* is a set of authenticated *users*, classified by common traits, defined in the Enterprise Server. For more information, read [“What Is a Group?”](#) on page 488.
  - A *principal* is an entity that can be authenticated by an authentication protocol in a security service that is deployed in an enterprise.

For more information on configuring users on the Enterprise Server, read [“Managing Users and Groups on the Enterprise Server”](#) on page 489.

- During deployment, the deployer takes the information provided in the application deployment descriptor and maps the roles specified for the application to users and groups defined on the server using the Enterprise Server deployment descriptors `sun-web.xml`, `sun-ejb-jar.xml`, or `sun-application.xml`.

For more information, read [“Mapping Security Roles to Enterprise Server Groups”](#) on page 547.

- The concepts of users, groups, principals, and realms exist in most application or web servers, but might use different names in different products. If you are using a server other than Enterprise Server, consult your product's documentation for the terminology specific to that server.

## Declaring Security Roles

You can declare security role names used in web applications using the `security-role` element of the deployment descriptor. Use this element to list all of the security roles that you have referenced in your application, and also in conjunction with the `security-role-ref` element (see [“Declaring and Linking Role References”](#) on page 553.)

The following snippet of a deployment descriptor is taken from the `simple` sample application. This snippet declares the roles that will be used in the application using the `security-role` element, and specifies which of these roles is authorized to access protected resources using the `auth-constraint` element.

```

<security-constraint>
  <web-resource-collection>
    <web-resource-name>Protected Area</web-resource-name>
    <url-pattern>/jsp/security/protected/*</url-pattern>
    <http-method>PUT</http-method>
    <http-method>DELETE</http-method>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>manager</role-name>
  </auth-constraint>
</security-constraint>

  <!-- Security roles used by this web application -->
<security-role>
  <role-name>manager</role-name>
</security-role>
<security-role>
  <role-name>employee</role-name>
</security-role>

```

In this example, the `security-role` element lists all of the security roles used in the application: `manager` and `employee`. This enables the deployer to map all of the roles defined in the application to users and groups defined on the Enterprise Server.

The `auth-constraint` element specifies the role, `manager`, that can access the HTTP methods `PUT`, `DELETE`, `GET`, `POST` located in the directory specified by the `url-pattern` element (`/jsp/security/protected/*`).

## Mapping Security Roles to Enterprise Server Groups

To map security roles defined in applications to Enterprise Server principals and groups, use the `security-role-mapping` element in the *runtime* deployment descriptor (DD). The runtime deployment descriptor is different from the *application* deployment descriptor file. The runtime deployment descriptor is an XML file that contains information such as the context root of the web application and the mapping of the portable names of an application's resources to the Enterprise Server resources. The Enterprise Server web application runtime DD is located in the `/WEB-INF/` directory along with the web application deployment descriptor. Runtime deployment descriptors are named `sun-web.xml`, `sun-application.xml`, or `sun-ejb-jar.xml`.

The following example demonstrates how to do this mapping in the `sun-web.xml` file, which is the file used for web applications:

```

<sun-web-app>

  <security-role-mapping>

```

```
        <role-name>CEO</role-name>
        <principal-name>Schwartz</principal-name>
    </security-role-mapping>

    <security-role-mapping>
        <role-name>Admin</role-name>
        <group-name>director</group-name>
    </security-role-mapping>

    ...

</sun-web-app>
```

A role can be mapped to specific principals, specific groups, or both. The principal or group names must be valid principals or groups in the current default realm, or in the realm specified in the `login-config` element. In this example, the role of CEO that is used in the application is mapped to a principal named Schwartz that exists on the application server. Mapping a role to a specific principal is useful when the person occupying that role may change. For this application, you would only need to modify the runtime deployment descriptor and not search and replace throughout the application for references to this principal.

Also in this example, the role of Admin is mapped to a group of users who are assigned the group name of director. This is useful because the group of people authorized to access director-level administrative data only has to be maintained on the Enterprise Server. The application developer does not need to know who these people are, just define the group of people who will be given access to the information.

The `role-name` must match the `role-name` in the `security-role` element of the corresponding application deployment descriptor (`web.xml`, `ejb-jar.xml`) or the role name defined in an `@DeclareRoles` annotation.

Sometimes the role names used in the application are the same as the group names defined on the Enterprise Server. Under these circumstances, you can use the Admin Console to define a default principal-to-role-mapping that applies to the entire Enterprise Server instance. From the Admin Console, select Configuration, then Security, then check the Enable box beside Default Principal to Role Mapping. For more information, read the [Sun GlassFish Enterprise Server v3 Administration Guide](#).

## Using Programmatic Security with Web Applications

Programmatic security is used by security-aware applications when declarative security alone is not sufficient to express the security model of the application. The following topics are discussed in this section:

- “Authenticating Users Programmatically” on page 549
- “Checking Caller Identity Programmatically” on page 551

- [“Example Code for Programmatic Security” on page 551](#)
- [“Declaring and Linking Role References” on page 553](#)

## Authenticating Users Programmatically

Servlet 3.0 specifies the following methods of the `HttpServletRequest` interface that enable you to authenticate users for a web application programmatically:

- `authenticate`  
The `authenticate` method allows an application to instigate authentication of the request caller by the container from within an unconstrained request context. A login dialog box displays and collects the user's name and password for authentication purposes.
- `login`  
The `login` method allows an application to collect username and password information as an alternative to specifying form-based authentication in an application deployment descriptor.
- `logout`  
The `logout` method is provided to allow an application to reset the caller identity of a request.

The following example code shows how to use the `login` and `logout` methods:

```
package test;

import java.io.IOException;
import java.io.PrintWriter;
import java.math.BigDecimal;
import javax.ejb.EJB;
import javax.servlet.ServletException;
import javax.servlet.annotation.WebServlet;
import javax.servlet.http.HttpServlet;
import javax.servlet.http.HttpServletRequest;
import javax.servlet.http.HttpServletResponse;

@WebServlet(name="TutorialServlet", urlPatterns={"/TutorialServlet"})
public class TutorialServlet extends HttpServlet {
    @EJB
    private ConverterBean converterBean;

    /**
     * Processes requests for both HTTP GET
     * and POST methods.
     * @param request servlet request
     * @param response servlet response
     */
}
```

```
* @throws ServletException if a servlet-specific error occurs
* @throws IOException if an I/O error occurs
*/
protected void processRequest(HttpServletRequest request,
    HttpServletResponse response)
    throws ServletException, IOException {
    response.setContentType("text/html;charset=UTF-8");
    PrintWriter out = response.getWriter();
    try {

        out.println("<html>");
        out.println("<head>");
        out.println("<title>Servlet TutorialServlet</title>");
        out.println("</head>");
        out.println("<body>");
        request.login("TutorialUser", "TutorialUser");
        BigDecimal result = converterBean.dollarToYen(new BigDecimal("1.0"));
        out.println("<h1>Servlet TutorialServlet result of dollarToYen= " + result + "</h1>");
        out.println("</body>");
        out.println("</html>");
    } catch (Exception e) {
        throw new ServletException(e);
    } finally {
        request.logout();
        out.close();
    }
}
}
```

This code sample shows how to use the `authenticate` method:

```
package com.sam.test;

import java.io.*;

import javax.servlet.*;
import javax.servlet.http.*;

public class TestServlet extends HttpServlet {

    protected void processRequest(HttpServletRequest request,
        HttpServletResponse response)
        throws ServletException, IOException {
        response.setContentType("text/html;charset=UTF-8");
        PrintWriter out = response.getWriter();
        try {
            request.authenticate(response);
            out.println("Authenticate Successful");
        }
    }
}
```

```

        } finally {
            out.close();
        }
    }
}

```

## Checking Caller Identity Programmatically

In general, security management should be enforced by the container in a manner that is transparent to the web component. The security API described in this section should be used only in the less frequent situations in which the web component methods need to access the security context information.

Servlet 3.0 specifies the following methods of the `HttpServletRequest` interface that enable you to verify a caller's identity programmatically, and to use that information to grant or deny access to data:

- The `HttpServletRequest` interface provides the following methods that enable you to access security information about the component's caller: `getRemoteUser`: Determines the user name with which the client authenticated. The `getRemoteUser` method returns the name of the remote user (that is, the caller) associated by the container with the request. If no user has been authenticated, this method returns `null`.
- `isUserInRole`: Determines whether a remote user is in a specific security role. If no user has been authenticated, this method returns `false`. This method expects a `String` `role-name` parameter.

The `<security-role-ref>` element should be declared in the deployment descriptor with a `<role-name>` sub-element containing the role name to be passed to the method. Using security role references is discussed in [“Declaring and Linking Role References” on page 553](#).

- `getUserPrincipal`: Determines the principal name of the current user and returns a `java.security.Principal` object. If no user has been authenticated, this method returns `null`. Calling the `getName` method on the `Principal` returned by `getUserPrincipal` returns the name of the remote user.

Your application can make business logic decisions based on the information obtained using these APIs.

## Example Code for Programmatic Security

The following code demonstrates the use of programmatic security for the purposes of programmatic login.

```

package enterprise.programmatic_login;

import java.io.*;

```

```
import java.net.*;

import javax.annotation.security.DeclareRoles;
import javax.servlet.*;
import javax.servlet.http.*;

@DeclareRoles("javaee6user")
public class LoginServlet extends HttpServlet {

    /**
     * Processes requests for both HTTP GET and POST methods.
     * @param request servlet request
     * @param response servlet response
     */
    protected void processRequest(HttpServletRequest request,
        HttpServletResponse response)
        throws ServletException, IOException {
        response.setContentType("text/html;charset=UTF-8");
        PrintWriter out = response.getWriter();
        try {
            String userName = request.getParameter("txtUserName");
            String password = request.getParameter("txtPassword");

            out.println("Before Login"+"<br><br>");
            out.println("IsUserInRole?.." + request.isUserInRole("javaee6user")+"<br>");
            out.println("getRemoteUser?.." + request.getRemoteUser()+"<br>");
            out.println("getUserPrincipal?.." + request.getUserPrincipal()+"<br>");
            out.println("getAuthType?.." + request.getAuthType()+"<br><br>");

            try {
                request.login(userName, password);
            }catch(ServletException ex) {
                out.println("Login Failed with a ServletException.." +
                    ex.getMessage());
                return;
            }

            out.println("After Login..."+"<br><br>");
            out.println("IsUserInRole?.." + request.isUserInRole("javaee6user")+"<br>");
            out.println("getRemoteUser?.." + request.getRemoteUser()+"<br>");
            out.println("getUserPrincipal?.." + request.getUserPrincipal()+"<br>");
            out.println("getAuthType?.." + request.getAuthType()+"<br><br>");

            request.logout();
            out.println("After Logout..."+"<br><br>");
            out.println("IsUserInRole?.." + request.isUserInRole("javaee6user")+"<br>");
            out.println("getRemoteUser?.." + request.getRemoteUser()+"<br>");
            out.println("getUserPrincipal?.." + request.getUserPrincipal()+"<br>");
            out.println("getAuthType?.." + request.getAuthType()+"<br>");
```

```

        } finally {
            out.close();
        }
    }

    /**
     * Handles the HTTP <code>GET</code> method.
     * @param request servlet request
     * @param response servlet response
     */
    protected void doGet(HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {
        processRequest(request, response);
    }

    /**
     * Handles the HTTP <code>POST</code> method.
     * @param request servlet request
     * @param response servlet response
     */
    protected void doPost(HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {
        processRequest(request, response);
    }

    /**
     * Returns a short description of the servlet.
     */
    public String getServletInfo() {
        return "Short description";
    }
}

```

This servlet displays the name of the user who is currently logged in, then prompts a user for login, and prints out the information again, then logs out the user and prints the information again in order to demonstrate the effect of the `login` and `logout` methods. The application deployment descriptor, `web.xml`, includes a `login-config` element that specifies basic authentication. The runtime deployment descriptor, `sun-web.xml`, includes a role-mapping of the user of the application to a group of the same name on the Enterprise Server.

## Declaring and Linking Role References

A *security role reference* defines a mapping between the name of a role that is called from a web component using `isUserInRole(String role)` and the name of a security role that has been defined for the application. If no `<security-role-ref>` element is declared in a deployment

descriptor, and the `isUserInRole` method is called, the container defaults to checking the provided role name against the list of all security roles defined for the web application. Using the default method instead of using the `security-role-ref` element limits your flexibility to change role names in an application without also recompiling the servlet making the call.

The `security-role-ref` element is used when an application uses the `HttpServletRequest.isUserInRole(String role)` method. The value passed to the `isUserInRole` method is a `String` representing the role name of the user. The value of the `role-name` element must be the `String` used as the parameter to the `HttpServletRequest.isUserInRole(String role)` method. The `role-link` must contain the name of one of the security roles defined in the `security-role` elements. The container uses the mapping of `security-role-ref` to `security-role` when determining the return value of the call.

For example, to map the security role reference `cust` to the security role with role name `bankCustomer`, the syntax would be:

```
<servlet>
...
  <security-role-ref>
    <role-name>cust</role-name>
    <role-link>bankCustomer</role-link>
  </security-role-ref>
...
</servlet>
```

In this case, if the servlet called by a user belonging to the `bankCustomer` security role made the API call, `isUserInRole("cust")` would return `true`.

The `<role-link>` element in the `security-role-ref` element must match a `<role-name>` defined in the `<security-role>` element of the same `web.xml` deployment descriptor, as shown here:

```
<security-role>
  <role-name>bankCustomer</role-name>
</security-role>
```

A security role reference, including the name defined by the reference, is scoped to the component whose deployment descriptor contains the `security-role-ref` deployment descriptor element.

## Using Message Security with Web Applications

Web Services Security: SOAP Message Security (WS-Security) is an international standard for interoperable Web Services Security that was collaboratively developed in OASIS by 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://www.oasis-open.org/committees/download.php/3281/WSS-SOAPMessageSecurity-17-082703-merged.pdf>

WS-Security incorporates security features in the header of a SOAP message, working in the application layer. Message security differs from transport layer security (which is what is discussed in this chapter) in that message security can be used to decouple message protection from message transport so that messages remain protected after transmission, ensuring end-to-end security.

Sun's implementation of WS-Security is part of Metro, a web service stack. In the past, web services have relied on transport-based security such as SSL to provide point-to-point security. Metro implements the WS-Security specification to provide interoperable message content integrity and confidentiality, even in the presence of intermediaries. Metro also provides an implementation of the WS-Trust specification as a means for issuing, renewing, and validating security tokens used by WS-Security, and to establish and broker trust relationships. That portion of Metro is known as WSIT (Web Services Interoperability Technologies). Metro's WSIT subsystem is an implementation of a number of open web services specifications to support enterprise features. In addition to security, reliable messaging, and atomic transactions, Metro includes a bootstrapping and configuration technology.

Message security is not part of the Java EE 6 platform, but can be used with and by Java EE 6 applications by following the instructions from the *Metro User's Guide* at <https://metro.dev.java.net/guide/>.

## Examples: Securing Web Applications

There is some basic setup required before any of the example applications will run correctly. Refer to “[Setting Up Your System for Running the Security Examples](#)” on page 556 to make sure you have completed these steps prior to continuing with the examples.

The following examples use annotations, programmatic security, and/or declarative security to demonstrate adding security to existing web applications:

- “[Setting Up Your System for Running the Security Examples](#)” on page 556
- “[Example: Basic Authentication with a Servlet](#)” on page 556

- “Example: Basic Authentication with JAX-WS” on page 562
- “Example: Form-Based Authentication with a Servlet” on page 566

Here are some links to other locations where you will find examples of securing different types of applications:

- “Example: Securing an Enterprise Bean” on page 512
- “Example: Using the `isCallerInRole` and `getCallerPrincipal` Methods” on page 516
- GlassFish samples: <https://glassfish-samples.dev.java.net/>

## Setting Up Your System for Running the Security Examples

To set up your system for running the security examples, you basically need to configure a user database that can be used by the application for authenticating users. Here are the steps you need to complete before continuing:

- If you have not already done so, make sure you have read and followed the directions for installing the tutorial examples, Ant, and NetBeans IDE, and that you understand how to start the Enterprise Server and Administration Console. These instructions can be found in [Chapter 2, “Using the Tutorial Examples.”](#)
- If you have not already done so, add an authorized user to the Enterprise Server. For this example, add users to the `file` realm of the Enterprise Server and assign the user to the group `TutorialUser`. Be sure to write down the user name and password for the user you create so that you can use it for testing the example applications. Authentication is case-sensitive for both the user name and password, so write down the user name and password exactly. This topic is discussed more in [“Managing Users and Groups on the Enterprise Server” on page 489.](#)
- Set up Default Principal to Role Mapping on the Enterprise Server. From the Admin Console, select Configuration, then Security, then check the enable box beside Default Principal to Role Mapping.

## Example: Basic Authentication with a Servlet

This example discusses how to use basic authentication with a servlet. With basic authentication of a servlet, the web browser presents a standard login dialog that is not customizable. When a user submits their name and password, the server determines if the user name and password are those of an authorized user and sends the requested web resource if the user is authorized to view it.

In general, the following steps are necessary for adding basic authentication to an unsecured servlet, such as the one described in [Chapter 3, “Getting Started with Web Applications.”](#) In the example application included with this tutorial, many of these steps have been completed for you and are listed here simply to show what needs to be done should you wish to create a similar application. The completed version of this example application can be found in the directory `tut-install/examples/web/hello2_basicauth/`.

1. Follow the steps in [“Setting Up Your System for Running the Security Examples”](#) on page 556.
2. Create a web module as described in [Chapter 3, “Getting Started with Web Applications,”](#) for the servlet example, `hello2`.
3. Add the appropriate security elements to the `web.xml` deployment descriptor. The deployment descriptor for the example application can be viewed at `tut-install/examples/web/hello2_basicauth/web/WEB-INF/web.xml`. The security elements are described in [“Specifying Security in the Deployment Descriptor”](#) on page 557.
4. Build, package, and deploy the web application by following the steps in [“Building, Packaging, and Deploying the Servlet Basic Authentication Example Using NetBeans IDE”](#) on page 558 or [“Building, Packaging, and Deploying the Servlet Basic Authentication Example Using Ant”](#) on page 559.
5. Run the web application by following the steps described in [“Running the Basic Authentication Servlet”](#) on page 559.
6. If you have any problems running this example, refer to the troubleshooting tips in [“Troubleshooting the Basic Authentication Example”](#) on page 561.

## Specifying Security in the Deployment Descriptor

The elements of the deployment descriptor that add basic authentication to this example tells the server or browser to perform the following tasks:

- Send a standard login dialog to collect user name and password data
- Verify that the user is authorized to access the application
- If authorized, display the servlet to the user

Deployment descriptors elements are described in [“Introduction to Web Application Deployment Descriptors”](#) on page 530.

The following sample code shows the security elements for the deployment descriptor used in this example of basic authentication, which can be found in `tut-install/examples/web/hello2_basicauth/web/WEB-INF/web.xml`.

```
<security-constraint>
  <display-name>SecurityConstraint</display-name>
  <web-resource-collection>
    <web-resource-name>WRCollection</web-resource-name>
    <url-pattern>/greeting</url-pattern>
```

```
</web-resource-collection>
<auth-constraint>
  <role-name>TutorialUser</role-name>
</auth-constraint>
<user-data-constraint>
  <transport-guarantee>CONFIDENTIAL</transport-guarantee>
</user-data-constraint>
</security-constraint>
<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>file</realm-name>
</login-config>
<security-role>
  <role-name>TutorialUser</role-name>
</security-role>
```

This deployment descriptor shows that all the request URI `/greeting` can only be accessed by users who have entered their user name and password and have been authorized to access this URL because they have been verified to be in the role `TutorialUser`. The data will be sent over a protected transport in order to keep the user name and password data from being read in transit.

## Building, Packaging, and Deploying the Servlet Basic Authentication Example Using NetBeans IDE

To build, package, and deploy the `web/hello2_basicauth` example application using NetBeans IDE, follow these steps:

1. Follow the steps in [“Setting Up Your System for Running the Security Examples”](#) on [page 556](#).
2. Open the project in NetBeans IDE by selecting `File→Open Project`.
3. Browse to the `tut-install/examples/web/hello2_basicauth/` directory.
4. Make sure that `Open as Main Project` is selected.
5. Select `Open Project`.
6. Right-click `hello2_basicauth` in the `Projects` pane, then select `Clean and Build`.
7. Right-click `hello2_basicauth` in the `Projects` pane, then select `Deploy`.
8. To run the servlet, follow the steps in [“Running the Basic Authentication Servlet”](#) on [page 559](#).

## Building, Packaging, and Deploying the Servlet Basic Authentication Example Using Ant

To build, package, and deploy the `web/hello2_basicauth` example using the Ant tool, follow these steps:

1. Follow the steps in “[Setting Up Your System for Running the Security Examples](#)” on [page 556](#).
2. From a terminal window or command prompt, change to the `tut-install/examples/web/hello2_basicauth/` directory.
3. Build and package the web application by entering the following command at the terminal window or command prompt:

```
ant
```

4. To deploy the example using Ant, enter the following command at the terminal window or command prompt:

```
ant deploy
```

The `deploy` target in this case gives you an incorrect URL to run the application. To run the application, please use the URL shown in “[Running the Basic Authentication Servlet](#)” on [page 559](#).

5. To run the web application, follow the steps in “[Running the Basic Authentication Servlet](#)” on [page 559](#).

## Running the Basic Authentication Servlet

To run the web client, follow these steps:

1. Open a web browser.
2. Enter the following URL in your web browser:

```
https://localhost:8181/hello2_basicauth/greeting
```

You may be prompted to accept the security certificate for the server. If so, accept the security certificate.

3. A default login form displays. Enter a user name and password combination that corresponds to a user that has already been created in the file realm of the Enterprise Server and has been assigned to the group of `TutorialUser`.

Basic authentication is case-sensitive for both the user name and password, so enter the user name and password exactly as defined for the Enterprise Server.

The server returns the requested resource if all of the following conditions are met:

- There is a user defined for the Enterprise Server with the user name you entered.
- The user with the user name you entered has the password you entered.

- The user name and password combination you entered is assigned to the group of TutorialUser on the Enterprise Server.
- The role of TutorialUser, as defined for the application, is mapped to the group of TutorialUser, as defined for the Enterprise Server.

When these conditions are met, and the server has authenticated the user, the application will display as shown in [Figure 28–6](#).

4. Enter your name and click the Submit button. Because you have already been authorized, the name you enter in this step does not have any limitations. You have unlimited access to the application now.

The application responds by saying “Hello” to you, as shown in [Figure 28–7](#).



FIGURE 28–6 Running the Application

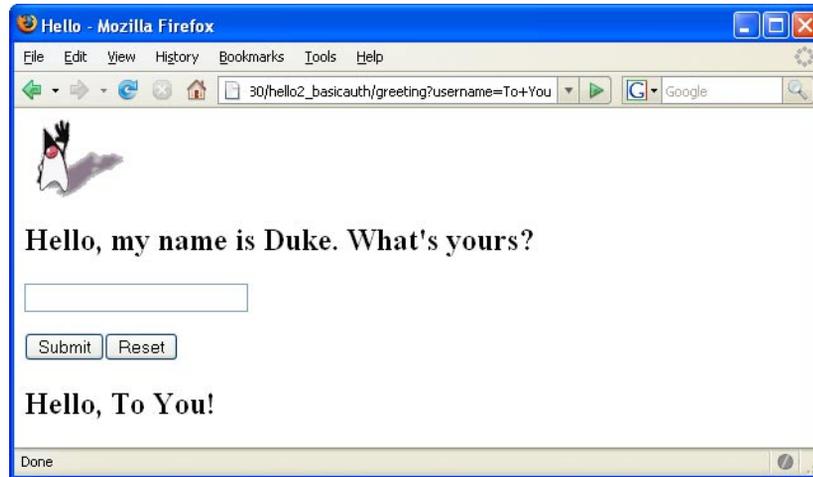


FIGURE 28-7 The Running Basic Authentication Response

---

**Note** – For repetitive testing of this example, you may need to close and reopen your browser. You should also run the `ant clean` and `ant undeploy` targets or the NetBeans IDE Clean and Build option to get a fresh start.

---

## Troubleshooting the Basic Authentication Example

When doing iterative development with this web application, follow these steps if you are using NetBeans IDE:

1. Close your web browser.
2. Clean and recompile the files from the previous build by right-clicking `hello2_basicauth` and selecting Clean and Build.
3. Redeploy the application by right-clicking `hello2_basicauth` and selecting Undeploy and Deploy.
4. Open your web browser and reload the following URL:

`https://localhost:8181/hello2_basicauth/greeting`

Follow these steps if you are using the Ant tool:

1. Close your web browser.
2. Undeploy the web application. To undeploy the application, use the following command in the directory:

```
ant undeploy
```

3. Clean out files from the previous build, using the following command:

```
ant clean
```

4. Recompile, repackage, and redeploy the application, using the following commands:

```
ant
```

```
ant deploy
```

5. Open your web browser and reload the following URL:

```
https://localhost:8181/hello2_basicauth/greeting
```

## Example: Basic Authentication with JAX-WS

This section discusses how to configure a JAX-WS-based web service for HTTP basic authentication. When a service that is constrained by *HTTP basic authentication* is requested, the server requests a user name and password from the client and verifies that the user name and password are valid by comparing them against a database of authorized users.

For this tutorial, you will add the security elements to the JAX-WS service; build, package, and deploy the service; and then build and run the client application.

This example does not include a finished application, but provides instructions in the event that you want to secure a JAX-WS web service, such as the one that can be found in the directory `tut-install/examples/jaxws/helloservice` and is discussed in [“Creating a Simple Web Service and Client with JAX-WS” on page 238](#). You build on this simple application by adding the necessary elements to secure the application using basic authentication.

In general, the following steps are necessary to add basic authentication to a JAX-WS web service.

1. Create an application like the one in [“Creating a Simple Web Service and Client with JAX-WS” on page 238](#).
2. Follow the steps in [“Setting Up Your System for Running the Security Examples” on page 556](#).
3. Add security elements that specify that basic authentication is to be performed to the application deployment descriptor, `web.xml`. This step is discussed in [“Adding Security Elements to the Deployment Descriptor” on page 563](#).
4. Build, package, and deploy the web service. See [“Building and Deploying `helloservice` with Basic Authentication Using NetBeans IDE” on page 564](#) or [“Building and Deploying `helloservice` with Basic Authentication Using Ant” on page 564](#) for the steps to accomplish this.
5. Build and run the client application. See [“Building and Running the `helloservice` Client Application with Basic Authentication Using NetBeans IDE” on page 565](#) or [“Building and Running the `helloservice` Client Application with Basic Authentication Using Ant” on page 565](#) for the steps to accomplish this.

## Adding Security Elements to the Deployment Descriptor

To enable basic authentication for the service, add security elements to the application deployment descriptor, `web.xml`. The security elements that need to be added to the deployment descriptor include the `<security-constraint>`, `<login-config>`, and `<security-role>` elements. These security elements are discussed in more detail in [“Introduction to Web Application Deployment Descriptors” on page 530](#) and in the Java Servlet Specification. The code is added to the original deployment descriptor to enable HTTP basic authentication. The resulting deployment descriptor looks like this:

```
<?xml version="1.0" encoding="UTF-8"?><web-app
  xmlns="http://java.sun.com/xml/ns/javaee" version="2.5"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
  http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd">
  <display-name>HelloService</display-name>
  <listener>
    <listener-class>
      com.sun.xml.ws.transport.http.servlet.WSServletContextListener
    </listener-class>
  </listener>
  <servlet>
    <display-name>HelloService</display-name>
    <servlet-name>HelloService</servlet-name>
    <servlet-class>com.sun.xml.ws.transport.http.servlet.WSServlet</servlet-class>
  </servlet>
  <servlet-mapping>
    <servlet-name>HelloService</servlet-name>
    <url-pattern>/hello</url-pattern>
  </servlet-mapping>
  <session-config>
    <session-timeout>30</session-timeout>
  </session-config>
  <security-constraint>
    <display-name>SecurityConstraint</display-name>
    <web-resource-collection>
      <web-resource-name>WRCollection</web-resource-name>
      <url-pattern>/hello</url-pattern>
    </web-resource-collection>
    <auth-constraint>
      <role-name>TutorialUser</role-name>
    </auth-constraint>
    <user-data-constraint>
      <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    </user-data-constraint>
  </security-constraint>
  <login-config>
    <auth-constraint>BASIC</auth-constraint>
```

```
<realm-name>file</realm-name>
</login-config>
<security-role>
  <role-name>TutorialUser</role-name>
</security-role>
</web-app>
```

This security constraint protects resources at the URI `/hello`. Anyone who tries to access this resource will be prompted for their user name and password and must be authenticated by the Enterprise Server before they will be granted access to the resource. The request is sent over a protected transport to ensure that the username and password are not intercepted in transit.

## Building and Deploying `helloworldservice` with Basic Authentication Using NetBeans IDE

To build, package, and deploy the example using NetBeans IDE, follow these steps, or the steps described in [“Building, Packaging, and Deploying the Service” on page 240](#).

1. Follow the steps in [“Setting Up Your System for Running the Security Examples” on page 556](#).
2. In NetBeans IDE, select `File`→`Open Project`.
3. In the Open Project dialog, navigate to the project.
4. Click `Open Project`.
5. In the Projects tab, right-click the project and select `Clean and Build`.
6. In the Projects tab, right-click the project and select `Deploy`.

This step builds and packages the application into a WAR file, and deploys this war file to your Enterprise Server instance.

## Building and Deploying `helloworldservice` with Basic Authentication Using Ant

To build, package, and deploy the project using the Ant tool, follow these steps, or the steps described in [“Building, Packaging, and Deploying the Service” on page 240](#).

1. Follow the steps in [“Setting Up Your System for Running the Security Examples” on page 556](#).
2. From a terminal window or command prompt, go to the project directory.
3. Build, package, and deploy the JAX-WS service by entering the following at the terminal window or command prompt in the project directory:

```
ant all
```

You can test the service in the Admin Console. For more information on how to do this, read [“Testing the Service without a Client” on page 242](#).

## Building and Running the `helloservice` Client Application with Basic Authentication Using NetBeans IDE

To build and run the client application using NetBeans IDE, follow these steps. The service must be deployed onto the Enterprise Server before compiling the client files. For information on deploying the service, read [“Building and Deploying `helloservice` with Basic Authentication Using NetBeans IDE”](#) on page 564.

1. In NetBeans IDE, select File→Open Project.
2. In the Open Project dialog, navigate to the project.
3. Click Open Project.
4. In the Projects tab, right-click the project and select Clean and Build.
5. In the Projects tab, right-click the project and select Run.

You will be prompted for your user name and password.

6. Enter the user name and password of a user that has been entered into the database of users for the file realm and has been assigned to the group of `TutorialUser`.

If the username and password you enter are authorized, you will see the output of the application client in the Output pane.

## Building and Running the `helloservice` Client Application with Basic Authentication Using Ant

To build and run the client application using the Ant tool, follow these steps. The secured service must be deployed onto the Enterprise Server before you can successfully compile the client application. For more information on deploying the service, read [“Building and Deploying `helloservice` with Basic Authentication Using Ant”](#) on page 564.

1. Build the client by changing to the project directory and entering the following at the terminal window or command prompt:

```
ant
```

This command calls the default target, which builds and packages the application into a JAR file.

2. Run the client by entering the following at the terminal window or command prompt:

```
ant run
```

A Login for User dialog displays.

3. Enter a user name and password that correspond to a user set up on the Enterprise Server with a group of `TutorialUser`. Click OK.

## Example: Form-Based Authentication with a Servlet

This example discusses how to use form-based authentication with a basic servlet. With form-based authentication, you can customize the login screen and error pages that are presented to the web client for authentication of their user name and password. When a user submits their name and password, the server determines if the user name and password are those of an authorized user and, if authorized, sends the requested web resource.

In general, the steps are necessary for adding form-based authentication to an unsecured servlet are similar to those described in [“Example: Basic Authentication with a Servlet” on page 556](#), so just follow all of the steps in [“Example: Basic Authentication with a Servlet” on page 556](#), except use the deployment descriptor described in [“Specifying Security in the Deployment Descriptor” on page 567](#) instead and create the login form and login error form pages as described in [“Creating the Login Form and the Error Page” on page 566](#). The completed version of this example application can be found in the directory `tut-install/examples/web/hello2_formauth/`.

### Creating the Login Form and the Error Page

When using form-based login mechanisms, you must specify a page that contains the form you want to use to obtain the user name and password, as well as which page to display if login authentication fails. This section discusses the login form and the error page used in this example. The section [“Specifying Security in the Deployment Descriptor” on page 567](#) shows how you specify these pages in the deployment descriptor.

The login page can be an HTML page, a JSP page, or a servlet, and it must return an HTML page containing a form that conforms to specific naming conventions (see the Java Servlet 3.0 specification for more information on these requirements). To do this, include the elements that accept user name and password information between `<form></form>` tags in your login page. The content of an HTML page, JSP page, or servlet for a login page should be coded as follows:

```
<form method=post action="j_security_check" >
  <input type="text" name="j_username" >
  <input type="password" name="j_password" >
</form>
```

The full code for the login page used in this example can be found at `tut-install/examples/web/hello2_formauth/web/loginform.html`. An example of the running login form page is shown later in [Figure 28–8](#). Here is the code for this page:

```
<html>
<head>
  <title>Login Page</title>
</head>
```

```

<h2>Hello, please log in:</h2>
<br><br>
<form action="j_security_check" method=post>
  <p><strong>Please Enter Your User Name: </strong>
    <input type="text" name="j_username" size="25">
  <p><p><strong>Please Enter Your Password: </strong>
    <input type="password" size="15" name="j_password">
  <p><p>
    <input type="submit" value="Submit">
    <input type="reset" value="Reset">
</form>
</html>

```

The login error page is displayed if the user enters a user name and password combination that is not authorized to access the protected URI. For this example, the login error page can be found at *tut-install/examples/web/hello2\_formauth/web/loginerror.html*. For this example, the login error page explains the reason for receiving the error page and provides a link that will allow the user to try again. Here is the code for this page:

```

<html>
<head>
  <title>Login Error</title>
</head>
<body>
  <c:url var="url" value="/index.jsp"/>
  <h2>Invalid user name or password.</h2>

  <p>Please enter a user name or password that is authorized to access this
  application. For this application, this means a user that has been created in the
  <code>file</code> realm and has been assigned to the <em>group</em> of
  <code>TutorialUser</code>. Click here to <a href="{$url}">Try Again</a></p>
</body>
</html>

```

## Specifying Security in the Deployment Descriptor

This example takes a very simple servlet-based web application and adds form-based security to this application. All security for this example is declared in the deployment descriptor for the application. A security constraint is defined in the deployment descriptor that tells the server to send a login form to collect user data, verify that the user is authorized to access the application, and, if so, display the JSP page to the user.

Deployment descriptor elements are described in [“Introduction to Web Application Deployment Descriptors”](#) on page 530.

The following sample code shows the deployment descriptor used in this example of form-based login authentication, which can be found in *tut-install/examples/web/hello2\_formauth/web/WEB-INF/web.xml*.

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

  <display-name>hello2_formauth</display-name>
  <servlet>
    <display-name>index</display-name>
    <servlet-name>index</servlet-name>
    <jsp-file>/index.jsp</jsp-file>
  </servlet>
  <security-constraint>
    <display-name>SecurityConstraint</display-name>
    <web-resource-collection>
      <web-resource-name>WRCollection</web-resource-name>
      <url-pattern>/*</url-pattern>
    </web-resource-collection>
    <auth-constraint>
      <role-name>TutorialUser</role-name>
    </auth-constraint>
    <user-data-constraint>
      <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    </user-data-constraint>
  </security-constraint>
  <login-config>
    <auth-method>FORM</auth-method>
    <form-login-config>
      <form-login-page>/loginform.html</form-login-page>
      <form-error-page>/loginerror.html</form-error-page>
    </form-login-config>
  </login-config>
  <security-role>
    <role-name>TutorialUser</role-name>
  </security-role>
</web-app>
```

## Building, Packaging, and Deploying the Form-Based Authentication Example Using NetBeans IDE

To build, package, and deploy this application using NetBeans IDE, follow these steps:

1. Follow the steps in “[Setting Up Your System for Running the Security Examples](#)” on [page 556](#).
2. Open the project in NetBeans IDE by selecting File→Open Project.
3. Browse to the `tut-install/examples/web/hello2_formauth/` directory.
4. Make sure that Open as Main Project is selected.

5. Select Open Project.
6. Right-click `hello2_formauth` in the Projects pane, then select Clean and Build.
7. Right-click `hello2_formauth` in the Projects pane, then select Deploy.
8. Follow the steps in [“Testing the Form-Based Authentication Web Client”](#) on page 569.

## Building, Packaging, and Deploying the Form-Based Authentication Example Using Ant

To build, package, and deploy this application using the Ant tool, follow these steps:

1. Follow the steps in [“Setting Up Your System for Running the Security Examples”](#) on page 556.
2. From a terminal window or command prompt, change to the `tut-install/examples/web/hello2_formauth/` directory.
3. Enter the following command at the terminal window or command prompt:

```
ant
```

This target will spawn any necessary compilations, copy files to the `tut-install/examples/web/hello2_formauth/build/` directory, create the WAR file, and copy it to the `tut-install/examples/web/hello2_formauth/dist/` directory.

4. Deploy the WAR named `hello2_formauth.war` onto the Enterprise Server using Ant by entering the following command at the terminal window or command prompt:

```
ant deploy
```

5. Follow the steps in [“Testing the Form-Based Authentication Web Client”](#) on page 569.

## Testing the Form-Based Authentication Web Client

To run the web client, follow these steps:

1. Open a web browser.
2. Enter the following URL in your web browser:

```
https://localhost:8181/hello2_formauth
```

The login form displays in the browser, as shown in [Figure 28–8](#).

3. Enter a user name and password combination that corresponds to a user that has already been created in the `file` realm of the Enterprise Server and has been assigned to the group of `TutorialUser`.
4. Click the Submit button. Form-based authentication is case-sensitive for both the user name and password, so enter the user name and password exactly as defined for the Enterprise Server.

If you entered `My_Name` as the name and `My_Pwd` for the password, the server returns the requested resource if all of the following conditions are met:

- There is a user defined for the Enterprise Server with the user name of My\_Name.
- The user with the user name of My\_Name has a password of My\_Pwd defined for the Enterprise Server.
- The user My\_Name with the password My\_Pwd is assigned to the group of TutorialUser on the Enterprise Server.
- The role of TutorialUser, as defined for the application, is mapped to the group of TutorialUser, as defined for the Enterprise Server.

When these conditions are met, and the server has authenticated the user, the application will display as shown in [Figure 28-9](#).

5. Enter your name and click the Submit button. Because you have already been authorized, the name you enter in this step does not have any limitations. You have unlimited access to the application now.

The application responds by saying “Hello” to you, as shown in [Figure 28-10](#).

For additional testing, close and reopen your browser, enter the application URL, and enter a username and password that are not authorized to see the login error page generated.



FIGURE 28-8 Form-Based Login Page

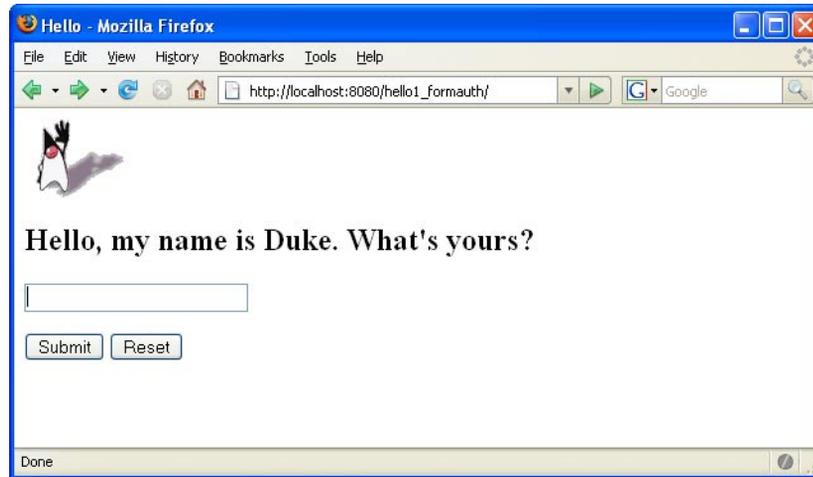


FIGURE 28-9 Running Web Application

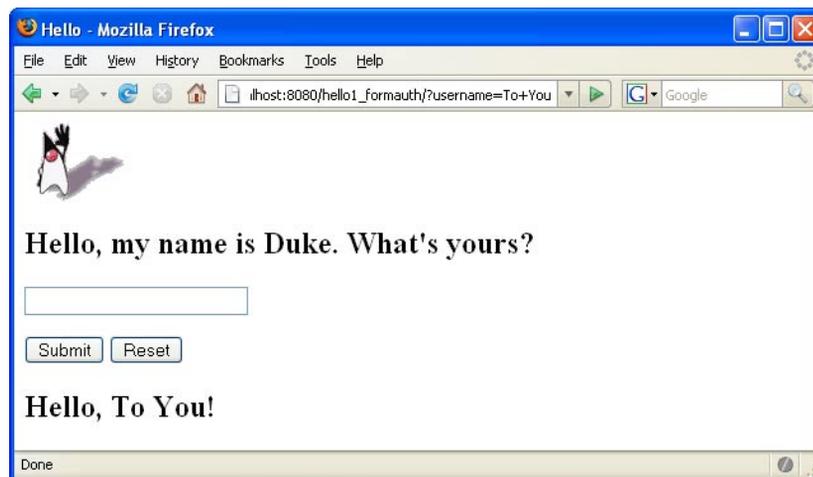


FIGURE 28-10 The Running Form-Based Authentication Example

---

**Note** – For repetitive testing of this example, you may need to close and reopen your browser. You should also run the `ant clean` and `ant undeploy` commands to ensure a fresh build if using the Ant tool, or select Clean and Build then Undeploy and Deploy if using NetBeans IDE.

---



## Advanced Concepts in Java EE Security

---

This chapter provides more information on securing applications.

### **Advanced Concepts in Java EE Security**

After you have read the introductory chapters in Java EE Security, read this chapter to learn more about using annotations with web applications and using deployment descriptors with enterprise beans.





PART VIII

## Java EE Supporting Technologies

Part Eight explores several technologies that support the Java EE platform.



# Introduction to Java™ EE Supporting Technologies

---

The Java EE platform includes several technologies and APIs that extend its functionality. These technologies allow applications to access a wide range of services in a uniform manner.

This chapter introduces the following Java EE technologies:

These technologies are explained in greater in subsequent chapters: [Chapter 31, “Transactions,”](#) and [Chapter 32, “Resource Connections.”](#)

The following topics are addressed here:

- “Transactions” on page 577
- “Resources” on page 578

## Transactions

A transaction is a series of actions in a Java EE application that must all complete successfully or else all the changes in each action are backed out. Transactions end in either a commit or a rollback.

The Java Transaction API (JTA) allows applications to access transactions in a manner that is independent of specific implementations. JTA specifies standard Java interfaces between a transaction manager and the parties involved in a distributed transaction system: the transactional application, the Java EE server, and the manager that controls access to the shared resources affected by the transactions.

The JTA defines the `UserTransaction` interface that is used by applications to start, and commit or abort transactions. Application components get a `UserTransaction` object through a JNDI lookup using the name `java:comp/UserTransaction` or by requesting injection of a `UserTransaction` object. A number of interfaces defined by JTA are used by an application server to communicate with a transaction manager, and for a transaction manager to interact with a resource manager.

See the [Chapter 31, “Transactions,”](#) chapter for a more detailed explanation. The JTA 1.1 specification is available at <http://java.sun.com/javaee/technologies/jta/index.jsp>.

## Resources

A resource is a program object that provides connections to systems such as database servers and messaging systems.

### The Java EE Connector Architecture and Resource Adapters

The Java EE Connector 1.6 architecture enables Java EE components to interact with enterprise information systems (EISs) and EISs to interact with Java EE components. EIS software includes various types of systems: enterprise resource planning (ERP), mainframe transaction processing, and nonrelational databases, among others. Connector architecture simplifies the integration of diverse EISs. Each EIS requires only one implementation of the Connector architecture. Because an implementation adheres to the Connector specification, it is portable across all compliant Java EE servers.

The specification defines the contracts for an application server as well as for resource adapters, which are system-level software drivers for specific EIS resources. These standard contracts provide pluggability between application servers and EISs. The Java EE Connector 1.6 specification defines new system contracts such as Generic Work Context and Security Inflow.

The Java EE Connector 1.6 specification is available at: <http://jcp.org/en/jsr/detail?id=322>

A resource adapter is a Java EE component that implements the Connector architecture for a specific EIS.

A resource adapter can choose to support the following levels of transactions:

- `NoTransaction`- No transaction support is provided.
- `LocalTransaction` - Resource manager local transactions are supported..
- `XATransaction` - Resource adapter supports XA and the JTA `XATransaction` interface..

See the [Chapter 32, “Resource Connections,”](#) chapter for a more detailed explanation of resource adapters.

---

## Java Message Service

Messaging is a method of communication between software components or applications. A messaging system is a peer-to-peer facility: A messaging client can send messages to, and receive messages from, any other client. Each client connects to a messaging agent that provides facilities for creating, sending, receiving, and reading messages.

The Java Message Service API is a Java API that allows applications to create, send, receive, and read messages. Designed by Sun and several partner companies, the JMS API defines a common set of interfaces and associated semantics that allow programs written in the Java programming language to communicate with other messaging implementations.

The JMS API minimizes the set of concepts a programmer must learn in order to use messaging products but provides enough features to support sophisticated messaging applications. It also strives to maximize the portability of JMS applications across JMS providers in the same messaging domain.

Sun GlassFish™ Enterprise Server v3 implements the Java Message Service (JMS) API by integrating the Sun GlassFish Message Queue software with the Enterprise Server software.

## Java DataBase Connectivity (JDBC™) Software

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

A JDBC resource (data source) provides applications with a means of connecting to a database. Typically, a JDBC resource is created for each database accessed by the applications deployed in a domain. Transactional access to JDBC resources is available from servlets, JSP pages, and enterprise beans. The connection pooling and distributed transaction features are intended for use by JDBC drivers to coordinate with an application server.

For more information, see [“DataSource Objects and Connection Pools” on page 594](#).



# Transactions

---

A typical enterprise application accesses and stores information in one or more databases. Because this information is critical for business operations, it must be accurate, current, and reliable. Data integrity would be lost if multiple programs were allowed to update the same information simultaneously. It would also be lost if a system that failed while processing a business transaction were to leave the affected data only partially updated. By preventing both of these scenarios, software transactions ensure data integrity. Transactions control the concurrent access of data by multiple programs. In the event of a system failure, transactions make sure that after recovery the data will be in a consistent state.

The following topics are addressed here:

- [“What Is a Transaction?” on page 581](#)
- [“Container-Managed Transactions” on page 582](#)
- [“Bean-Managed Transactions” on page 587](#)
- [“Transaction Timeouts” on page 589](#)
- [“Updating Multiple Databases” on page 590](#)
- [“Transactions in Web Components” on page 591](#)

## What Is a Transaction?

To emulate a business transaction, a program may need to perform several steps. A financial program, for example, might transfer funds from a checking account to a savings account using the steps listed in the following pseudocode:

```
begin transaction
    debit checking account
    credit savings account
    update history log
commit transaction
```

Either all three of these steps must complete, or none of them at all. Otherwise, data integrity is lost. Because the steps within a transaction are a unified whole, a *transaction* is often defined as an indivisible unit of work.

A transaction can end in two ways: with a commit or with a rollback. When a transaction commits, the data modifications made by its statements are saved. If a statement within a transaction fails, the transaction rolls back, undoing the effects of all statements in the transaction. In the pseudocode, for example, if a disk drive were to crash during the `credit` step, the transaction would roll back and undo the data modifications made by the `debit` statement. Although the transaction fails, data integrity would be intact because the accounts still balance.

In the preceding pseudocode, the `begin` and `commit` statements mark the boundaries of the transaction. When designing an enterprise bean, you determine how the boundaries are set by specifying either container-managed or bean-managed transactions.

## Container-Managed Transactions

In an enterprise bean with *container-managed transaction demarcation*, the EJB container sets the boundaries of the transactions. You can use container-managed transactions with any type of enterprise bean: session, or message-driven. Container-managed transactions simplify development because the enterprise bean code does not explicitly mark the transaction's boundaries. The code does not include statements that begin and end the transaction.

By default if no transaction demarcation is specified enterprise beans use container-managed transaction demarcation.

Typically, the container begins a transaction immediately before an enterprise bean method starts. It commits the transaction just before the method exits. Each method can be associated with a single transaction. Nested or multiple transactions are not allowed within a method.

Container-managed transactions do not require all methods to be associated with transactions. When developing a bean, you can specify which of the bean's methods are associated with transactions by setting the transaction attributes.

Enterprise beans that use container-managed transaction demarcation must not use any transaction management methods that interfere with the container's transaction demarcation boundaries. Examples of such methods are the `commit`, `setAutoCommit`, and `rollback` methods of `java.sql.Connection` or the `commit` and `rollback` methods of `javax.jms.Session`. If you require control over the transaction demarcation, you must use application-managed transaction demarcation.

Enterprise beans that use container-managed transaction demarcation also must not use the `javax.transaction.UserTransaction` interface.

## Transaction Attributes

A *transaction attribute* controls the scope of a transaction. [Figure 31-1](#) illustrates why controlling the scope is important. In the diagram, `method-A` begins a transaction and then invokes `method-B` of `Bean-2`. When `method-B` executes, does it run within the scope of the transaction started by `method-A`, or does it execute with a new transaction? The answer depends on the transaction attribute of `method-B`.

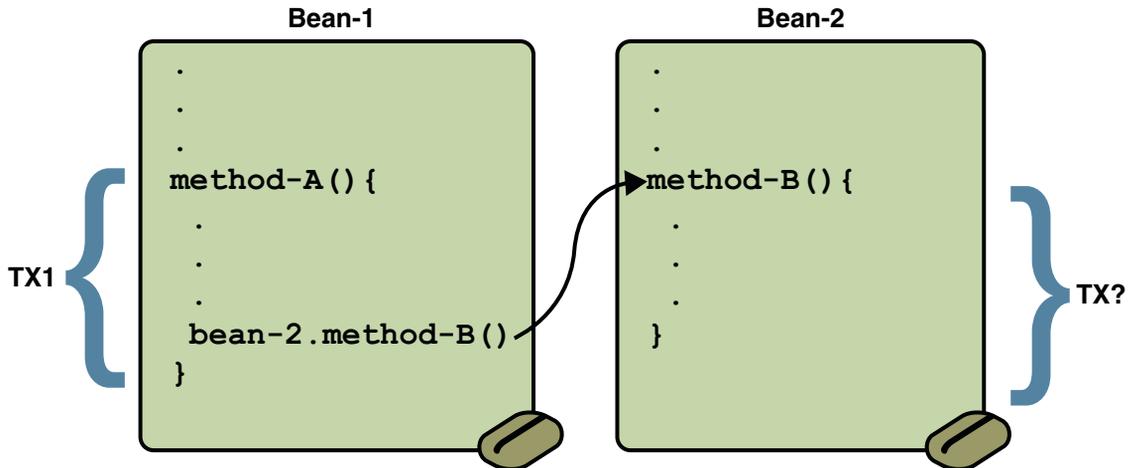


FIGURE 31-1 Transaction Scope

A transaction attribute can have one of the following values:

- Required
- RequiresNew
- Mandatory
- NotSupported
- Supports
- Never

### Required Attribute

If the client is running within a transaction and invokes the enterprise bean's method, the method executes within the client's transaction. If the client is not associated with a transaction, the container starts a new transaction before running the method.

The `Required` attribute is the implicit transaction attribute for all enterprise bean methods running with container-managed transaction demarcation. You typically do not set the `Required` attribute unless you need to override another transaction attribute. Because transaction attributes are declarative, you can easily change them later.

### RequiresNew **Attribute**

If the client is running within a transaction and invokes the enterprise bean's method, the container takes the following steps:

1. Suspends the client's transaction
2. Starts a new transaction
3. Delegates the call to the method
4. Resumes the client's transaction after the method completes

If the client is not associated with a transaction, the container starts a new transaction before running the method.

You should use the `RequiresNew` attribute when you want to ensure that the method always runs within a new transaction.

### Mandatory **Attribute**

If the client is running within a transaction and invokes the enterprise bean's method, the method executes within the client's transaction. If the client is not associated with a transaction, the container throws the `TransactionRequiredException`.

Use the `Mandatory` attribute if the enterprise bean's method must use the transaction of the client.

### NotSupported **Attribute**

If the client is running within a transaction and invokes the enterprise bean's method, the container suspends the client's transaction before invoking the method. After the method has completed, the container resumes the client's transaction.

If the client is not associated with a transaction, the container does not start a new transaction before running the method.

Use the `NotSupported` attribute for methods that don't need transactions. Because transactions involve overhead, this attribute may improve performance.

### Supports **Attribute**

If the client is running within a transaction and invokes the enterprise bean's method, the method executes within the client's transaction. If the client is not associated with a transaction, the container does not start a new transaction before running the method.

Because the transactional behavior of the method may vary, you should use the `Supports` attribute with caution.

## Never Attribute

If the client is running within a transaction and invokes the enterprise bean's method, the container throws a `RemoteException`. If the client is not associated with a transaction, the container does not start a new transaction before running the method.

## Summary of Transaction Attributes

[Table 31-1](#) summarizes the effects of the transaction attributes. Both the T1 and the T2 transactions are controlled by the container. A T1 transaction is associated with the client that calls a method in the enterprise bean. In most cases, the client is another enterprise bean. A T2 transaction is started by the container just before the method executes.

In the last column of [Table 31-1](#), the word *None* means that the business method does not execute within a transaction controlled by the container. However, the database calls in such a business method might be controlled by the transaction manager of the DBMS.

TABLE 31-1 Transaction Attributes and Scope

Transaction Attribute	Client's Transaction	Business Method's Transaction
Required	None	T2
	T1	T1
RequiresNew	None	T2
	T1	T2
Mandatory	None	error
	T1	T1
NotSupported	None	None
	T1	None
Supports	None	None
	T1	T1
Never	None	None
	T1	Error

## Setting Transaction Attributes

Transaction attributes are specified by decorating the enterprise bean class or method with a `javax.ejb.TransactionAttribute` annotation, and setting it to one of the `javax.ejb.TransactionAttributeType` constants.

If you decorate the enterprise bean class with `@TransactionAttribute`, the specified `TransactionAttributeType` is applied to all the business methods in the class. Decorating a business method with `@TransactionAttribute` applies the `TransactionAttributeType` only to that method. If a `@TransactionAttribute` annotation decorates both the class and the method, the method `TransactionAttributeType` overrides the class `TransactionAttributeType`.

The `TransactionAttributeType` constants encapsulate the transaction attributes described earlier in this section.

- `Required`: `TransactionAttributeType.REQUIRED`
- `RequiresNew`: `TransactionAttributeType.REQUIRES_NEW`
- `Mandatory`: `TransactionAttributeType.MANDATORY`
- `NotSupported`: `TransactionAttributeType.NOT_SUPPORTED`
- `Supports`: `TransactionAttributeType.SUPPORTS`
- `Never`: `TransactionAttributeType.NEVER`

The following code snippet demonstrates how to use the `@TransactionAttribute` annotation:

```
@TransactionAttribute(NOT_SUPPORTED)
@Stateful
public class TransactionBean implements Transaction {
    ...
    @TransactionAttribute(REQUIRES_NEW)
    public void firstMethod() {...}

    @TransactionAttribute(REQUIRED)
    public void secondMethod() {...}

    public void thirdMethod() {...}

    public void fourthMethod() {...}
}
```

In this example, the `TransactionBean` class's transaction attribute has been set to `NotSupported`. `firstMethod` has been set to `RequiresNew`, and `secondMethod` has been set to `Required`. Because a `@TransactionAttribute` set on a method overrides the class `@TransactionAttribute`, calls to `firstMethod` will create a new transaction, and calls to `secondMethod` will either run in the current transaction, or start a new transaction. Calls to `thirdMethod` or `fourthMethod` do not take place within a transaction.

## Rolling Back a Container-Managed Transaction

There are two ways to roll back a container-managed transaction. First, if a system exception is thrown, the container will automatically roll back the transaction. Second, by invoking the `setRollbackOnly` method of the `EJBContext` interface, the bean method instructs the

container to roll back the transaction. If the bean throws an application exception, the rollback is not automatic but can be initiated by a call to `setRollbackOnly`.

## Synchronizing a Session Bean's Instance Variables

The `SessionSynchronization` interface, which is optional, allows stateful session bean instances to receive transaction synchronization notifications. For example, you could synchronize the instance variables of an enterprise bean with their corresponding values in the database. The container invokes the `SessionSynchronization` methods (`afterBegin`, `beforeCompletion`, and `afterCompletion`) at each of the main stages of a transaction.

The `afterBegin` method informs the instance that a new transaction has begun. The container invokes `afterBegin` immediately before it invokes the business method.

The container invokes the `beforeCompletion` method after the business method has finished, but just before the transaction commits. The `beforeCompletion` method is the last opportunity for the session bean to roll back the transaction (by calling `setRollbackOnly`).

The `afterCompletion` method indicates that the transaction has completed. It has a single boolean parameter whose value is `true` if the transaction was committed and `false` if it was rolled back.

## Methods Not Allowed in Container-Managed Transactions

You should not invoke any method that might interfere with the transaction boundaries set by the container. The list of prohibited methods follows:

- The `commit`, `setAutoCommit`, and `rollback` methods of `java.sql.Connection`
- The `getUserTransaction` method of `javax.ejb.EJBContext`
- Any method of `javax.transaction.UserTransaction`

You can, however, use these methods to set boundaries in application-managed transactions.

## Bean-Managed Transactions

In *bean-managed transaction demarcation*, the code in the session or message-driven bean explicitly marks the boundaries of the transaction. Although beans with container-managed transactions require less coding, they have one limitation: When a method is executing, it can be associated with either a single transaction or no transaction at all. If this limitation will make coding your bean difficult, you should consider using bean-managed transactions.

The following pseudocode illustrates the kind of fine-grained control you can obtain with application-managed transactions. By checking various conditions, the pseudocode decides whether to start or stop different transactions within the business method.

```
begin transaction
...
  update table-a
...
  if (condition-x)
    commit transaction
  else if (condition-y)
    update table-b
    commit transaction
  else
    rollback transaction
begin transaction
  update table-c
  commit transaction
```

When coding a application-managed transaction for session or message-driven beans, you must decide whether to use Java™ DataBase Connectivity (JDBC™) or JTA transactions. The sections that follow discuss both types of transactions.

## JTA Transactions

JTA is the abbreviation for the Java Transaction API. This API allows you to demarcate transactions in a manner that is independent of the transaction manager implementation. Sun GlassFish™ Enterprise Server v3 implements the transaction manager with the Java Transaction Service (JTS). But your code doesn't call the JTS methods directly. Instead, it invokes the JTA methods, which then call the lower-level JTS routines.

A *JTA transaction* is controlled by the Java EE transaction manager. You may want to use a JTA transaction because it can span updates to multiple databases from different vendors. A particular DBMS's transaction manager may not work with heterogeneous databases. However, the Java EE transaction manager does have one limitation: it does not support nested transactions. In other words, it cannot start a transaction for an instance until the preceding transaction has ended.

To demarcate a JTA transaction, you invoke the `begin`, `commit`, and `rollback` methods of the `javax.transaction.UserTransaction` interface.

## Returning without Committing

In a stateless session bean with bean-managed transactions, a business method must commit or roll back a transaction before returning. However, a stateful session bean does not have this restriction.

In a stateful session bean with a JTA transaction, the association between the bean instance and the transaction is retained across multiple client calls. Even if each business method called by the client opens and closes the database connection, the association is retained until the instance completes the transaction.

In a stateful session bean with a JDBC transaction, the JDBC connection retains the association between the bean instance and the transaction across multiple calls. If the connection is closed, the association is not retained.

## Methods Not Allowed in Bean-Managed Transactions

Do not invoke the `getRollbackOnly` and `setRollbackOnly` methods of the `EJBContext` interface in bean-managed transactions. These methods should be used only in container-managed transactions. For bean-managed transactions, invoke the `getStatus` and `rollback` methods of the `UserTransaction` interface.

## Transaction Timeouts

For container-managed transactions, you can use the Admin Console to configure the transaction timeout interval. See [“Starting the Administration Console” on page 64](#).

1. In the Admin Console, expand the Configuration node and select Transaction Service.
2. On the Transaction Service page, set the value of the Transaction Timeout field to the value of your choice (for example, 5).

With this setting, if the transaction has not completed within 5 seconds, the EJB container rolls it back.

The default value is 0, meaning that the transaction will not time out.

3. Click Save.

For enterprise beans with bean-managed JTA transactions, you invoke the `setTransactionTimeout` method of the `UserTransaction` interface.

## Updating Multiple Databases

The Java EE transaction manager controls all enterprise bean transactions except for bean-managed JDBC transactions. The Java EE transaction manager allows an enterprise bean to update multiple databases within a transaction. The figures that follow show two scenarios for updating multiple databases in a single transaction.

In [Figure 31-2](#), the client invokes a business method in Bean-A. The business method begins a transaction, updates Database X, updates Database Y, and invokes a business method in Bean-B. The second business method updates Database Z and returns control to the business method in Bean-A, which commits the transaction. All three database updates occur in the same transaction.

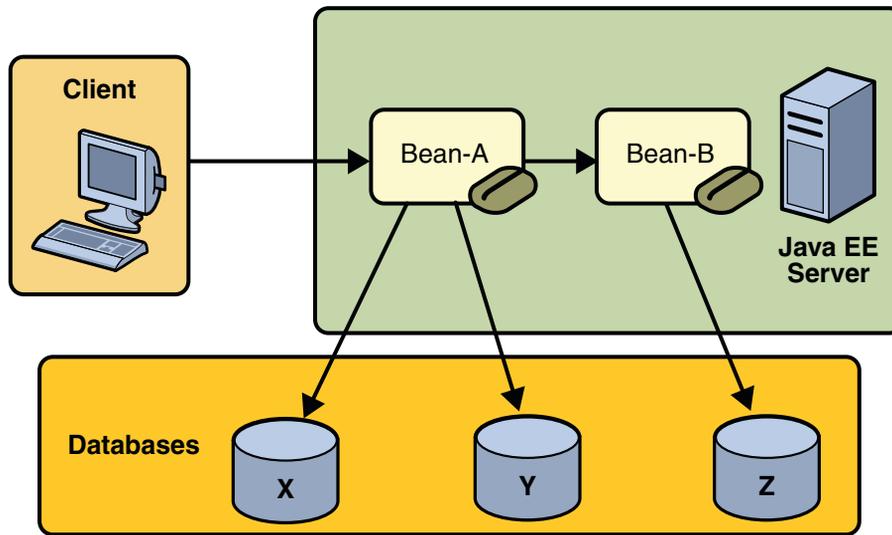


FIGURE 31-2 Updating Multiple Databases

In [Figure 31-3](#), the client calls a business method in Bean-A, which begins a transaction and updates Database X. Then Bean-A invokes a method in Bean-B, which resides in a remote Java EE server. The method in Bean-B updates Database Y. The transaction managers of the Java EE servers ensure that both databases are updated in the same transaction.

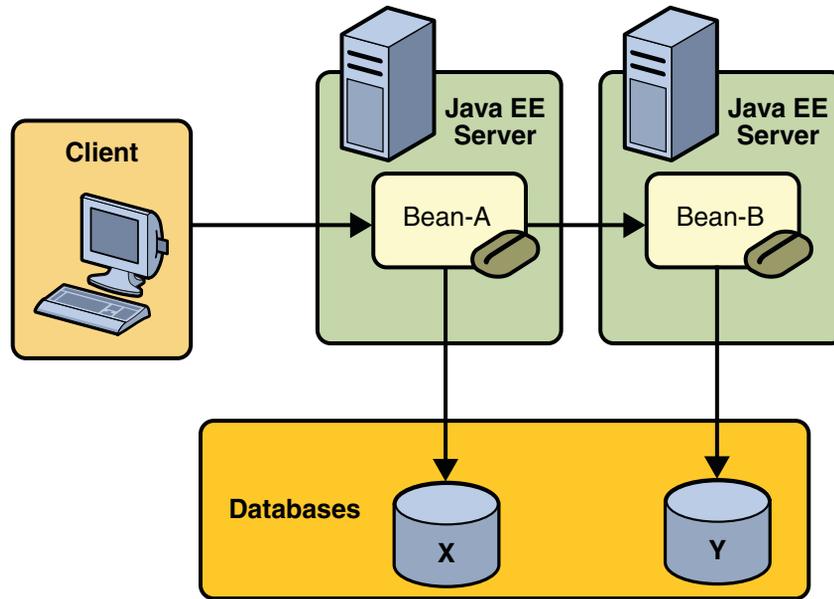


FIGURE 31-3 Updating Multiple Databases across Java EE Servers

## Transactions in Web Components

You can demarcate a transaction in a web component by using either the `java.sql.Connection` or `javax.transaction.UserTransaction` interface. These are the same interfaces that a session bean with bean-managed transactions can use. Transactions demarcated with the `UserTransaction` interface are discussed in the section “[JTA Transactions](#)” on page 588.



# Resource Connections

---

Java™ EE components can access a wide variety of resources, including databases, mail sessions, Java Message Service objects, and URLs. The Java EE 6 platform provides mechanisms that allow you to access all these resources in a similar manner. This chapter describes how to get connections to several types of resources.

The following topics are addressed here:

- “Resources and JNDI Naming” on page 593
- “DataSource Objects and Connection Pools” on page 594
- “Resource Injection” on page 595
- “Resource Adapters” on page 598
- “Metadata Annotations” on page 603
- “Replacing Deployment Descriptors With Metadata Annotations” on page 604
- “Common Client Interface” on page 608
- “Further Information about Resources” on page 609

## Resources and JNDI Naming

In a distributed application, components need to access other components and resources such as databases. For example, a servlet might invoke remote methods on an enterprise bean that retrieves information from a database. In the Java EE platform, the Java Naming and Directory Interface™ (JNDI) naming service enables components to locate other components and resources.

A resource is a program object that provides connections to systems, such as database servers and messaging systems. (A Java DataBase Connectivity (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.

For example, the JNDI name of the JDBC resource for the Java DB database that is shipped with Sun GlassFish™ Enterprise Server v3 (Enterprise Server) is `jdbc/_default`.

An administrator creates resources in a JNDI namespace. In the Enterprise Server, you can use either the Administration Console or the `asadmin` command to create resources. Applications then use annotations to inject the resources. If an application uses resource injection, the Enterprise Server invokes the JNDI API, and the application is not required to do so. However, it is also possible for an application to locate resources by making direct calls to the JNDI API.

A resource object and its JNDI name are bound together by the naming and directory service. To create a new resource, a new name-object binding is entered into the JNDI namespace. You inject resources by using the `@Resource` annotation in an application.

You can use a deployment descriptor to override the resource mapping that you specify in an annotation. Using a deployment descriptor allows you to change an application by repackaging it, rather than by both recompiling the source files and repackaging. However, for most applications, a deployment descriptor is not necessary.

## DataSource **Objects and Connection Pools**

To store, organize, and retrieve data, most applications use a relational database. Java EE 6 components may access relational databases through the JDBC API. For information on this API, see <http://java.sun.com/javase/technologies/database/>.

In the JDBC API, databases are accessed by using DataSource objects. A DataSource has a set of properties that identify and describe the real world data source that it represents. These properties include information such as the location of the database server, the name of the database, the network protocol to use to communicate with the server, and so on. In the Sun GlassFish Enterprise Server, a data source is called a JDBC resource.

Applications access a data source using a connection, and a DataSource object can be thought of as a factory for connections to the particular data source that the DataSource instance represents. In a basic DataSource implementation, a call to the `getConnection` method returns a connection object that is a physical connection to the data source.

If a DataSource object is registered with a JNDI naming service, an application can use the JNDI API to access that DataSource object, which can then be used to connect to the data source it represents.

DataSource objects that implement connection pooling also produce a connection to the particular data source that the DataSource class represents. The connection object that the `getConnection` method returns is a handle to a `PooledConnection` object rather than being a physical connection. An application uses the connection object in the same way that it uses a connection. Connection pooling has no effect on application code except that a pooled connection, like all connections, should always be explicitly closed. When an application closes a connection that is pooled, the connection is returned to a pool of reusable connections. The next time `getConnection` is called, a handle to one of these pooled connections will be returned if one is available. Because connection pooling avoids creating a new physical connection every time one is requested, applications can run significantly faster.

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.

Applications that use the Persistence API specify the `DataSource` object they are using in the `jta-data-source` element of the `persistence.xml` file.

```
<jta-data-source>jdbc/MyOrderDB</jta-data-source>
```

This is typically the only reference to a JDBC object for a persistence unit. The application code does not refer to any JDBC objects.

## Resource Injection

The `javax.annotation.Resource` annotation is used to declare a reference to a resource. `@Resource` can decorate a class, a field, or a method. The container will inject the resource referred to by `@Resource` into the component either at runtime or when the component is initialized, depending on whether field/method injection or class injection is used. With field and method-based injection, the container will inject the resource when the application is initialized. For class-based injection, the resource is looked up by the application at runtime.

`@Resource` has the following elements:

- `name`: The JNDI name of the resource
- `type`: The Java language type of the resource
- `authenticationType`: The authentication type to use for the resource
- `shareable`: Indicates whether the resource can be shared
- `mappedName`: A non-portable, implementation-specific name to which the resource should be mapped
- `description`: The description of the resource

The `name` element is the JNDI name of the resource, and is optional for field- and method-based injection. For field-based injection, the default name is the field name qualified by the class name. For method-based injection, the default name is the `JavaBeans™` property name based on the method qualified by the class name. The `name` element must be specified for class-based injection.

The type of resource is determined by one of the following:

- The type of the field the `@Resource` annotation is decorating for field-based injection
- The type of the JavaBeans property the `@Resource` annotation is decorating for method-based injection
- The type element of `@Resource`

For class-based injection, the type element is required.

The `authenticationType` element is used only for connection factory resources, such as the resources of a connector (also called the resource adapter or data source). This element can be set to one of the `javax.annotation.Resource.AuthenticationType` enumerated type values: `CONTAINER`, the default, and `APPLICATION`.

The `shareable` element is used only for Object Resource Broker (ORB) instance resources or connection factory resource. It indicates whether the resource can be shared between this component and other components, and may be set to `true`, the default, or `false`.

The `mappedName` element is a non-portable, implementation-specific name to which the resource should be mapped. Because the name element, when specified or defaulted, is local only to the application, many Java EE servers provide a way of referring to resources across the application server. This is done by setting the `mappedName` element. Use of the `mappedName` element is non-portable across Java EE server implementations.

The `description` element is the description of the resource, typically in the default language of the system on which the application is deployed. It is used to help identify resources, and to help application developers choose the correct resource.

## Field-Based Injection

To use field-based resource injection, declare a field and decorate it with the `@Resource` annotation. The container will infer the name and type of the resource if the name and type elements are not specified. If you do specify the type element, it must match the field's type declaration.

```
package com.example;

public class SomeClass {
    @Resource
    private javax.sql.DataSource myDB;
    ...
}
```

In the code above, the container infers the name of the resource based on the class name and the field name: `com.example.SomeClass/myDB`. The inferred type is `javax.sql.DataSource.class`.

```
package com.example;

public class SomeClass {
    @Resource(name="customerDB")
    private javax.sql.DataSource myDB;
    ...
}
```

In the code above, the JNDI name is `customerDB`, and the inferred type is `javax.sql.DataSource.class`.

## Method-Based Injection

To use method-based injection, declare a setter method and decorate it with the `@Resource` annotation. The container will infer the name and type of the resource if the name and type elements are not specified. The setter method must follow the JavaBeans conventions for property names: the method name must begin with `set`, have a `void` return type, and only one parameter. If you do specify the type element, it must match the field's type declaration.

```
package com.example;

public class SomeClass {

    private javax.sql.DataSource myDB;
    ...
    @Resource
    private void setMyDB(javax.sql.DataSource ds) {
        myDB = ds;
    }
    ...
}
```

In the code above, the container infers the name of the resource based on the class name and the field name: `com.example.SomeClass/myDB`. The inferred type is `javax.sql.DataSource.class`.

```
package com.example;

public class SomeClass {

    private javax.sql.DataSource myDB;
    ...
    @Resource(name="customerDB")
    private void setMyDB(javax.sql.DataSource ds) {
        myDB = ds;
    }
}
```

```
...  
}
```

In the code above, the JNDI name is `customerDB`, and the inferred type is `javax.sql.DataSource.class`.

## Class-Based Injection

To use class-based injection, decorate the class with a `@Resource` annotation, and set the required name and type elements.

```
@Resource(name="myMessageQueue",  
          type="javax.jms.ConnectionFactory")  
public class SomeMessageBean {  
    ...  
}
```

## Declaring Multiple Resources

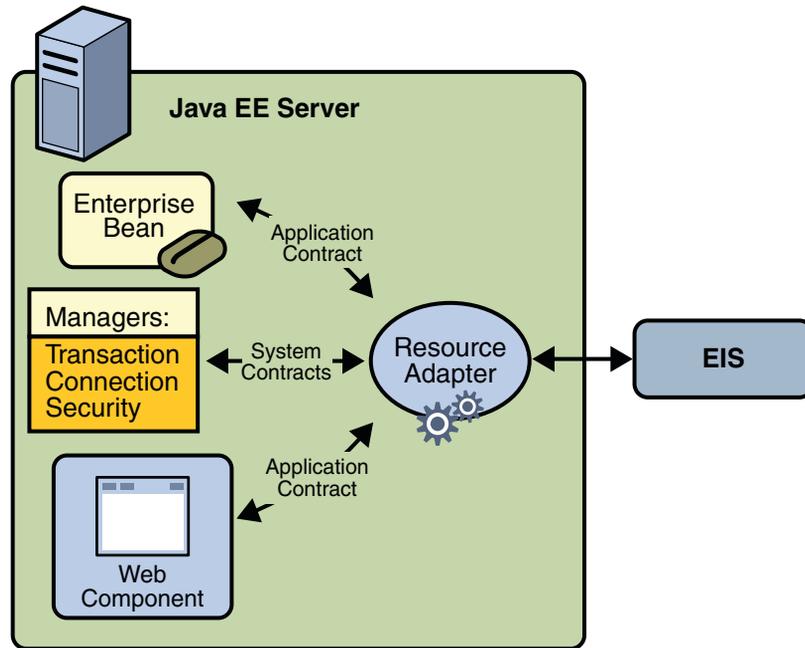
The `@Resources` annotation is used to group together multiple `@Resource` declarations for class-based injection.

```
@Resources({  
    @Resource(name="myMessageQueue",  
              type="javax.jms.ConnectionFactory"),  
    @Resource(name="myMailSession",  
              type="javax.mail.Session")  
})  
public class SomeMessageBean {  
    ...  
}
```

The code above shows the `@Resources` annotation containing two `@Resource` declarations. One is a Java Message Service (JMS) message queue, and the other is a JavaMail™ session.

## Resource Adapters

A resource adapter is a Java EE component that implements the Java EE Connector Architecture for a specific Enterprise Information System (EIS). Examples of EISs include Enterprise Resource Planning (ERP), mainframe transaction processing (TP), and database systems. As illustrated in Figure 35-1, the resource adapter facilitates communication between a Java EE application and an EIS.



Stored in a Resource Adapter Archive (RAR) file, a resource adapter can be deployed on any Java EE server, much like a Java EE application. An RAR file may be contained in an Enterprise Archive (EAR) file, or it may exist as a separate file.

A resource adapter is analogous to a JDBC driver. Both provide a standard API through which an application can access a resource that is outside the Java EE server. For a resource adapter, the target system is an EIS; for a JDBC driver, it is a DBMS. Resource adapters and JDBC drivers are rarely created by application developers. In most cases, both types of software are built by vendors that sell products such as tools, servers, or integration software.

## Resource Adapter Contracts

The resource adapter mediates communication between the Java EE server and the EIS by means of contracts. The application contract defines the API through which a Java EE component such as an enterprise bean accesses the EIS. This API is the only view that the component has of the EIS. The system contracts link the resource adapter to important services that are managed by the Java EE server. The resource adapter itself and its system contracts are transparent to the Java EE component.

### Management Contracts

The Java EE Connector Architecture defines system contracts that enable resource adapter life cycle and thread management.

## Life cycle Management

The Connector Architecture specifies a life cycle management contract that allows an application server to manage the life cycle of a resource adapter. This contract provides a mechanism for the application server to bootstrap a resource adapter instance during the deployment or application server startup. It also provides a means for the application server to notify the resource adapter instance when it is undeployed or when an orderly shutdown of the application server takes place.

## Work Management Contract

The Connector Architecture work management contract ensures that resource adapters use threads in the proper, recommended manner. It also enables an application server to manage threads for resource adapters.

Resource adapters that improperly use threads can jeopardize the entire application server environment. For example, a resource adapter might create too many threads or it might not properly release threads it has created. Poor thread handling inhibits application server shutdown. It also impacts the application server's performance because creating and destroying threads are expensive operations.

The work management contract establishes a means for the application server to pool and reuse threads, similar to pooling and reusing connections. By adhering to this contract, the resource adapter does not have to manage threads itself. Instead, the resource adapter has the application server create and provide needed threads. When the resource adapter is finished with a given thread, it returns the thread to the application server. The application server manages the thread: It can return the thread to a pool and reuse it later, or it can destroy the thread. Handling threads in this manner results in increased application server performance and more efficient use of resources.

In addition to moving thread management to the application server, the Connector Architecture provides a flexible model for a resource adapter that uses threads:

- The requesting thread can choose to block (stop its own execution) until the work thread completes.
- Or the requesting thread can block while it waits to get the work thread. When the application server provides a work thread, the requesting thread and the work thread execute in parallel.
- The resource adapter can opt to submit the work for the thread to a queue. The thread executes the work from the queue at some later point. The resource adapter continues its own execution from the point it submitted the work to the queue, no matter when the thread executes it.

With the latter two approaches, the submitting thread and the work thread may execute simultaneously or independently from each other. For these approaches, the contract specifies a listener mechanism to notify the resource adapter that the thread has completed its operation.

The resource adapter can also specify the execution context for the thread, and the work management contract controls the context in which the thread executes.

## Generic Work Context Contract

A generic work context contract enables a resource adapter to control the contexts in which the Work instances submitted by it are executed by the application server's `WorkManager`. A Generic Work context mechanism also enables an application server to support new message inflow and delivery schemes. It also provides a richer contextual Work execution environment to the resource adapter while still maintaining control over concurrent behavior in a managed environment.

The Generic Work Context contract standardizes the following contexts: Transaction Context and Security Context.

## Transaction Context

The transaction context between the resource adapter and the application server leverages the Generic Work Context mechanism by describing a standard `WorkContext`, the `TransactionContext`. It represents the standard interface that a resource adapter can use to propagate transaction context information from the EIS to the application server.

## Security Context

The security context between the resource adapter and the application server leverages the Generic Work Context mechanism by describing a standard `WorkContext`, the `SecurityContext`, that can be provided by the resource adapter while submitting a Work for execution.

The `SecurityContext` provides a portable mechanism for the resource adapter to pass security context information to the application server. This work context enables an EIS or resource adapter to flow-in security context information while submitting a Work to a `WorkContext` for execution.

## Work Security Map

A work security map matches EIS identities to the application server domain's identities.

## Hints Context

The propagation of Quality of Service hints to a `WorkManager` for the execution of a Work instance is standardized through the `HintsContext` class. It provides a mechanism for the resource adapter to pass quality-of-service metadata to the `WorkManager` during the submission of a Work instance. The application server can use the specified hints to control the execution of the Work instance.

## Outbound Contracts

The Connector Architecture defines system-level contracts between an application server and an EIS that enable outbound connectivity to an EIS: connection management, transaction management, and security.

### Connection Management Contract

The connection management contract supports connection pooling, a technique that enhances application performance and scalability. Connection pooling is transparent to the application, which simply obtains a connection to the EIS.

### Transaction Management Contract

The Connector Architecture supports the concept of transactions - a number of operations that must be committed together or not at all for the data to remain consistent and to maintain data integrity.

A local transaction is limited in scope to a single EIS system, and the EIS resource manager itself manages such transaction. An XA transaction or global transaction can span multiple resource managers. This form of transaction requires transaction coordination by an external transaction manager, typically bundled with an application server. A transaction manager uses a two-phase commit protocol to manage a transaction that spans multiple resource managers or EISs. It uses one-phase commit optimization if only one resource manager is participating in an XA transaction.

The Connector Architecture defines a transaction management contract between an application server and a resource adapter . The transaction management contract extends the connection management contract and provides support for management of both local and XA transactions. These contracts enable an application server to provide the infrastructure and runtime environment for transaction management. Application components rely on this transaction infrastructure to support the component-level transaction model.

An application server is required to support all three levels of transactions:

- No transaction support at all - this is typical of legacy applications and back-end systems.
- Support for only local transactions
- Support for both local and XA transactions

### Security Management Contract

The security management contract provides mechanisms for authentication, authorization, and secure communication between a Java EE server and an EIS to protect the information in the EIS.

## Inbound Contracts

The Java EE Connector Architecture defines system contracts between a Java EE server and an EIS that enable inbound connectivity from the EIS: pluggability contracts for message providers and contracts for importing transactions.

## Messaging Contracts

To enable external systems to connect to a Java EE application server, the Connector Architecture extends the capabilities of message-driven beans to handle messages from any message provider. That is, message-driven beans are no longer limited to handling JMS messages. Instead, EISs and message providers can plug any message provider, including their own custom or proprietary message providers, into a Java EE server.

To provide this feature, a message provider or an EIS resource adapter implements the messaging contract, which details APIs for message handling and message delivery. A conforming resource adapter is assured of the ability to send messages from any provider to a message-driven bean, and it also can be plugged into a Java EE server in a standard manner.

## Transaction Inflow

The Connector Architecture supports importing transactions from an EIS to a Java EE server. The architecture specifies how to propagate the transaction context from the EIS. For example, a transaction can be started by the EIS, such as the Customer Information Control System (CICS). Within the same CICS transaction, a connection can be made through a resource adapter to an enterprise bean on the application server. The enterprise bean does its work under the CICS transaction context and commits within that transaction context.

The Connector Architecture also specifies how the container participates in transaction completion and how it handles crash recovery to ensure that data integrity is not lost.

# Metadata Annotations

Java EE Connector Architecture 1.6 introduces a set of annotations to minimize the need for deployment descriptors.

- The `@Connector` annotation can be used by the resource adapter developer to specify that the JavaBeans component is a resource adapter JavaBeans component. This annotation is used for providing metadata about the capabilities of the resource adapter. Optionally, you can provide a JavaBeans component implementing the `ResourceAdapter` interface.
- The `@ConnectionDefinition` annotation defines a set of connection interfaces and classes pertaining to a particular connection type. The role of this annotation is identical to the role played by the connection-definition element in the deployment descriptor.
- The `@AdministeredObject` annotation designates a JavaBeans component as an administered object.

- The `@Activation` annotation contains the configuration information pertaining to inbound connectivity from an EIS instance.
- The `@ConfigProperty` annotation can be used on JavaBeans components to indicate to the application server that a specific JavaBeans property is a configuration property for that JavaBeans component. A configuration property may be used by the deployer and resource adapter provider to provide additional configuration information. The application server provides configuration tools to automatically discover the configuration properties of a JavaBeans component through JavaBeans introspection and need not be specified by using a deployment descriptor.

The specification allows a resource adapter to be developed in mixed-mode form, that is the ability for a resource adapter developer to use both metadata annotations and deployment descriptors in applications. An application assembler or deployer may use the deployment descriptor to override the metadata annotations specified by the resource adapter developer.

A new attribute, `metadata-complete`, is introduced in the Connector 1.6 deployment descriptor (the `ra.xml` file). The `metadata-complete` attribute defines whether the deployment descriptor for the resource adapter module is complete, or whether the class files available to the module and packaged with the resource adapter need to be examined for annotations that specify deployment information.

For the complete list of annotations and JavaBeans components introduced in Java EE 6, see <http://jcp.org/en/jsr/detail?id=322>.

## Replacing Deployment Descriptors With Metadata Annotations

The use of annotations reduces or completely eliminates the need to deal with a deployment descriptor in many cases. The use of annotations also reduces the need to keep the deployment descriptor synchronized with changes to source code.

The following examples provide sample deployment descriptors and equivalent annotations.

### Example 1: `@Connector` Annotation

The following deployment descriptor defines a connector:

```
<connector xmlns="http://java.sun.com/xml/ns/j2ee"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://java.sun.com/xml/ns/j2ee
http://java.sun.com/xml/ns/j2ee/connector_1_6.xsd"
version="1.6">
```

```

<description>Sample adapter using the JavaMail API</description>
<display-name>InboundResourceAdapter</display-name>
<icon></icon>
<vendor-name>Sun Microsystems, Inc</vendor-name>

<eis-type>MAIL</eis-type>
<resourceadapter-version>1.0</resourceadapter-version>
...
...
...
<authentication-mechanism>
  <authentication-mechanism-type>BasicPassword</authentication-mechanism-type>
<credential-interface>javax.resource.spi.security.PasswordCredential</credential-interface>
</authentication-mechanism>
<reauthentication-support>false</reauthentication-support>
...
...
</connector>

```

The equivalent metadata annotation is as follows:

```

@Connector(
  description = "Sample adapter using the JavaMail API",
  displayName = "InboundResourceAdapter",
  vendorName = "Sun Microsystems, Inc.",
  eisType = "MAIL",
  version = "1.0",
  authMechanisms = {
    @AuthenticationMechanism(
      authMechanism = "BasicPassword",
      credentialInterface = AuthenticationMechanism.CredentialInterface.PasswordCredential
    )
  }
)

/*
// Since the following attribute values denote the default values of the annotation,
// they need not be specified explicitly

  transactionSupport = TransactionSupport.TransactionSupportLevel.NoTransaction,
  reauthenticationSupport = false
*/
)

public class ResourceAdapterImpl
implements ResourceAdapter, java.io.Serializable
{
  ...
  ...
}

```

## Example 2: @ConnectionDefinition Annotation

The following deployment descriptor snippet describes a connection definition:

```
<connection-definition>
<managedconnectionfactory-class>
samples.mailra.ra.outbound.ManagedConnectionFactoryImpl
</managedconnectionfactory-class>
<config-property><config-property-name>serverName</config-property-name>
<config-property-type>java.lang.String</config-property-type>
<config-property-value>UnknownHostName</config-property-value>

...
...
<connectionfactory-interface>samples.mailra.api.JavaMailConnectionFactory</connectionfactory-interface>
<connectionfactory-impl-class>samples.mailra.ra.outbound.JavaMailConnectionFactoryImpl</connectionfactory-impl-class>
<connection-interface>samples.mailra.api.JavaMailConnection</connection-interface>
<credential-interface>javax.resource.spi.security.PasswordCredential</credential-interface>
<connection-impl-class>samples.mailra.ra.outbound.JavaMailConnectionImpl</connection-impl-class>
</connectiondefinition>
```

The equivalent metadata annotation is as follows:

```
@ConnectionDefinition(
    connectionFactory =
        samples.mailra..api.JavaMailConnectionFactory.class,
    connectionFactoryImpl =
        samples.mailra.ra.outbound.JavaMailConnectionFactoryImpl.class,
    connection = samples.connectors.mailconnector.api.JavaMailConnection.class,
    connectionImpl =
        samples.mailra..ra.outbound.JavaMailConnectionImpl.class
)
public class ManagedConnectionFactoryImpl implements
    ManagedConnectionFactory, Serializable
{
    ...
    ...
    @ConfigProperty(
        defaultValue = "UnknownHostName"
    )
    public void setServerName(String serverName)
    {
        ...
    }
}
```

## Example 3: @Activation Annotation

The following deployment descriptor snippet describes the messaging capabilities of a resource adapter:

```
<messageadapter>
<messagelistener>
<messagelistener-type>samples.mailra.api.JavaMailMessageListener</messagelistener-type>
<activationspec>
<activationspec-class>samples.mailra.ra.inbound.ActivationSpecImpl</activationspec-class>
  <required-config-property>
    <config-property-name>serverName</config-property-name>
  </required-config-property>
  <required-config-property>
    <config-property-name>userName</config-property-value>
  </required-config-property>
  <required-config-property>
    <config-property-name>password</config-property-value>
  </required-config-property>
  <required-config-property>
    <config-property-name>folderName</config-property-value>
  </required-config-property>
  <required-config-property>
    <description>Normally imap or pop3</description>
    <config-property-name>protocol</config-property-name>
    <config-property-value>IMAP</config-property-value>
  </required-config-property>
</activationspec>
</messagelistener>
</messageadapter>
```

The equivalent metadata annotation is as follows:

```
@Activation(
    messageListeners = {samples.mailra.api.JavaMailMessageListener.class}
)
public class ActivationSpecImpl implements javax.resource.spi.ActivationSpec,
    java.io.Serializable
{
    ...
    @ConfigProperty()
    // serverName property value
    private String serverName = new String("");

    @ConfigProperty()
    // userName property value
    private String userName = new String("");
```

```
@ConfigProperty()  
// password property value  
private String password = new String("");  
  
@ConfigProperty()  
// folderName property value  
private String folderName = new String("Inbox");  
  
// protocol property value  
// Normally imap or pop3  
@ConfigProperty(  
    description = "Normally imap or pop3"  
)  
private String protocol = new String("imap");  
  
...  
...  
}
```

## Common Client Interface

This section describes how components use the Connector Architecture Common Client Interface (CCI) API and a resource adapter to access data from an EIS.

Defined by the Java EE Connector Architecture specification, the CCI defines a set of interfaces and classes whose methods allow a client to perform typical data access operations. The CCI interfaces and classes are as follows:

- **ConnectionFactory**: Provides an application component with a `Connection` instance to an EIS.
- **Connection**: Represents the connection to the underlying EIS.
- **ConnectionSpec**: Provides a means for an application component to pass connection-request-specific properties to the `ConnectionFactory` when making a connection request.
- **Interaction**: Provides a means for an application component to execute EIS functions, such as database stored procedures.
- **InteractionSpec**: Holds properties pertaining to an application component's interaction with an EIS.
- **Record**: The superclass for the various kinds of record instances. Record instances can be `MappedRecord`, `IndexedRecord`, or `ResultSet` instances, all of which inherit from the `Record` interface.
- **RecordFactory**: Provides an application component with a `Record` instance.

- `IndexedRecord`: Represents an ordered collection of `Record` instances based on the `java.util.List` interface.

A client or application component that uses the CCI to interact with an underlying EIS does so in a prescribed manner. The component must establish a connection to the EIS's resource manager, and it does so using the `ConnectionFactory`. The `Connection` object represents the actual connection to the EIS and is used for subsequent interactions with the EIS.

The component performs its interactions with the EIS, such as accessing data from a specific table, using an `Interaction` object. The application component defines the `Interaction` object using an `InteractionSpec` object. When the application component reads data from the EIS (such as from database tables) or writes to those tables, it does so using a particular type of `Record` instance: either a `MappedRecord`, an `IndexedRecord`, or a `ResultSet` instance. Just as the `ConnectionFactory` creates `Connection` instances, a `RecordFactory` creates `Record` instances.

Note, too, that a client application that relies on a CCI resource adapter is very much like any other Java EE client that uses enterprise bean methods.

## Further Information about Resources

For more information about resources and annotations, see:

- The Java EE 6 Platform Specification (JSR 316):  
<http://jcp.org/en/jsr/detail?id=316>
- The Java EE Connector Architecture Specification 1.6 (JSR 322):  
<http://jcp.org/en/jsr/detail?id=322>
- EJB 3.1 specification (JSR 220):  
<http://jcp.org/en/jsr/detail?id=318>
- Common Annotations for the Java Platform (JSR 250):  
<http://www.jcp.org/en/jsr/detail?id=250>



# Java Message Service Concepts

---

This chapter provides an introduction to the Java Message Service (JMS) API, a Java API that allows applications to create, send, receive, and read messages using reliable, asynchronous, loosely coupled communication. It covers the following topics:

- “Overview of the JMS API” on page 611
- “Basic JMS API Concepts” on page 614
- “The JMS API Programming Model” on page 618
- “Creating Robust JMS Applications” on page 628
- “Using the JMS API in Java EE Applications” on page 638
- “Further Information about JMS” on page 644

## Overview of the JMS API

This overview of the JMS API answers the following questions.

- “What Is Messaging?” on page 611
- “What Is the JMS API?” on page 612
- “When Can You Use the JMS API?” on page 612
- “How Does the JMS API Work with the Java EE Platform?” on page 613

## What Is Messaging?

Messaging is a method of communication between software components or applications. A messaging system is a peer-to-peer facility: A messaging client can send messages to, and receive messages from, any other client. Each client connects to a messaging agent that provides facilities for creating, sending, receiving, and reading messages.

Messaging enables distributed communication that is *loosely coupled*. A component sends a message to a destination, and the recipient can retrieve the message from the destination. However, the sender and the receiver do not have to be available at the same time in order to

communicate. In fact, the sender does not need to know anything about the receiver; nor does the receiver need to know anything about the sender. The sender and the receiver need to know only which message format and which destination to use. In this respect, messaging differs from tightly coupled technologies, such as Remote Method Invocation (RMI), which require an application to know a remote application's methods.

Messaging also differs from electronic mail (email), which is a method of communication between people or between software applications and people. Messaging is used for communication between software applications or software components.

## What Is the JMS API?

The Java Message Service is a Java API that allows applications to create, send, receive, and read messages. Designed by Sun and several partner companies, the JMS API defines a common set of interfaces and associated semantics that allow programs written in the Java programming language to communicate with other messaging implementations.

The JMS API minimizes the set of concepts a programmer must learn in order to use messaging products but provides enough features to support sophisticated messaging applications. It also strives to maximize the portability of JMS applications across JMS providers in the same messaging domain.

The JMS API enables communication that is not only loosely coupled but also

- **Asynchronous:** A JMS provider can deliver messages to a client as they arrive; a client does not have to request messages in order to receive them.
- **Reliable:** The JMS API can ensure that a message is delivered once and only once. Lower levels of reliability are available for applications that can afford to miss messages or to receive duplicate messages.

The JMS specification was first published in August 1998. The latest version is Version 1.1, which was released in April 2002. You can download a copy of the specification from the JMS web site: <http://java.sun.com/products/jms/>.

## When Can You Use the JMS API?

An enterprise application provider is likely to choose a messaging API over a tightly coupled API, such as remote procedure call (RPC), under the following circumstances.

- The provider wants the components not to depend on information about other components' interfaces, so that components can be easily replaced.
- The provider wants the application to run whether or not all components are up and running simultaneously.
- The application business model allows a component to send information to another and to continue to operate without receiving an immediate response.

For example, components of an enterprise application for an automobile manufacturer can use the JMS API in situations like these:

- The inventory component can send a message to the factory component when the inventory level for a product goes below a certain level so that the factory can make more cars.
- The factory component can send a message to the parts components so that the factory can assemble the parts it needs.
- The parts components in turn can send messages to their own inventory and order components to update their inventories and to order new parts from suppliers.
- Both the factory and the parts components can send messages to the accounting component to update their budget numbers.
- The business can publish updated catalog items to its sales force.

Using messaging for these tasks allows the various components to interact with one another efficiently, without tying up network or other resources. [Figure 33–1](#) illustrates how this simple example might work.

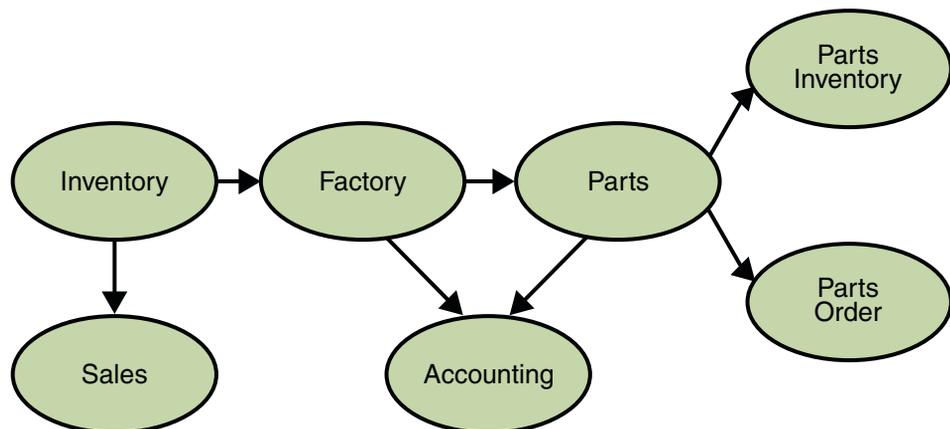


FIGURE 33–1 Messaging in an Enterprise Application

Manufacturing is only one example of how an enterprise can use the JMS API. Retail applications, financial services applications, health services applications, and many others can make use of messaging.

## How Does the JMS API Work with the Java EE Platform?

When the JMS API was introduced in 1998, its most important purpose was to allow Java applications to access existing messaging-oriented middleware (MOM) systems, such as

MQSeries from IBM. Since that time, many vendors have adopted and implemented the JMS API, so a JMS product can now provide a complete messaging capability for an enterprise.

Beginning with the 1.3 release of the Java EE platform, the JMS API has been an integral part of the platform, and application developers can use messaging with Java EE components.

The JMS API in the Java EE platform has the following features.

- Application clients, Enterprise JavaBeans (EJB) components, and web components can send or synchronously receive a JMS message. Application clients can in addition receive JMS messages asynchronously. (Applets, however, are not required to support the JMS API.)
- Message-driven beans, which are a kind of enterprise bean, enable the asynchronous consumption of messages. A JMS provider can optionally implement concurrent processing of messages by message-driven beans.
- Message send and receive operations can participate in distributed transactions, which allow JMS operations and database accesses to take place within a single transaction.

The JMS API enhances the Java EE platform by simplifying enterprise development, allowing loosely coupled, reliable, asynchronous interactions among Java EE components and legacy systems capable of messaging. A developer can easily add new behavior to a Java EE application that has existing business events by adding a new message-driven bean to operate on specific business events. The Java EE platform, moreover, enhances the JMS API by providing support for distributed transactions and allowing for the concurrent consumption of messages. For more information, see the Enterprise JavaBeans specification, v3.1.

The JMS provider can be integrated with the application server using the Java EE Connector architecture. You access the JMS provider through a resource adapter. This capability allows vendors to create JMS providers that can be plugged in to multiple application servers, and it allows application servers to support multiple JMS providers. For more information, see the Java EE Connector architecture specification, v1.6.

## Basic JMS API Concepts

This section introduces the most basic JMS API concepts, the ones you must know to get started writing simple application clients that use the JMS API:

- [“JMS API Architecture” on page 615](#)
- [“Messaging Domains” on page 616](#)
- [“Message Consumption” on page 618](#)

The next section introduces the JMS API programming model. Later sections cover more advanced concepts, including the ones you need to write applications that use message-driven beans.

## JMS API Architecture

A JMS application is composed of the following parts.

- A *JMS provider* is a messaging system that implements the JMS interfaces and provides administrative and control features. An implementation of the Java EE platform includes a JMS provider.
- *JMS clients* are the programs or components, written in the Java programming language, that produce and consume messages. Any Java EE application component can act as a JMS client.
- *Messages* are the objects that communicate information between JMS clients.
- *Administered objects* are preconfigured JMS objects created by an administrator for the use of clients. The two kinds of JMS administered objects are destinations and connection factories, which are described in “[JMS Administered Objects](#)” on page 619.

Figure 33–2 illustrates the way these parts interact. Administrative tools allow you to bind destinations and connection factories into a JNDI namespace. A JMS client can then use resource injection to access the administered objects in the namespace and then establish a logical connection to the same objects through the JMS provider.

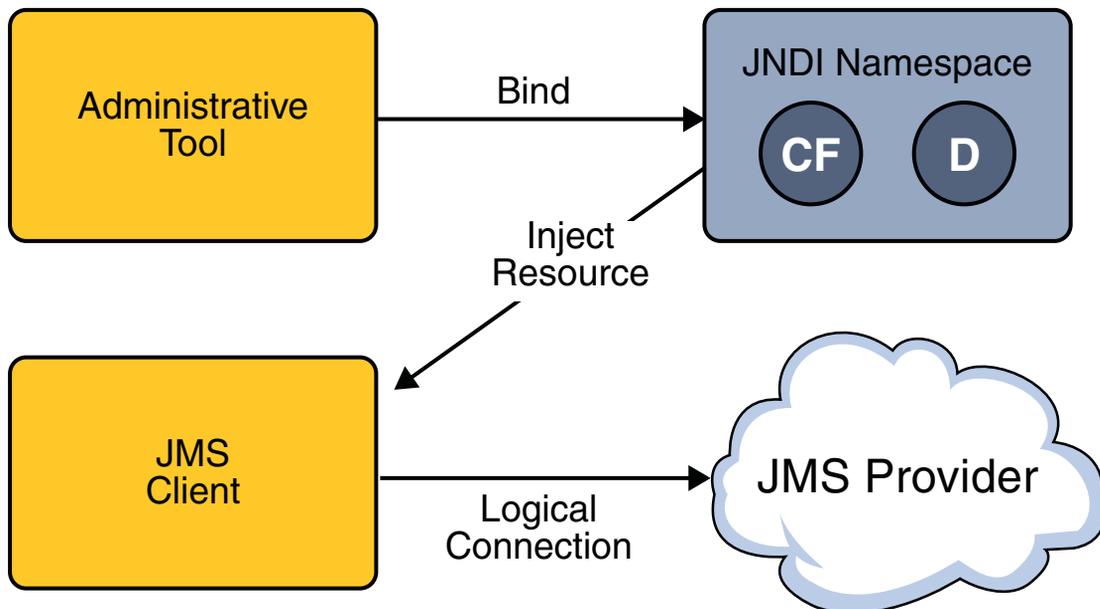


FIGURE 33–2 JMS API Architecture

## Messaging Domains

Before the JMS API existed, most messaging products supported either the *point-to-point* or the *publish/subscribe* approach to messaging. The JMS specification provides a separate domain for each approach and defines compliance for each domain. A stand-alone JMS provider can implement one or both domains. A Java EE provider must implement both domains.

In fact, most implementations of the JMS API support both the point-to-point and the publish/subscribe domains, and some JMS clients combine the use of both domains in a single application. In this way, the JMS API has extended the power and flexibility of messaging products.

The JMS specification goes one step further: It provides common interfaces that enable you to use the JMS API in a way that is not specific to either domain. The following subsections describe the two messaging domains and then describe the use of the common interfaces.

### Point-to-Point Messaging Domain

A point-to-point (PTP) product or application is built on the concept of message queues, senders, and receivers. Each message is addressed to a specific queue, and receiving clients extract messages from the queues established to hold their messages. Queues retain all messages sent to them until the messages are consumed or until the messages expire.

PTP messaging has the following characteristics and is illustrated in [Figure 33-3](#).

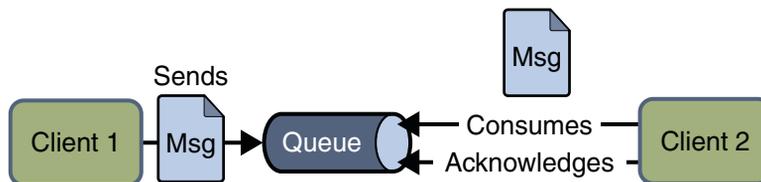


FIGURE 33-3 Point-to-Point Messaging

- Each message has only one consumer.
- A sender and a receiver of a message have no timing dependencies. The receiver can fetch the message whether or not it was running when the client sent the message.
- The receiver acknowledges the successful processing of a message.

Use PTP messaging when every message you send must be processed successfully by one consumer.

## Publish/Subscribe Messaging Domain

In a publish/subscribe (pub/sub) product or application, clients address messages to a *topic*, which functions somewhat like a bulletin board. Publishers and subscribers are generally anonymous and can dynamically publish or subscribe to the content hierarchy. The system takes care of distributing the messages arriving from a topic's multiple publishers to its multiple subscribers. Topics retain messages only as long as it takes to distribute them to current subscribers.

Pub/sub messaging has the following characteristics.

- Each message can have multiple consumers.
- Publishers and subscribers have a timing dependency. A client that subscribes to a topic can consume only messages published after the client has created a subscription, and the subscriber must continue to be active in order for it to consume messages.

The JMS API relaxes this timing dependency to some extent by allowing subscribers to create *durable subscriptions*, which receive messages sent while the subscribers are not active. Durable subscriptions provide the flexibility and reliability of queues but still allow clients to send messages to many recipients. For more information about durable subscriptions, see [“Creating Durable Subscriptions” on page 633](#).

Use pub/sub messaging when each message can be processed by zero, one, or many consumers. [Figure 33–4](#) illustrates pub/sub messaging.

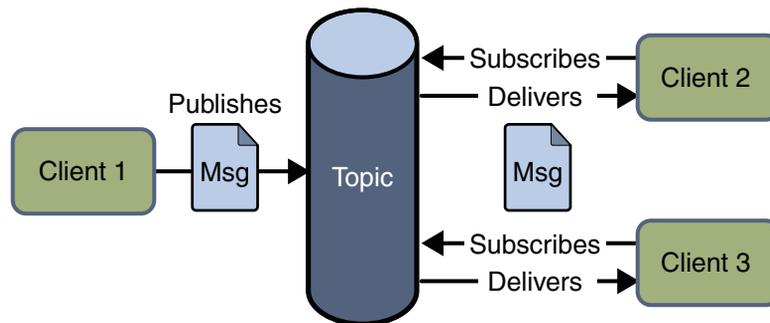


FIGURE 33–4 Publish/Subscribe Messaging

## Programming with the Common Interfaces

Version 1.1 of the JMS API allows you to use the same code to send and receive messages under either the PTP or the pub/sub domain. The destinations that you use remain domain-specific, and the behavior of the application will depend in part on whether you are using a queue or a topic. However, the code itself can be common to both domains, making your applications flexible and reusable. This tutorial describes and illustrates these common interfaces.

## Message Consumption

Messaging products are inherently asynchronous: There is no fundamental timing dependency between the production and the consumption of a message. However, the JMS specification uses this term in a more precise sense. Messages can be consumed in either of two ways:

- **Synchronously:** A subscriber or a receiver explicitly fetches the message from the destination by calling the `receive` method. The `receive` method can block until a message arrives or can time out if a message does not arrive within a specified time limit.
- **Asynchronously:** A client can register a *message listener* with a consumer. A message listener is similar to an event listener. Whenever a message arrives at the destination, the JMS provider delivers the message by calling the listener's `onMessage` method, which acts on the contents of the message.

## The JMS API Programming Model

The basic building blocks of a JMS application consist of

- Administered objects: connection factories and destinations
- Connections
- Sessions
- Message producers
- Message consumers
- Messages

Figure 33–5 shows how all these objects fit together in a JMS client application.

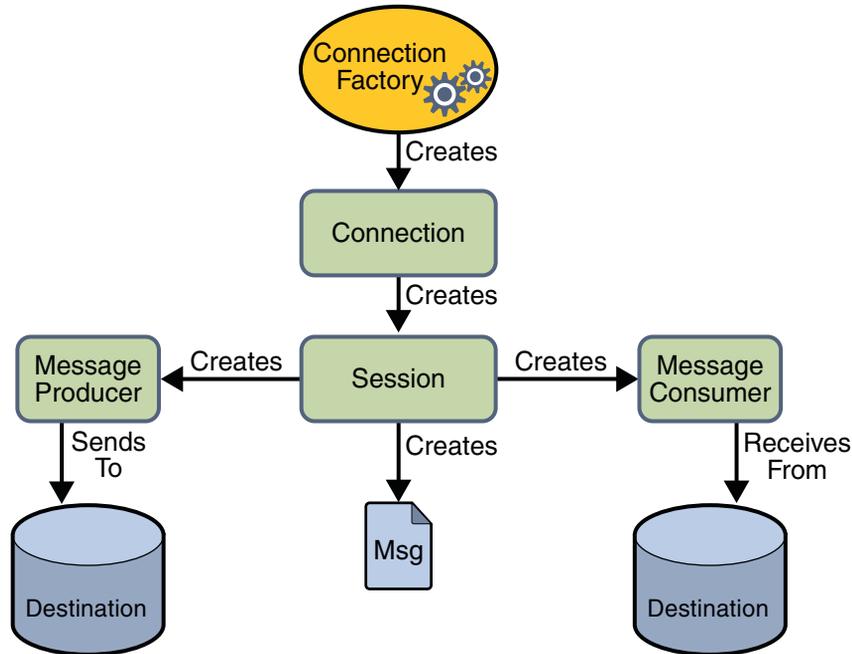


FIGURE 33-5 The JMS API Programming Model

This section describes all these objects briefly and provides sample commands and code snippets that show how to create and use the objects. The last subsection briefly describes JMS API exception handling.

Examples that show how to combine all these objects in applications appear in later sections. For more details, see the JMS API documentation, which is part of the Java EE API documentation.

## JMS Administered Objects

Two parts of a JMS application, destinations and connection factories, are best maintained administratively rather than programmatically. The technology underlying these objects is likely to be very different from one implementation of the JMS API to another. Therefore, the management of these objects belongs with other administrative tasks that vary from provider to provider.

JMS clients access these objects through interfaces that are portable, so a client application can run with little or no change on more than one implementation of the JMS API. Ordinarily, an administrator configures administered objects in a JNDI namespace, and JMS clients then access them by using resource injection.

With Sun GlassFish Enterprise Server, you use the `asadmin` command or the Administration Console to create JMS administered objects in the form of resources.

## JMS Connection Factories

A *connection factory* is the object a client uses to create a connection to a provider. A connection factory encapsulates a set of connection configuration parameters that has been defined by an administrator. Each connection factory is an instance of the `ConnectionFactory`, `QueueConnectionFactory`, or `TopicConnectionFactory` interface. To learn how to create connection factories, see [“Creating JMS Administered Objects for the Synchronous Receive Example” on page 650](#).

At the beginning of a JMS client program, you usually inject a connection factory resource into a `ConnectionFactory` object. For example, the following code fragment specifies a resource whose JNDI name is `jms/ConnectionFactory` and assigns it to a `ConnectionFactory` object:

```
@Resource(mappedName="jms/ConnectionFactory")
private static ConnectionFactory connectionFactory;
```

In a Java EE application, JMS administered objects are normally placed in the `jms` naming subcontext.

---

**Note** – The `mappedName` element of the `@Resource` annotation is specific to the Enterprise Server.

---

## JMS Destinations

A *destination* is the object a client uses to specify the target of messages it produces and the source of messages it consumes. In the PTP messaging domain, destinations are called queues. In the pub/sub messaging domain, destinations are called topics. A JMS application can use multiple queues or topics (or both). To learn how to create destination resources, see [“Creating JMS Administered Objects for the Synchronous Receive Example” on page 650](#).

To create a destination using the Enterprise Server, you create a JMS destination resource that specifies a JNDI name for the destination.

In the Enterprise Server implementation of JMS, each destination resource refers to a physical destination. You can create a physical destination explicitly, but if you do not, the Application Server creates it when it is needed and deletes it when you delete the destination resource.

In addition to injecting a connection factory resource into a client program, you usually inject a destination resource. Unlike connection factories, destinations are specific to one domain or the other. To create an application that allows you to use the same code for both topics and queues, you assign the destination to a `Destination` object.

The following code specifies two resources, a queue and a topic. The resource names are mapped to destinations created in the JNDI namespace.

```
@Resource(mappedName="jms/Queue")
private static Queue queue;
```

```
@Resource(mappedName="jms/Topic")
private static Topic topic;
```

With the common interfaces, you can mix or match connection factories and destinations. That is, in addition to using the `ConnectionFactory` interface, you can inject a `QueueConnectionFactory` resource and use it with a `Topic`, and you can inject a `TopicConnectionFactory` resource and use it with a `Queue`. The behavior of the application will depend on the kind of destination you use and not on the kind of connection factory you use.

## JMS Connections

A *connection* encapsulates a virtual connection with a JMS provider. A connection could represent an open TCP/IP socket between a client and a provider service daemon. You use a connection to create one or more sessions.

Connections implement the `Connection` interface. When you have a `ConnectionFactory` object, you can use it to create a `Connection`:

```
Connection connection = connectionFactory.createConnection();
```

Before an application completes, you must close any connections that you have created. Failure to close a connection can cause resources not to be released by the JMS provider. Closing a connection also closes its sessions and their message producers and message consumers.

```
connection.close();
```

Before your application can consume messages, you must call the connection's `start` method; for details, see [“JMS Message Consumers” on page 623](#). If you want to stop message delivery temporarily without closing the connection, you call the `stop` method.

## JMS Sessions

A *session* is a single-threaded context for producing and consuming messages. You use sessions to create the following:

- Message producers
- Message consumers
- Messages
- Queue browsers
- Temporary queues and topics (see [“Creating Temporary Destinations” on page 632](#))

Sessions serialize the execution of message listeners; for details, see [“JMS Message Listeners” on page 623](#).

A session provides a transactional context with which to group a set of sends and receives into an atomic unit of work. For details, see [“Using JMS API Local Transactions” on page 635](#).

Sessions implement the `Session` interface. After you create a `Connection` object, you use it to create a `Session`:

```
Session session = connection.createSession(false,  
    Session.AUTO_ACKNOWLEDGE);
```

The first argument means that the session is not transacted; the second means that the session automatically acknowledges messages when they have been received successfully. (For more information, see [“Controlling Message Acknowledgment” on page 629](#).)

To create a transacted session, use the following code:

```
Session session = connection.createSession(true, 0);
```

Here, the first argument means that the session is transacted; the second indicates that message acknowledgment is not specified for transacted sessions. For more information on transactions, see [“Using JMS API Local Transactions” on page 635](#). For information about the way JMS transactions work in Java EE applications, see [“Using the JMS API in Java EE Applications” on page 638](#).

## JMS Message Producers

A *message producer* is an object that is created by a session and used for sending messages to a destination. It implements the `MessageProducer` interface.

You use a `Session` to create a `MessageProducer` for a destination. The following examples show that you can create a producer for a `Destination` object, a `Queue` object, or a `Topic` object:

```
MessageProducer producer = session.createProducer(dest);  
MessageProducer producer = session.createProducer(queue);  
MessageProducer producer = session.createProducer(topic);
```

You can create an unidentified producer by specifying `null` as the argument to `createProducer`. With an unidentified producer, you do not specify a destination until you send a message.

After you have created a message producer, you can use it to send messages by using the `send` method:

```
producer.send(message);
```

You must first create the messages; see [“JMS Messages” on page 625](#).

If you created an unidentified producer, use an overloaded `send` method that specifies the destination as the first parameter. For example:

```
MessageProducer anon_prod = session.createProducer(null);
anon_prod.send(dest, message);
```

## JMS Message Consumers

A *message consumer* is an object that is created by a session and used for receiving messages sent to a destination. It implements the `MessageConsumer` interface.

A message consumer allows a JMS client to register interest in a destination with a JMS provider. The JMS provider manages the delivery of messages from a destination to the registered consumers of the destination.

For example, you could use a `Session` to create a `MessageConsumer` for a `Destination` object, a `Queue` object, or a `Topic` object:

```
MessageConsumer consumer = session.createConsumer(dest);
MessageConsumer consumer = session.createConsumer(queue);
MessageConsumer consumer = session.createConsumer(topic);
```

You use the `Session.createDurableSubscriber` method to create a durable topic subscriber. This method is valid only if you are using a topic. For details, see [“Creating Durable Subscriptions” on page 633](#).

After you have created a message consumer, it becomes active, and you can use it to receive messages. You can use the `close` method for a `MessageConsumer` to make the message consumer inactive. Message delivery does not begin until you start the connection you created by calling its `start` method. (Remember always to call the `start` method; forgetting to start the connection is one of the most common JMS programming errors.)

You use the `receive` method to consume a message synchronously. You can use this method at any time after you call the `start` method:

```
connection.start();
Message m = consumer.receive();
connection.start();
Message m = consumer.receive(1000); // time out after a second
```

To consume a message asynchronously, you use a message listener, described in the next section.

## JMS Message Listeners

A *message listener* is an object that acts as an asynchronous event handler for messages. This object implements the `MessageListener` interface, which contains one method, `onMessage`. In the `onMessage` method, you define the actions to be taken when a message arrives.

You register the message listener with a specific `MessageConsumer` by using the `setMessageListener` method. For example, if you define a class named `Listener` that implements the `MessageListener` interface, you can register the message listener as follows:

```
Listener myListener = new Listener();
consumer.setMessageListener(myListener);
```

After you register the message listener, you call the `start` method on the `Connection` to begin message delivery. (If you call `start` before you register the message listener, you are likely to miss messages.)

When message delivery begins, the JMS provider automatically calls the message listener's `onMessage` method whenever a message is delivered. The `onMessage` method takes one argument of type `Message`, which your implementation of the method can cast to any of the other message types (see [“Message Bodies” on page 626](#)).

A message listener is not specific to a particular destination type. The same listener can obtain messages from either a queue or a topic, depending on the type of destination for which the message consumer was created. A message listener does, however, usually expect a specific message type and format.

Your `onMessage` method should handle all exceptions. It must not throw checked exceptions, and throwing a `RuntimeException` is considered a programming error.

The session used to create the message consumer serializes the execution of all message listeners registered with the session. At any time, only one of the session's message listeners is running.

In the Java EE platform, a message-driven bean is a special kind of message listener. For details, see [“Using Message-Driven Beans to Receive Messages Asynchronously” on page 640](#).

## JMS Message Selectors

If your messaging application needs to filter the messages it receives, you can use a JMS API message selector, which allows a message consumer to specify the messages it is interested in. Message selectors assign the work of filtering messages to the JMS provider rather than to the application. For an example of an application that uses a message selector, see [“An Application That Uses the JMS API with a Session Bean” on page 680](#).

A message selector is a `String` that contains an expression. The syntax of the expression is based on a subset of the SQL92 conditional expression syntax. The message selector in the example selects any message that has a `NewsType` property that is set to the value `'Sports'` or `'Opinion'`:

```
NewsType = 'Sports' OR NewsType = 'Opinion'
```

The `createConsumer`, `createDurableSubscriber` methods allow you to specify a message selector as an argument when you create a message consumer.

The message consumer then receives only messages whose headers and properties match the selector. (See “[Message Headers](#)” on page 625, and “[Message Properties](#)” on page 626.) A message selector cannot select messages on the basis of the content of the message body.

## JMS Messages

The ultimate purpose of a JMS application is to produce and to consume messages that can then be used by other software applications. JMS messages have a basic format that is simple but highly flexible, allowing you to create messages that match formats used by non-JMS applications on heterogeneous platforms.

A JMS message has three parts: a header, properties, and a body. Only the header is required. The following sections describe these parts:

- “[Message Headers](#)” on page 625
- “[Message Properties](#)” on page 626
- “[Message Bodies](#)” on page 626

For complete documentation of message headers, properties, and bodies, see the documentation of the `Message` interface in the API documentation.

## Message Headers

A JMS message header contains a number of predefined fields that contain values that both clients and providers use to identify and to route messages. [Table 33–1](#) lists the JMS message header fields and indicates how their values are set. For example, every message has a unique identifier, which is represented in the header field `JMSMessageID`. The value of another header field, `JMSDestination`, represents the queue or the topic to which the message is sent. Other fields include a timestamp and a priority level.

Each header field has associated setter and getter methods, which are documented in the description of the `Message` interface. Some header fields are intended to be set by a client, but many are set automatically by the `send` or the `publish` method, which overrides any client-set values.

TABLE 33–1 How JMS Message Header Field Values Are Set

Header Field	Set By
<code>JMSDestination</code>	send or publish method
<code>JMSDeliveryMode</code>	send or publish method
<code>JMSExpiration</code>	send or publish method
<code>JMSPriority</code>	send or publish method

TABLE 33-1 How JMS Message Header Field Values Are Set (Continued)

Header Field	Set By
JMSMessageID	send or publish method
JMSTimestamp	send or publish method
JMSCorrelationID	Client
JMSReplyTo	Client
JMSType	Client
JMSRedelivered	JMS provider

## Message Properties

You can create and set properties for messages if you need values in addition to those provided by the header fields. You can use properties to provide compatibility with other messaging systems, or you can use them to create message selectors (see “[JMS Message Selectors](#)” on page 624). For an example of setting a property to be used as a message selector, see “[An Application That Uses the JMS API with a Session Bean](#)” on page 680.

The JMS API provides some predefined property names that a provider can support. The use either of these predefined properties or of user-defined properties is optional.

## Message Bodies

The JMS API defines five message body formats, also called message types, which allow you to send and to receive data in many different forms and provide compatibility with existing messaging formats. [Table 33-2](#) describes these message types.

TABLE 33-2 JMS Message Types

Message Type	Body Contains
TextMessage	A <code>java.lang.String</code> object (for example, the contents of an XML file).
MapMessage	A set of name-value pairs, with names as <code>String</code> objects and values as primitive types in the Java programming language. The entries can be accessed sequentially by enumerator or randomly by name. The order of the entries is undefined.
BytesMessage	A stream of uninterpreted bytes. This message type is for literally encoding a body to match an existing message format.
StreamMessage	A stream of primitive values in the Java programming language, filled and read sequentially.
ObjectMessage	A <code>Serializable</code> object in the Java programming language.

TABLE 33-2 JMS Message Types (Continued)

Message Type	Body Contains
Message	Nothing. Composed of header fields and properties only. This message type is useful when a message body is not required.

The JMS API provides methods for creating messages of each type and for filling in their contents. For example, to create and send a `TextMessage`, you might use the following statements:

```
TextMessage message = session.createTextMessage();
message.setText(msg_text);    // msg_text is a String
producer.send(message);
```

At the consuming end, a message arrives as a generic `Message` object and must be cast to the appropriate message type. You can use one or more getter methods to extract the message contents. The following code fragment uses the `getText` method:

```
Message m = consumer.receive();
if (m instanceof TextMessage) {
    TextMessage message = (TextMessage) m;
    System.out.println("Reading message: " + message.getText());
} else {
    // Handle error
}
```

## JMS Queue Browsers

You can create a `QueueBrowser` object to inspect the messages in a queue. Messages sent to a queue remain in the queue until the message consumer for that queue consumes them. Therefore, the JMS API provides an object that allows you to browse the messages in the queue and display the header values for each message. To create a `QueueBrowser` object, use the `Session.createBrowser` method. For example:

```
QueueBrowser browser = session.createBrowser(queue);
```

See [“A Simple Example of Browsing Messages in a Queue” on page 660](#) for an example of the use of a `QueueBrowser` object.

The `createBrowser` method allows you to specify a message selector as a second argument when you create a `QueueBrowser`. For information on message selectors, see [“JMS Message Selectors” on page 624](#).

The JMS API provides no mechanism for browsing a topic. Messages usually disappear from a topic as soon as they appear: if there are no message consumers to consume them, the JMS provider removes them. Although durable subscriptions allow messages to remain on a topic while the message consumer is not active, no facility exists for examining them.

## JMS Exception Handling

The root class for exceptions thrown by JMS API methods is `JMSException`. Catching `JMSException` provides a generic way of handling all exceptions related to the JMS API.

The `JMSException` class includes the following subclasses, which are described in the API documentation:

- `IllegalStateException`
- `InvalidClientIDException`
- `InvalidDestinationException`
- `InvalidSelectorException`
- `JMSSecurityException`
- `MessageEOFException`
- `MessageFormatException`
- `MessageNotReadableException`
- `MessageNotWriteableException`
- `ResourceAllocationException`
- `TransactionInProgressException`
- `TransactionRolledBackException`

All the examples in the tutorial catch and handle `JMSException` when it is appropriate to do so.

## Creating Robust JMS Applications

This section explains how to use features of the JMS API to achieve the level of reliability and performance your application requires. Many people choose to implement JMS applications because they cannot tolerate dropped or duplicate messages and require that every message be received once and only once. The JMS API provides this functionality.

The most reliable way to produce a message is to send a `PERSISTENT` message within a transaction. JMS messages are `PERSISTENT` by default. A *transaction* is a unit of work into which you can group a series of operations, such as message sends and receives, so that the operations either all succeed or all fail. For details, see [“Specifying Message Persistence” on page 631](#) and [“Using JMS API Local Transactions” on page 635](#).

The most reliable way to consume a message is to do so within a transaction, either from a queue or from a durable subscription to a topic. For details, see [“Creating Temporary Destinations” on page 632](#), [“Creating Durable Subscriptions” on page 633](#), and [“Using JMS API Local Transactions” on page 635](#).

For other applications, a lower level of reliability can reduce overhead and improve performance. You can send messages with varying priority levels (see [“Setting Message Priority Levels” on page 631](#)) and you can set them to expire after a certain length of time (see [“Allowing Messages to Expire” on page 632](#)).

The JMS API provides several ways to achieve various kinds and degrees of reliability. This section divides them into two categories:

- [“Using Basic Reliability Mechanisms” on page 629](#)
- [“Using Advanced Reliability Mechanisms” on page 633](#)

The following sections describe these features as they apply to JMS clients. Some of the features work differently in Java EE applications; in these cases, the differences are noted here and are explained in detail in [“Using the JMS API in Java EE Applications” on page 638](#).

## Using Basic Reliability Mechanisms

The basic mechanisms for achieving or affecting reliable message delivery are as follows:

- **Controlling message acknowledgment:** You can specify various levels of control over message acknowledgment.
- **Specifying message persistence:** You can specify that messages are persistent, meaning that they must not be lost in the event of a provider failure.
- **Setting message priority levels:** You can set various priority levels for messages, which can affect the order in which the messages are delivered.
- **Allowing messages to expire:** You can specify an expiration time for messages so that they will not be delivered if they are obsolete.
- **Creating temporary destinations:** You can create temporary destinations that last only for the duration of the connection in which they are created.

### Controlling Message Acknowledgment

Until a JMS message has been acknowledged, it is not considered to be successfully consumed. The successful consumption of a message ordinarily takes place in three stages.

1. The client receives the message.
2. The client processes the message.
3. The message is acknowledged. Acknowledgment is initiated either by the JMS provider or by the client, depending on the session acknowledgment mode.

In transacted sessions (see [“Using JMS API Local Transactions” on page 635](#)), acknowledgment happens automatically when a transaction is committed. If a transaction is rolled back, all consumed messages are redelivered.

In nontransacted sessions, when and how a message is acknowledged depend on the value specified as the second argument of the `createSession` method. The three possible argument values are as follows:

- `Session.AUTO_ACKNOWLEDGE`: The session automatically acknowledges a client's receipt of a message either when the client has successfully returned from a call to `receive` or when the `MessageListener` it has called to process the message returns successfully. A synchronous `receive` in an `AUTO_ACKNOWLEDGE` session is the one exception to the rule that message consumption is a three-stage process as described earlier.

In this case, the receipt and acknowledgment take place in one step, followed by the processing of the message.

- `Session.CLIENT_ACKNOWLEDGE`: A client acknowledges a message by calling the message's `acknowledge` method. In this mode, acknowledgment takes place on the session level: Acknowledging a consumed message automatically acknowledges the receipt of *all* messages that have been consumed by its session. For example, if a message consumer consumes ten messages and then acknowledges the fifth message delivered, all ten messages are acknowledged.
- `Session.DUPS_OK_ACKNOWLEDGE`: This option instructs the session to lazily acknowledge the delivery of messages. This is likely to result in the delivery of some duplicate messages if the JMS provider fails, so it should be used only by consumers that can tolerate duplicate messages. (If the JMS provider redelivers a message, it must set the value of the `JMSRedelivered` message header to `true`.) This option can reduce session overhead by minimizing the work the session does to prevent duplicates.

If messages have been received from a queue but not acknowledged when a session terminates, the JMS provider retains them and redelivers them when a consumer next accesses the queue. The provider also retains unacknowledged messages for a terminated session that has a durable `TopicSubscriber`. (See [“Creating Durable Subscriptions” on page 633](#).) Unacknowledged messages for a nondurable `TopicSubscriber` are dropped when the session is closed.

If you use a queue or a durable subscription, you can use the `Session.recover` method to stop a nontransacted session and restart it with its first unacknowledged message. In effect, the session's series of delivered messages is reset to the point after its last acknowledged message. The messages it now delivers may be different from those that were originally delivered, if messages have expired or if higher-priority messages have arrived. For a nondurable `TopicSubscriber`, the provider may drop unacknowledged messages when its session is recovered.

The sample program in XREF the next section demonstrates two ways to ensure that a message will not be acknowledged until processing of the message is complete.

## Specifying Message Persistence

The JMS API supports two delivery modes for messages to specify whether messages are lost if the JMS provider fails. These delivery modes are fields of the `DeliveryMode` interface.

- The `PERSISTENT` delivery mode, which is the default, instructs the JMS provider to take extra care to ensure that a message is not lost in transit in case of a JMS provider failure. A message sent with this delivery mode is logged to stable storage when it is sent.
- The `NON_PERSISTENT` delivery mode does not require the JMS provider to store the message or otherwise guarantee that it is not lost if the provider fails.

You can specify the delivery mode in either of two ways.

- You can use the `setDeliveryMode` method of the `MessageProducer` interface to set the delivery mode for all messages sent by that producer. For example, the following call sets the delivery mode to `NON_PERSISTENT` for a producer:

```
producer.setDeliveryMode(DeliveryMode.NON_PERSISTENT);
```

- You can use the long form of the `send` or the `publish` method to set the delivery mode for a specific message. The second argument sets the delivery mode. For example, the following `send` call sets the delivery mode for message to `NON_PERSISTENT`:

```
producer.send(message, DeliveryMode.NON_PERSISTENT, 3, 10000);
```

The third and fourth arguments set the priority level and expiration time, which are described in the next two subsections.

If you do not specify a delivery mode, the default is `PERSISTENT`. Using the `NON_PERSISTENT` delivery mode may improve performance and reduce storage overhead, but you should use it only if your application can afford to miss messages.

## Setting Message Priority Levels

You can use message priority levels to instruct the JMS provider to deliver urgent messages first. You can set the priority level in either of two ways.

- You can use the `setPriority` method of the `MessageProducer` interface to set the priority level for all messages sent by that producer. For example, the following call sets a priority level of 7 for a producer:

```
producer.setPriority(7);
```

- You can use the long form of the `send` or the `publish` method to set the priority level for a specific message. The third argument sets the priority level. For example, the following `send` call sets the priority level for message to 3:

```
producer.send(message, DeliveryMode.NON_PERSISTENT, 3, 10000);
```

The ten levels of priority range from 0 (lowest) to 9 (highest). If you do not specify a priority level, the default level is 4. A JMS provider tries to deliver higher-priority messages before lower-priority ones but does not have to deliver messages in exact order of priority.

## Allowing Messages to Expire

By default, a message never expires. If a message will become obsolete after a certain period, however, you may want to set an expiration time. You can do this in either of two ways.

- You can use the `setTimeToLive` method of the `MessageProducer` interface to set a default expiration time for all messages sent by that producer. For example, the following call sets a time to live of one minute for a producer:

```
producer.setTimeToLive(60000);
```

- You can use the long form of the `send` or the `publish` method to set an expiration time for a specific message. The fourth argument sets the expiration time in milliseconds. For example, the following `send` call sets a time to live of 10 seconds:

```
producer.send(message, DeliveryMode.NON_PERSISTENT, 3, 10000);
```

If the specified `timeToLive` value is `0`, the message never expires.

When the message is sent, the specified `timeToLive` is added to the current time to give the expiration time. Any message not delivered before the specified expiration time is destroyed. The destruction of obsolete messages conserves storage and computing resources.

## Creating Temporary Destinations

Normally, you create JMS destinations (queues and topics) administratively rather than programmatically. Your JMS provider includes a tool that you use to create and remove destinations, and it is common for destinations to be long-lasting.

The JMS API also enables you to create destinations (`TemporaryQueue` and `TemporaryTopic` objects) that last only for the duration of the connection in which they are created. You create these destinations dynamically using the `Session.createTemporaryQueue` and the `Session.createTemporaryTopic` methods.

The only message consumers that can consume from a temporary destination are those created by the same connection that created the destination. Any message producer can send to the temporary destination. If you close the connection that a temporary destination belongs to, the destination is closed and its contents are lost.

You can use temporary destinations to implement a simple request/reply mechanism. If you create a temporary destination and specify it as the value of the `JMSReplyTo` message header field when you send a message, then the consumer of the message can use the value of the `JMSReplyTo` field as the destination to which it sends a reply. The consumer can also reference the original request by setting the `JMSCorrelationID` header field of the reply message to the value of the `JMSMessageID` header field of the request. For example, an `onMessage` method can create a session so that it can send a reply to the message it receives. It can use code such as the following:

```
producer = session.createProducer(msg.getJMSReplyTo());
replyMsg = session.createTextMessage("Consumer " +
    "processed message: " + msg.getText());
replyMsg.setJMSCorrelationID(msg.getJMSMessageID());
producer.send(replyMsg);
```

For more examples, see [Chapter 34, “Java Message Service Examples.”](#)

## Using Advanced Reliability Mechanisms

The more advanced mechanisms for achieving reliable message delivery are the following:

- **Creating durable subscriptions:** You can create durable topic subscriptions, which receive messages published while the subscriber is not active. Durable subscriptions offer the reliability of queues to the publish/subscribe message domain.
- **Using local transactions:** You can use local transactions, which allow you to group a series of sends and receives into an atomic unit of work. Transactions are rolled back if they fail at any time.

### Creating Durable Subscriptions

To ensure that a pub/sub application receives all published messages, use `PERSISTENT` delivery mode for the publishers. In addition, use durable subscriptions for the subscribers.

The `Session.createConsumer` method creates a nondurable subscriber if a topic is specified as the destination. A nondurable subscriber can receive only messages that are published while it is active.

At the cost of higher overhead, you can use the `Session.createDurableSubscriber` method to create a durable subscriber. A durable subscription can have only one active subscriber at a time.

A durable subscriber registers a durable subscription by specifying a unique identity that is retained by the JMS provider. Subsequent subscriber objects that have the same identity resume the subscription in the state in which it was left by the preceding subscriber. If a durable subscription has no active subscriber, the JMS provider retains the subscription’s messages until they are received by the subscription or until they expire.

You establish the unique identity of a durable subscriber by setting the following:

- A client ID for the connection
- A topic and a subscription name for the subscriber

You set the client ID administratively for a client-specific connection factory using either the command line or the Administration Console.

After using this connection factory to create the connection and the session, you call the `createDurableSubscriber` method with two arguments: the topic and a string that specifies the name of the subscription:

```
String subName = "MySub";
MessageConsumer topicSubscriber =
    session.createDurableSubscriber(myTopic, subName);
```

The subscriber becomes active after you start the `Connection` or `TopicConnection`. Later, you might close the subscriber:

```
topicSubscriber.close();
```

The JMS provider stores the messages sent or published to the topic, as it would store messages sent to a queue. If the program or another application calls `createDurableSubscriber` using the same connection factory and its client ID, the same topic, and the same subscription name, the subscription is reactivated, and the JMS provider delivers the messages that were published while the subscriber was inactive.

To delete a durable subscription, first close the subscriber, and then use the `unsubscribe` method, with the subscription name as the argument:

```
topicSubscriber.close();
session.unsubscribe("MySub");
```

The `unsubscribe` method deletes the state that the provider maintains for the subscriber.

[Figure 33–6](#) and [Figure 33–7](#) show the difference between a nondurable and a durable subscriber. With an ordinary, nondurable subscriber, the subscriber and the subscription begin and end at the same point and are, in effect, identical. When a subscriber is closed, the subscription also ends. Here, `create` stands for a call to `Session.createConsumer` with a `Topic` argument, and `close` stands for a call to `MessageConsumer.close`. Any messages published to the topic between the time of the first `close` and the time of the second `create` are not consumed by the subscriber. In [Figure 33–6](#), the subscriber consumes messages M1, M2, M5, and M6, but messages M3 and M4 are lost.

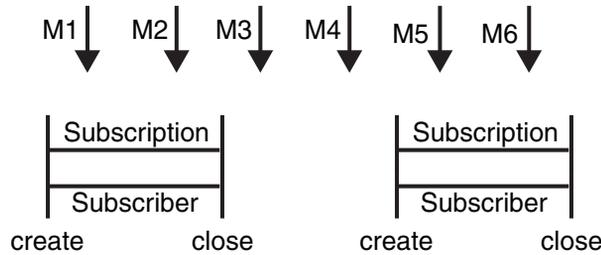


FIGURE 33-6 Nondurable Subscribers and Subscriptions

With a durable subscriber, the subscriber can be closed and re-created, but the subscription continues to exist and to hold messages until the application calls the `unsubscribe` method. In [Figure 33-7](#), `create` stands for a call to `Session.createDurableSubscriber`, `close` stands for a call to `MessageConsumer.close`, and `unsubscribe` stands for a call to `Session.unsubscribe`. Messages published while the subscriber is closed are received when the subscriber is created again. So even though messages M2, M4, and M5 arrive while the subscriber is closed, they are not lost.

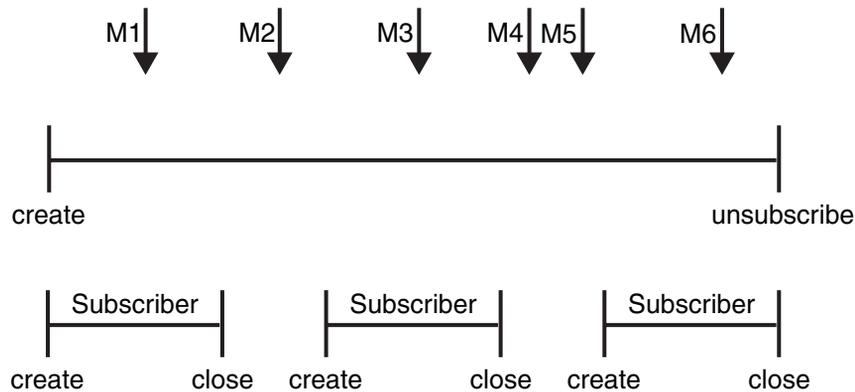


FIGURE 33-7 A Durable Subscriber and Subscription

See [“A Message Acknowledgment Example”](#) on page 670, [“A Durable Subscription Example”](#) on page 672, and [“An Application That Uses the JMS API with a Session Bean”](#) on page 680 for examples of Java EE applications that use durable subscriptions.

## Using JMS API Local Transactions

You can group a series of operations into an atomic unit of work called a transaction. If any one of the operations fails, the transaction can be rolled back, and the operations can be attempted again from the beginning. If all the operations succeed, the transaction can be committed.

In a JMS client, you can use local transactions to group message sends and receives. The JMS API `Session` interface provides `commit` and `rollback` methods that you can use in a JMS client. A transaction commit means that all produced messages are sent and all consumed messages are acknowledged. A transaction rollback means that all produced messages are destroyed and all consumed messages are recovered and redelivered unless they have expired (see [“Allowing Messages to Expire” on page 632](#)).

A transacted session is always involved in a transaction. As soon as the `commit` or the `rollback` method is called, one transaction ends and another transaction begins. Closing a transacted session rolls back its transaction in progress, including any pending sends and receives.

In an Enterprise JavaBeans component, you cannot use the `Session.commit` and `Session.rollback` methods. Instead, you use distributed transactions, which are described in [“Using the JMS API in Java EE Applications” on page 638](#).

You can combine several sends and receives in a single JMS API local transaction. If you do so, you need to be careful about the order of the operations. You will have no problems if the transaction consists of all sends or all receives or if the receives come before the sends. But if you try to use a request/reply mechanism, whereby you send a message and then try to receive a reply to the sent message in the same transaction, the program will hang, because the send cannot take place until the transaction is committed. The following code fragment illustrates the problem:

```
// Don't do this!
outMsg.setJMSReplyTo(replyQueue);
producer.send(outQueue, outMsg);
consumer = session.createConsumer(replyQueue);
inMsg = consumer.receive();
session.commit();
```

Because a message sent during a transaction is not actually sent until the transaction is committed, the transaction cannot contain any receives that depend on that message's having been sent.

In addition, the production and the consumption of a message cannot both be part of the same transaction. The reason is that the transactions take place between the clients and the JMS provider, which intervenes between the production and the consumption of the message. [Figure 33–8](#) illustrates this interaction.

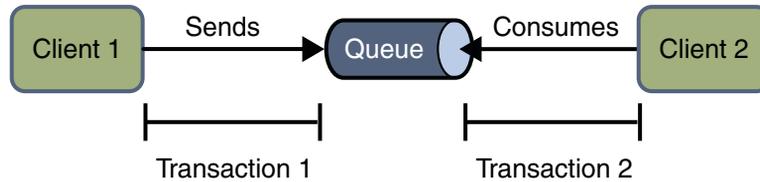


FIGURE 33-8 Using JMS API Local Transactions

The sending of one or more messages to one or more destinations by client 1 can form a single transaction, because it forms a single set of interactions with the JMS provider using a single session. Similarly, the receiving of one or more messages from one or more destinations by client 2 also forms a single transaction using a single session. But because the two clients have no direct interaction and are using two different sessions, no transactions can take place between them.

Another way of putting this is that the act of producing and/or consuming messages in a session can be transactional, but the act of producing and consuming a specific message across different sessions cannot be transactional.

This is the fundamental difference between messaging and synchronized processing. Instead of tightly coupling the sending and receiving of data, message producers and consumers use an alternative approach to reliability, one that is built on a JMS provider's ability to supply a once-and-only-once message delivery guarantee.

When you create a session, you specify whether it is transacted. The first argument to the `createSession` method is a boolean value. A value of `true` means that the session is transacted; a value of `false` means that it is not transacted. The second argument to this method is the acknowledgment mode, which is relevant only to nontransacted sessions (see [“Controlling Message Acknowledgment” on page 629](#)). If the session is transacted, the second argument is ignored, so it is a good idea to specify `0` to make the meaning of your code clear. For example:

```
session = connection.createSession(true, 0);
```

The `commit` and the `rollback` methods for local transactions are associated with the session. You can combine queue and topic operations in a single transaction if you use the same session to perform the operations. For example, you can use the same session to receive a message from a queue and send a message to a topic in the same transaction.

You can pass a client program's session to a message listener's constructor function and use it to create a message producer. In this way, you can use the same session for receives and sends in asynchronous message consumers.

[“A Local Transaction Example” on page 674](#) provides an example of the use of JMS API local transactions.

## Using the JMS API in Java EE Applications

This section describes the ways in which using the JMS API in enterprise bean applications or web applications differs from using it in application clients:

- “Using `@Resource` Annotations in Enterprise Bean or Web Components” on page 638
- “Using Session Beans to Produce and to Synchronously Receive Messages” on page 639
- “Using Message-Driven Beans to Receive Messages Asynchronously” on page 640
- “Managing Distributed Transactions” on page 642
- “Using the JMS API with Application Clients and Web Components” on page 644

A general rule in the Java EE platform specification applies to all Java EE components that use the JMS API within EJB or web containers:

Application components in the web and EJB containers must not attempt to create more than one active (not closed) `Session` object per connection.

This rule does not apply to application clients. The application client container supports the creation of multiple sessions for each connection.

### Using `@Resource` Annotations in Enterprise Bean or Web Components

When you use the `@Resource` annotation in an application client component, you normally declare the JMS resource static:

```
@Resource(mappedName="jms/ConnectionFactory")
private static ConnectionFactory connectionFactory;
```

```
@Resource(mappedName="jms/Queue")
private static Queue queue;
```

However, when you use this annotation in a session bean, a message-driven bean, or a web component, do *not* declare the resource static:

```
@Resource(mappedName="jms/ConnectionFactory")
private ConnectionFactory connectionFactory;
```

```
@Resource(mappedName="jms/Topic")
private Topic topic;
```

If you declare the resource static, runtime errors will result.

## Using Session Beans to Produce and to Synchronously Receive Messages

An application that produces messages or synchronously receives them can use a session bean to perform these operations. The example in [“An Application That Uses the JMS API with a Session Bean” on page 680](#) uses a stateless session bean to publish messages to a topic.

Because a blocking synchronous receive ties up server resources, it is not a good programming practice to use such a receive call in an enterprise bean. Instead, use a timed synchronous receive, or use a message-driven bean to receive messages asynchronously. For details about blocking and timed synchronous receives, see [“Writing the Clients for the Synchronous Receive Example” on page 647](#).

Using the JMS API in an enterprise bean or web application is in many ways similar to using it in an application client. The main differences are in resource management and transactions.

### Resource Management

The JMS API resources are a JMS API connection and a JMS API session. In general, it is important to release JMS resources when they are no longer being used. Here are some useful practices to follow.

- If you wish to maintain a JMS API resource only for the life span of a business method, it is a good idea to close the resource in a `finally` block within the method.
- If you would like to maintain a JMS API resource for the life span of an enterprise bean instance, it is a good idea to use a `@PostConstruct` callback method to create the resource and to use a `@PreDestroy` callback method to close the resource. If you use a stateful session bean and you wish to maintain the JMS API resource in a cached state, you must close the resource in a `@PrePassivate` callback method and set its value to `null`, and you must create it again in a `@PostActivate` callback method.

### Transactions

Instead of using local transactions, you use container-managed transactions for bean methods that perform sends or receives, allowing the EJB container to handle transaction demarcation. Because container-managed transactions are the default, you do not have to use an annotation to specify them.

You can use bean-managed transactions and the `javax.transaction.UserTransaction` interface's transaction demarcation methods, but you should do so only if your application has special requirements and you are an expert in using transactions. Usually, container-managed transactions produce the most efficient and correct behavior. This tutorial does not provide any examples of bean-managed transactions.

## Using Message-Driven Beans to Receive Messages Asynchronously

The sections “[What Is a Message-Driven Bean?](#)” on page 285 and “[How Does the JMS API Work with the Java EE Platform?](#)” on page 613 describe how the Java EE platform supports a special kind of enterprise bean, the message-driven bean, which allows Java EE applications to process JMS messages asynchronously. Session beans allow you to send messages and to receive them synchronously but not asynchronously.

A message-driven bean is a message listener that can reliably consume messages from a queue or a durable subscription. The messages can be sent by any Java EE component (from an application client, another enterprise bean, or a web component) or from an application or a system that does not use Java EE technology.

Like a message listener in an application client, a message-driven bean contains an `onMessage` method that is called automatically when a message arrives. Like a message listener, a message-driven bean class can implement helper methods invoked by the `onMessage` method to aid in message processing.

A message-driven bean, however, differs from an application client’s message listener in the following ways:

- Certain setup tasks are performed by the EJB container.
- The bean class uses the `@MessageDriven` annotation to specify properties for the bean or the connection factory, such as a destination type, a durable subscription, a message selector, or an acknowledgment mode. The examples in [Chapter 34, “Java Message Service Examples,”](#) show how the JMS resource adapter works in the Enterprise Server.

The EJB container automatically performs several setup tasks that a stand-alone client has to do:

- Creating a message consumer to receive the messages. Instead of creating a message consumer in your source code, you associate the message-driven bean with a destination and a connection factory at deployment time. If you want to specify a durable subscription or use a message selector, you do this at deployment time also.
- Registering the message listener. You must not call `setMessageListener`.
- Specifying a message acknowledgment mode. The default mode, `AUTO_ACKNOWLEDGE`, is used unless it is overridden by a property setting.

If JMS is integrated with the application server using a resource adapter, the JMS resource adapter handles these tasks for the EJB container.

Your message-driven bean class must implement the `javax.jms.MessageListener` interface and the `onMessage` method.

It may implement a `@PostConstruct` callback method to create a connection, and a `@PreDestroy` callback method to close the connection. Typically, it implements these methods if it produces messages or does synchronous receives from another destination.

The bean class commonly injects a `MessageDrivenContext` resource, which provides some additional methods that you can use for transaction management.

The main difference between a message-driven bean and a session bean is that a message-driven bean has no local or remote interface. Instead, it has only a bean class.

A message-driven bean is similar in some ways to a stateless session bean: Its instances are relatively short-lived and retain no state for a specific client. The instance variables of the message-driven bean instance can contain some state across the handling of client messages: for example, a JMS API connection, an open database connection, or an object reference to an enterprise bean object.

Like a stateless session bean, a message-driven bean can have many interchangeable instances running at the same time. The container can pool these instances to allow streams of messages to be processed concurrently. The container attempts to deliver messages in chronological order when it does not impair the concurrency of message processing, but no guarantees are made as to the exact order in which messages are delivered to the instances of the message-driven bean class. Because concurrency can affect the order in which messages are delivered, you should write your applications to handle messages that arrive out of sequence.

For example, your application could manage conversations by using application-level sequence numbers. An application-level conversation control mechanism with a persistent conversation state could cache later messages until earlier messages have been processed.

Another way to ensure order is to have each message or message group in a conversation require a confirmation message that the sender blocks on receipt of. This forces the responsibility for order back on the sender and more tightly couples senders to the progress of message-driven beans.

To create a new instance of a message-driven bean, the container does the following:

- Instantiates the bean
- Performs any required resource injection
- Calls the `@PostConstruct` callback method, if it exists

To remove an instance of a message-driven bean, the container calls the `@PreDestroy` callback method.

Figure 33–9 shows the life cycle of a message-driven bean.

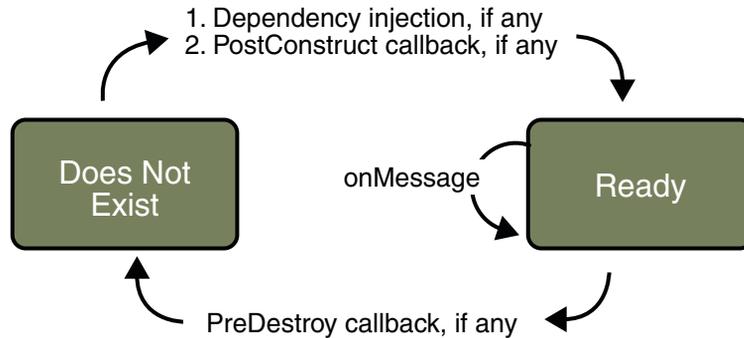


FIGURE 33-9 Life Cycle of a Message-Driven Bean

## Managing Distributed Transactions

JMS client applications use JMS API local transactions (described in “[Using JMS API Local Transactions](#)” on page 635), which allow the grouping of sends and receives within a specific JMS session. Java EE applications commonly use distributed transactions to ensure the integrity of accesses to external resources. For example, distributed transactions allow multiple applications to perform atomic updates on the same database, and they allow a single application to perform atomic updates on multiple databases.

In a Java EE application that uses the JMS API, you can use transactions to combine message sends or receives with database updates and other resource manager operations. You can access resources from multiple application components within a single transaction. For example, a servlet can start a transaction, access multiple databases, invoke an enterprise bean that sends a JMS message, invoke another enterprise bean that modifies an EIS system using the Connector architecture, and finally commit the transaction. Your application cannot, however, both send a JMS message and receive a reply to it within the same transaction; the restriction described in “[Using JMS API Local Transactions](#)” on page 635 still applies.

Distributed transactions within the EJB container can be either of two kinds:

- Container-managed transactions:** The EJB container controls the integrity of your transactions without your having to call `commit` or `rollback`. Container-managed transactions are recommended for Java EE applications that use the JMS API. You can specify appropriate transaction attributes for your enterprise bean methods.

Use the `Required` transaction attribute (the default) to ensure that a method is always part of a transaction. If a transaction is in progress when the method is called, the method will be part of that transaction; if not, a new transaction will be started before the method is called and will be committed when the method returns.

- **Bean-managed transactions:** You can use these in conjunction with the `javax.transaction.UserTransaction` interface, which provides its own `commit` and `rollback` methods that you can use to delimit transaction boundaries. Bean-managed transactions are recommended only for those who are experienced in programming transactions.

You can use either container-managed transactions or bean-managed transactions with message-driven beans. To ensure that all messages are received and handled within the context of a transaction, use container-managed transactions and use the `Required` transaction attribute (the default) for the `onMessage` method. This means that if there is no transaction in progress, a new transaction will be started before the method is called and will be committed when the method returns.

When you use container-managed transactions, you can call the following `MessageDrivenContext` methods:

- `setRollbackOnly`: Use this method for error handling. If an exception occurs, `setRollbackOnly` marks the current transaction so that the only possible outcome of the transaction is a rollback.
- `getRollbackOnly`: Use this method to test whether the current transaction has been marked for rollback.

If you use bean-managed transactions, the delivery of a message to the `onMessage` method takes place outside the distributed transaction context. The transaction begins when you call the `UserTransaction.begin` method within the `onMessage` method, and it ends when you call `UserTransaction.commit` or `UserTransaction.rollback`. Any call to the `Connection.createSession` method must take place within the transaction. If you call `UserTransaction.rollback`, the message is not redelivered, whereas calling `setRollbackOnly` for container-managed transactions does cause a message to be redelivered.

Neither the JMS API specification nor the Enterprise JavaBeans specification (available from <http://java.sun.com/products/ejb/>) specifies how to handle calls to JMS API methods outside transaction boundaries. The Enterprise JavaBeans specification does state that the EJB container is responsible for acknowledging a message that is successfully processed by the `onMessage` method of a message-driven bean that uses bean-managed transactions. Using bean-managed transactions allows you to process the message by using more than one transaction or to have some parts of the message processing take place outside a transaction context. In most cases, however, container-managed transactions provide greater reliability and are therefore preferable.

When you create a session in an enterprise bean, the container ignores the arguments you specify, because it manages all transactional properties for enterprise beans. It is still a good idea to specify arguments of `true` and `0` to the `createSession` method to make this situation clear:

```
session = connection.createSession(true, 0);
```

When you use container-managed transactions, you normally use the `Required` transaction attribute (the default) for your enterprise bean's business methods.

You do not specify a message acknowledgment mode when you create a message-driven bean that uses container-managed transactions. The container acknowledges the message automatically when it commits the transaction.

If a message-driven bean uses bean-managed transactions, the message receipt cannot be part of the bean-managed transaction, so the container acknowledges the message outside the transaction.

If the `onMessage` method throws a `RuntimeException`, the container does not acknowledge processing the message. In that case, the JMS provider will redeliver the unacknowledged message in the future.

## Using the JMS API with Application Clients and Web Components

An application client in a Java EE application can use the JMS API in much the same way that a stand-alone client program does. It can produce messages, and it can consume messages by using either synchronous receives or message listeners. See [Chapter 19, “A Message-Driven Bean Example,”](#) for an example of an application client that produces messages.

The Java EE platform specification does not impose strict constraints on how web components should use the JMS API. In the Enterprise Server, a web component can send messages and consume them synchronously but cannot consume them asynchronously.

Because a blocking synchronous receive ties up server resources, it is not a good programming practice to use such a receive call in a web component. Instead, use a timed synchronous receive. For details about blocking and timed synchronous receives, see [“Writing the Clients for the Synchronous Receive Example”](#) on page 647.

## Further Information about JMS

For more information about JMS, see:

- Java Message Service web site:  
<http://java.sun.com/products/jms/>
- Java Message Service specification, version 1.1, available from  
<http://java.sun.com/products/jms/docs.html>

## Java Message Service Examples

---

This chapter provides examples that show how to use the JMS API in various kinds of Java EE applications. It covers the following topics:

- “Writing Simple JMS Applications” on page 646
- “Writing Robust JMS Applications” on page 670
- “An Application That Uses the JMS API with a Session Bean” on page 680
- “An Application That Uses the JMS API with an Entity” on page 686

The examples are in the following directory:

*tut-install/examples/jms/*

To build and run the examples, you will do the following:

1. Use NetBeans IDE or the Ant tool to compile and package the example.
2. Use the Ant tool to create resources.
3. Use NetBeans IDE or the Ant tool to deploy the example.
4. Use NetBeans IDE or the Ant tool to run the client.

Each example has a `build.xml` file that refers to files in the following directory:

*tut-install/examples/bp-project/*

See [Chapter 19, “A Message-Driven Bean Example,”](#) for a simpler example of a Java EE application that uses the JMS API.

## Writing Simple JMS Applications

This section shows how to create, package, and run simple JMS clients that are packaged as application clients and deployed to a Java EE server. The clients demonstrate the basic tasks that a JMS application must perform:

- Creating a connection and a session
- Creating message producers and consumers
- Sending and receiving messages

In a Java EE application, some of these tasks are performed, in whole or in part, by the container. If you learn about these tasks, you will have a good basis for understanding how a JMS application works on the Java EE platform.

This section covers the following topics:

- [“A Simple Example of Synchronous Message Receives” on page 646](#)
- [“A Simple Example of Asynchronous Message Consumption” on page 655](#)
- [“A Simple Example of Browsing Messages in a Queue” on page 660](#)

Each example uses two clients: one that sends messages and one that receives them. You can run the clients in NetBeans IDE or in two terminal windows.

When you write a JMS client to run in an enterprise bean application, you use many of the same methods in much the same sequence as you do for an application client. However, there are some significant differences. [“Using the JMS API in Java EE Applications” on page 638](#) describes these differences, and this chapter provides examples that illustrate them.

The examples for this section are in the following directory:

```
tut-install/examples/jms/simple/
```

The examples are in the following four subdirectories:

```
producer  
synchconsumer  
asynchconsumer  
messagebrowser
```

### A Simple Example of Synchronous Message Receives

This section describes the sending and receiving clients in an example that uses the `receive` method to consume messages synchronously. This section then explains how to compile, package, and run the clients using the Enterprise Server.

The following sections describe the steps in creating and running the example:

- “Writing the Clients for the Synchronous Receive Example” on page 647
- “Starting the JMS Provider” on page 649
- “Creating JMS Administered Objects for the Synchronous Receive Example” on page 650
- “Compiling and Packaging the Clients for the Synchronous Receive Example” on page 650
- “Deploying and Running the Clients for the Synchronous Receive Example” on page 651

## Writing the Clients for the Synchronous Receive Example

The sending client, `producer/src/java/Producer.java`, performs the following steps:

1. Injects resources for a connection factory, queue, and topic:

```
@Resource(mappedName="jms/ConnectionFactory")
private static ConnectionFactory connectionFactory;
@Resource(mappedName="jms/Queue")private static Queue queue;
@Resource(mappedName="jms/Topic")private static Topic topic;
```

2. Retrieves and verifies command-line arguments that specify the destination type and the number of arguments:

```
final int NUM_MSGS;
String destType = args[0];
System.out.println("Destination type is " + destType);
if ( ! ( destType.equals("queue") || destType.equals("topic") ) ) {
    System.err.println("Argument must be \"queue\" or \"topic\"");
    System.exit(1);
}
if (args.length == 2){
    NUM_MSGS = (new Integer(args[1])).intValue();
}
else {
    NUM_MSGS = 1;
}
```

3. Assigns either the queue or topic to a destination object, based on the specified destination type:

```
Destination dest = null;
try {
    if (destType.equals("queue")) {
        dest = (Destination) queue;
    } else {
        dest = (Destination) topic;
    }
}
catch (Exception e) {
    System.err.println("Error setting destination: " + e.toString());
    e.printStackTrace();
}
```

```
        System.exit(1);  
    }
```

4. Creates a Connection and a Session:

```
Connection connection = connectionFactory.createConnection();  
Session session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE);
```

5. Creates a MessageProducer and a TextMessage:

```
MessageProducer producer = session.createProducer(dest);  
TextMessage message = session.createTextMessage();
```

6. Sends one or more messages to the destination:

```
for (int i = 0; i < NUM_MSGS; i++) {  
    message.setText("This is message " + (i + 1) + " from producer");  
    System.out.println("Sending message: " + message.getText());  
    producer.send(message);  
}
```

7. Sends an empty control message to indicate the end of the message stream:

```
producer.send(session.createMessage());
```

Sending an empty message of no specified type is a convenient way to indicate to the consumer that the final message has arrived.

8. Closes the connection in a finally block, automatically closing the session and MessageProducer:

```
} finally {  
    if (connection != null) {  
        try { connection.close(); }  
        catch (JMSEException e) { }  
    }  
}
```

The receiving client, `synchconsumer/src/java/SynchConsumer.java`, performs the following steps:

1. Injects resources for a connection factory, queue, and topic.
2. Assigns either the queue or topic to a destination object, based on the specified destination type.
3. Creates a Connection and a Session.
4. Creates a MessageConsumer:

```
consumer = session.createConsumer(dest);
```

5. Starts the connection, causing message delivery to begin:

```
connection.start();
```

6. Receives the messages sent to the destination until the end-of-message-stream control message is received:

```
while (true) {
    Message m = consumer.receive(1);
    if (m != null) {
        if (m instanceof TextMessage) {
            message = (TextMessage) m;
            System.out.println("Reading message: " + message.getText());
        } else {
            break;
        }
    }
}
```

Because the control message is not a `TextMessage`, the receiving client terminates the `while` loop and stops receiving messages after the control message arrives.

7. Closes the connection in a `finally` block, automatically closing the session and `MessageConsumer`.

The `receive` method can be used in several ways to perform a synchronous receive. If you specify no arguments or an argument of `0`, the method blocks indefinitely until a message arrives:

```
Message m = consumer.receive();
Message m = consumer.receive(0);
```

For a simple client, this may not matter. But if you do not want your application to consume system resources unnecessarily, use a timed synchronous receive. Do one of the following:

- Call the `receive` method with a timeout argument greater than `0`:

```
Message m = consumer.receive(1); // 1 millisecond
```

- Call the `receiveNowait` method, which receives a message only if one is available:

```
Message m = consumer.receiveNowait();
```

The `SynchConsumer` client uses an indefinite `while` loop to receive messages, calling `receive` with a timeout argument. Calling `receiveNowait` would have the same effect.

## Starting the JMS Provider

When you use the Enterprise Server, your JMS provider is the Enterprise Server. Start the server as described in [“Starting and Stopping the Enterprise Server”](#) on page 63.

## Creating JMS Administered Objects for the Synchronous Receive Example

Creating the JMS administered objects for this section involves the following:

- Creating a connection factory
- Creating two destination resources

If you built and ran the `SimpleMessage` example in [Chapter 19, “A Message-Driven Bean Example,”](#) and did not delete the resources afterward, you need to create only the topic resource.

You can create these objects using the Ant tool. To create all the resources, do the following:

1. In a terminal window, go to the producer directory:

```
cd producer
```

2. To create all the resources, type the following command:

```
ant create-resources
```

To create only the topic resource, type the following command:

```
ant create-topic
```

These Ant targets use the `asadmin create-jms-resource` command to create the connection factory and the destination resources.

To verify that the resources have been created, use the following command:

```
asadmin list-jms-resources
```

The output looks like this:

```
jms/Queue  
jms/Topic  
jms/ConnectionFactory  
Command list-jms-resources executed successfully.
```

## Compiling and Packaging the Clients for the Synchronous Receive Example

To run these examples using the Enterprise Server, package each one in an application client JAR file. The application client JAR file requires a manifest file, located in the `src/conf` directory for each example, along with the `.class` file.

The `build.xml` file for each example contains Ant targets that compile and package the example. The targets place the `.class` file for the example in the `build/jar` directory. Then the targets use the `jar` command to package the class file and the manifest file in an application client JAR file.

To compile and package the `Producer` and `SynchConsumer` examples using NetBeans IDE, follow these steps:

1. In NetBeans IDE, choose `Open Project` from the `File` menu.
2. In the `Open Project` dialog, navigate to `tut-install/examples/jms/simple/`.
3. Select the `producer` folder.
4. Select the `Open as Main Project` check box.
5. Click `Open Project`.
6. Right-click the project and choose `Build`.
7. In NetBeans IDE, choose `Open Project` from the `File` menu.
8. In the `Open Project` dialog, navigate to `tut-install/examples/jms/simple/`.
9. Select the `synchconsumer` folder.
10. Select the `Open as Main Project` check box.
11. Click `Open Project`.
12. Right-click the project and choose `Build`.

To compile and package the `Producer` and `SynchConsumer` examples using Ant, follow these steps:

1. In a terminal window, go to the `producer` directory:  

```
cd producer
```
2. Type the following command:  

```
ant
```
3. In a terminal window, go to the `synchconsumer` directory:  

```
cd ../synchconsumer
```
4. Type the following command:  

```
ant
```

The targets place the application client JAR file in the `dist` directory for each example.

## Deploying and Running the Clients for the Synchronous Receive Example

To deploy and run the clients using NetBeans IDE, follow these steps.

1. Run the `Producer` example:
  - a. Right-click the `producer` project and choose `Properties`.
  - b. Select `Run` from the `Categories` tree.
  - c. In the `Arguments` field, type the following:  

```
queue 3
```
  - d. Click `OK`.
  - e. Right-click the project and choose `Run`.

The output of the program looks like this:

```
Destination type is queue
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

The messages are now in the queue, waiting to be received.

2. Now run the SynchConsumer example:
  - a. Right-click the synchconsumer project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

**queue**

- d. Click OK.
- e. Right-click the project and choose Run.

The output of the program looks like this:

```
Destination type is queue
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
```

3. Now try running the programs in the opposite order. Right-click the synchconsumer project and choose Run.

The Output pane displays the destination type and then appears to hang, waiting for messages.

4. Right-click the producer project and choose Run.

The Output pane shows the output of both programs, in two different tabs.

5. Now run the Producer example using a topic instead of a queue.
  - a. Right-click the producer project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

**topic 3**

- d. Click OK.
- e. Right-click the project and choose Run.

The output looks like this:

```
Destination type is topic
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

6. Now run the SynchConsumer example using the topic.
  - a. Right-click the synchconsumer project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:
 

```
topic
```
  - d. Click OK.
  - e. Right-click the project and choose Run.

The result, however, is different. Because you are using a topic, messages that were sent before you started the consumer cannot be received. (See “[Publish/Subscribe Messaging Domain](#)” on page 617, for details.) Instead of receiving the messages, the program appears to hang.

7. Run the Producer example again. Right-click the producer project and choose Run. Now the SynchConsumer example receives the messages:

```
Destination type is topic
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
```

You can also run the clients using the `appclient` command. The `build.xml` file for each project includes a target that deploys the client and then retrieves the client stubs that the `appclient` command uses. Each of the clients takes one or more command-line arguments: a destination type and, for `Producer`, a number of messages.

To run the clients using the `appclient` command, follow these steps:

1. In a terminal window, go to the producer directory:
 

```
cd ../producer
```
2. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

3. Run the Producer program, sending three messages to the queue:

```
appclient -client client-jar/producerClient.jar queue 3
```

The output of the program looks like this (along with some application client container output):

```
Destination type is queue
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

The messages are now in the queue, waiting to be received.

4. In the same window, go to the `synchconsumer` directory:

```
cd ../synchconsumer
```

5. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

6. Run the `SynchConsumer` client, specifying the queue:

```
appclient -client client-jar/synchconsumerClient.jar queue
```

The output of the client looks like this (along with some application client container output):

```
Destination type is queue
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
```

7. Now try running the clients in the opposite order. Run the `SynchConsumer` client. It displays the destination type and then appears to hang, waiting for messages.

```
appclient -client client-jar/synchconsumerClient.jar queue
```

8. In a different terminal window, run the `Producer` client.

```
cd tut-install/examples/jms/simple/producer
appclient -client client-jar/producerClient.jar queue 3
```

When the messages have been sent, the `SynchConsumer` client receives them and exits.

9. Now run the `Producer` client using a topic instead of a queue:

```
appclient -client client-jar/producerClient.jar topic 3
```

The output of the client looks like this (along with some application client container output):

```
Destination type is topic
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

10. Now run the `SynchConsumer` client using the topic:

```
appclient -client client-jar/synchconsumerClient.jar topic
```

The result, however, is different. Because you are using a topic, messages that were sent before you started the consumer cannot be received. (See [“Publish/Subscribe Messaging Domain” on page 617](#), for details.) Instead of receiving the messages, the client appears to hang.

11. Run the Producer client again. Now the SynchronConsumer client receives the messages (along with some application client container output):

```
Destination type is topic
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
```

Because the examples use the common interfaces, you can run them using either a queue or a topic.

## A Simple Example of Asynchronous Message Consumption

This section describes the receiving clients in an example that uses a message listener to consume messages asynchronously. This section then explains how to compile and run the clients using the Enterprise Server.

The following sections describe the steps in creating and running the example:

- [“Writing the Clients for the Asynchronous Receive Example” on page 655](#)
- [“Compiling and Packaging the AsynchConsumer Client” on page 656](#)
- [“Running the Clients for the Asynchronous Receive Example” on page 657](#)

### Writing the Clients for the Asynchronous Receive Example

The sending client is `producer/src/java/Producer.java`, the same client used in the example in [“A Simple Example of Synchronous Message Receives” on page 646](#).

An asynchronous consumer normally runs indefinitely. This one runs until the user types the letter `q` or `Q` to stop the client.

The receiving client, `asynchconsumer/src/java/AsynchConsumer.java`, performs the following steps:

1. Injects resources for a connection factory, queue, and topic.
2. Assigns either the queue or topic to a destination object, based on the specified destination type.
3. Creates a `Connection` and a `Session`.
4. Creates a `MessageConsumer`.
5. Creates an instance of the `TextListener` class and registers it as the message listener for the `MessageConsumer`:

```
listener = new TextListener(); consumer.setMessageListener(listener);
```

6. Starts the connection, causing message delivery to begin.

7. Listens for the messages published to the destination, stopping when the user types the character q or Q:

```
System.out.println("To end program, type Q or q, " + "then <return>");
inputStreamReader = new InputStreamReader(System.in);
while (!(answer == 'q' || answer == 'Q')) {
    try {
        answer = (char) inputStreamReader.read();
    } catch (IOException e) {
        System.out.println("I/O exception: " + e.toString());
    }
}
```

8. Closes the connection, which automatically closes the session and `MessageConsumer`.

The message listener, `asynchconsumer/src/java/TextListener.java`, follows these steps:

1. When a message arrives, the `onMessage` method is called automatically.
2. The `onMessage` method converts the incoming message to a `TextMessage` and displays its content. If the message is not a text message, it reports this fact:

```
public void onMessage(Message message) {
    TextMessage msg = null;
    try {
        if (message instanceof TextMessage) {
            msg = (TextMessage) message;
            System.out.println("Reading message: " + msg.getText());
        } else {
            System.out.println("Message is not a " + "TextMessage");
        }
    } catch (JMSEException e) {
        System.out.println("JMSEException in onMessage(): " + e.toString());
    } catch (Throwable t) {
        System.out.println("Exception in onMessage(): " + t.getMessage());
    }
}
```

You will use the connection factory and destinations you created in [“Creating JMS Administered Objects for the Synchronous Receive Example”](#) on page 650.

## Compiling and Packaging the `AsynchConsumer` Client

To compile and package the `AsynchConsumer` example using NetBeans IDE, follow these steps:

1. In NetBeans IDE, choose `Open Project` from the `File` menu.
2. In the `Open Project` dialog, navigate to `tut-install/examples/jms/simple/`.
3. Select the `asynchconsumer` folder.
4. Select the `Open as Main Project` check box.
5. Click `Open Project`.

6. Right-click the project and choose Build.

To compile and package the `AsynchConsumer` example using Ant, follow these steps:

1. In a terminal window, go to the `asynchconsumer` directory:

```
cd ../asynchconsumer
```

2. Type the following command:

```
ant
```

The targets package both the main class and the message listener class in the JAR file and place the file in the `dist` directory for the example.

## Running the Clients for the Asynchronous Receive Example

To run the clients using NetBeans IDE, follow these steps.

1. Run the `AsynchConsumer` example:
  - a. Right-click the `asynchconsumer` project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

```
topic
```

- d. Click OK.
- e. Right-click the project and choose Run.

The client displays the following lines and appears to hang:

```
Destination type is topic
To end program, type Q or q, then <return>
```

2. Now run the Producer example:
  - a. Right-click the `producer` project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

```
topic 3
```

- d. Click OK.
- e. Right-click the project and choose Run.

The output of the client looks like this:

```
Destination type is topic
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

In the other window, the `AsynchConsumer` client displays the following:

```
Destination type is topic
To end program, type Q or q, then <return>
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
Message is not a TextMessage
```

The last line appears because the client has received the non-text control message sent by the Producer client.

3. Type Q or q in the Output window and press Return to stop the client.
4. Now run the clients using a queue. In this case, as with the synchronous example, you can run the Producer client first, because there is no timing dependency between the sender and receiver.
  - a. Right-click the producer project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

```
queue 3
```

- d. Click OK.
- e. Right-click the project and choose Run.

The output of the client looks like this:

```
Destination type is queue
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

5. Run the AsyncConsumer client.
  - a. Right-click the asynchconsumer project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

```
queue
```

- d. Click OK.
- e. Right-click the project and choose Run.

The output of the client looks like this:

```
Destination type is queue
To end program, type Q or q, then <return>
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
Message is not a TextMessage
```

6. Type Q or q in the Output window and press Return to stop the client.

To run the clients using the `appclient` command, follow these steps:

1. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

2. Run the `AsynchConsumer` client, specifying the topic destination type.

```
appclient -client client-jar/asynchconsumerClient.jar topic
```

The client displays the following lines (along with some application client container output) and appears to hang:

```
Destination type is topic
To end program, type Q or q, then <return>
```

3. In the terminal window where you ran the `Producer` client previously, run the client again, sending three messages. The command looks like this:

```
appclient -client client-jar/producerClient.jar topic 3
```

The output of the client looks like this (along with some application client container output):

```
Destination type is topic
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

In the other window, the `AsynchConsumer` client displays the following (along with some application client container output):

```
Destination type is topic
To end program, type Q or q, then <return>
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
Message is not a TextMessage
```

The last line appears because the client has received the non-text control message sent by the `Producer` client.

4. Type `Q` or `q` and press `Return` to stop the client.
5. Now run the clients using a queue. In this case, as with the synchronous example, you can run the `Producer` client first, because there is no timing dependency between the sender and receiver:

```
appclient -client client-jar/producerClient.jar queue 3
```

The output of the client looks like this:

```
Destination type is queue
Sending message: This is message 1 from producer
Sending message: This is message 2 from producer
Sending message: This is message 3 from producer
```

6. Run the AsynchConsumer client:

```
appclient -client client-jar/asynchconsumerClient.jar queue
```

The output of the client looks like this (along with some application client container output):

```
Destination type is queue
To end program, type Q or q, then <return>
Reading message: This is message 1 from producer
Reading message: This is message 2 from producer
Reading message: This is message 3 from producer
Message is not a TextMessage
```

7. Type Q or q to stop the client.

## A Simple Example of Browsing Messages in a Queue

This section describes an example that creates a `QueueBrowser` object to examine messages on a queue, as described in “[JMS Queue Browsers](#)” on page 627. This section then explains how to compile, package, and run the example using the Enterprise Server.

The following sections describe the steps in creating and running the example:

- “[Writing the Client for the Queue Browser Example](#)” on page 660
- “[Compiling and Packaging the MessageBrowser Client](#)” on page 661
- “[Running the Clients for the Queue Browser Example](#)” on page 662

### Writing the Client for the Queue Browser Example

To create a `QueueBrowser` for a queue, you call the `Session.createBrowser` method with the queue as the argument. You obtain the messages in the queue as an `Enumeration` object. You can then iterate through the `Enumeration` object and display the contents of each message.

The `messagebrowser/src/java/MessageBrowser.java` client performs the following steps:

1. Injects resources for a connection factory and a queue.
2. Creates a `Connection` and a `Session`.
3. Creates a `QueueBrowser`:

```
QueueBrowser browser = session.createBrowser(queue);
```

4. Retrieves the `Enumeration` that contains the messages:

```
Enumeration msgs = browser.getEnumeration();
```

5. Verifies that the Enumeration contains messages, then displays the contents of the messages:

```
if ( !msgs.hasMoreElements() ) {
    System.out.println("No messages in queue");
} else {
    while (msgs.hasMoreElements()) {
        Message tempMsg = (Message)msgs.nextElement();
        System.out.println("Message: " + tempMsg);
    }
}
```

6. Closes the connection, which automatically closes the session and QueueBrowser.

The format in which the message contents appear is implementation-specific. In the Enterprise Server, the message format looks like this:

Message contents:

```
Text: This is message 3 from producer
Class: com.sun.messaging.jmq.jmsclient.TextMessageImpl
getJMSMessageID(): ID:14-129.148.71.199(f9:86:a2:d5:46:9b)-40814-1255980521747
getJMSTimestamp(): 1129061034355
getJMSCorrelationID(): null
JMSReplyTo: null
JMSDestination: PhysicalQueue
getJMSDeliveryMode(): PERSISTENT
getJMSRedelivered(): false
getJMSType(): null
getJMSExpiration(): 0
getJMSPriority(): 4
Properties: null
```

You will use the connection factory and queue you created in [“Creating JMS Administered Objects for the Synchronous Receive Example”](#) on page 650.

## Compiling and Packaging the MessageBrowser Client

To compile and package the MessageBrowser example using NetBeans IDE, follow these steps:

1. In NetBeans IDE, choose Open Project from the File menu.
2. In the Open Project dialog, navigate to *tut-install/examples/jms/simple/*.
3. Select the messagebrowser folder.
4. Select the Open as Main Project check box.
5. Click Open Project.
6. Right-click the project and choose Build.

To compile and package the MessageBrowser example using Ant, follow these steps:

1. In a terminal window, go to the messagebrowser directory.

```
cd ../messagebrowser
```

2. Type the following command:

```
ant
```

The targets place the application client JAR file in the `dist` directory for the example.

You also need the Producer example to send the message to the queue, and one of the consumer clients to consume the messages after you inspect them. If you did not do so already, package these examples.

## Running the Clients for the Queue Browser Example

To run the clients using NetBeans IDE, follow these steps.

1. Run the Producer client, sending one message to the queue:
  - a. Right-click the producer project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

```
queue
```

- d. Click OK.
- e. Right-click the project and choose Run.

The output of the client looks like this:

```
Destination type is queue  
Sending message: This is message 1 from producer
```

2. Run the MessageBrowser client. Right-click the messagebrowser project and choose Run.

The output of the client looks like this:

```
Message:  
Text: This is message 1 from producer  
Class: com.sun.messaging.jmq.jmsclient.TextMessageImpl  
getJMSMessageID(): ID:12-129.148.71.199(8c:34:4a:1a:1b:b8) -40883-1255980521747  
getJMSTimestamp(): 1129062957611  
getJMSCorrelationID(): null  
JMSReplyTo: null  
JMSDestination: PhysicalQueue  
getJMSDeliveryMode(): PERSISTENT  
getJMSRedelivered(): false  
getJMSType(): null  
getJMSExpiration(): 0  
getJMSPriority(): 4  
Properties: null  
Message:  
Class: com.sun.messaging.jmq.jmsclient.MessageImpl  
getJMSMessageID(): ID:13-129.148.71.199(8c:34:4a:1a:1b:b8) -40883-1255980521747  
getJMSTimestamp(): 1129062957616
```

```

getJMSCorrelationID(): null
JMSReplyTo: null
JMSDestination: PhysicalQueue
getJMSDeliveryMode(): PERSISTENT
getJMSRedelivered(): false
getJMSType(): null
getJMSExpiration(): 0
getJMSPriority(): 4
Properties: null

```

3. The first message is the `TextMessage`, and the second is the non-text control message.
4. Run the `SynchConsumer` client to consume the messages.
  - a. Right-click the `synchconsumer` project and choose Properties.
  - b. Select Run from the Categories tree.
  - c. In the Arguments field, type the following:

```
queue
```
  - d. Click OK.
  - e. Right-click the project and choose Run.

The output of the client looks like this:

```

Destination type is queue
Reading message: This is message 1 from producer

```

To run the clients using the `appclient` command, follow these steps. You may want to use two terminal windows.

1. Go to the producer directory.
2. Run the Producer client, sending one message to the queue:

```
appclient -cclient client-jar/producerClient.jar queue
```

The output of the client looks like this (along with some application client container output):

```

Destination type is queue
Sending message: This is message 1 from producer

```

3. Go to the `messagebrowser` directory.
4. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

5. Run the `MessageBrowser` client:

```
appclient -cclient client-jar/messagebrowserClient.jar
```

The output of the client looks like this (along with some application client container output):

```
Message:
Text: This is message 1 from producer
Class: com.sun.messaging.jmq.jmsclient.TextMessageImpl
getJMSMessageID(): ID:12-129.148.71.199(8c:34:4a:1a:1b:b8) -40883-1255980521747
getJMSTimestamp(): 1255980521747
getJMSCorrelationID(): null
JMSReplyTo: null
JMSDestination: PhysicalQueue
getJMSDeliveryMode(): PERSISTENT
getJMSRedelivered(): false
getJMSType(): null
getJMSExpiration(): 0
getJMSPriority(): 4
Properties: null
Message:
Class: com.sun.messaging.jmq.jmsclient.MessageImpl
getJMSMessageID(): ID:13-129.148.71.199(8c:34:4a:1a:1b:b8) -40883-1255980521767
getJMSTimestamp(): 1255980521767
getJMSCorrelationID(): null
JMSReplyTo: null
JMSDestination: PhysicalQueue
getJMSDeliveryMode(): PERSISTENT
getJMSRedelivered(): false
getJMSType(): null
getJMSExpiration(): 0
getJMSPriority(): 4
Properties: null
```

The first message is the `TextMessage`, and the second is the non-text control message.

6. Go to the `synchconsumer` directory.
7. Run the `SynchConsumer` client to consume the messages:

```
appclient -client client-jar/synchconsumerClient.jar queue
```

The output of the client looks like this (along with some application client container output):

```
Destination type is queue
Reading message: This is message 1 from producer
```

## Running JMS Clients on Multiple Systems

JMS clients that use the Enterprise Server can exchange messages with each other when they are running on different systems in a network. The systems must be visible to each other by name

(the UNIX host name or the Microsoft Windows computer name) and must both be running the Enterprise Server. You do not have to install the tutorial examples on both systems; you can use the examples installed on one system if you can access its file system from the other system.

---

**Note** – Any mechanism for exchanging messages between systems is specific to the Java EE server implementation. This tutorial describes how to use the Enterprise Server for this purpose.

---

Suppose that you want to run the Producer client on one system, `earth`, and the SynchronConsumer client on another system, `jupiter`. Before you can do so, you need to perform these tasks:

- Create two new connection factories
- Change the name of the default JMS host on one system
- Edit the source code for the two examples
- Recompile and repackage the examples

---

**Note** – A limitation in the JMS provider in the Enterprise Server may cause a runtime failure to create a connection to systems that use the Dynamic Host Configuration Protocol (DHCP) to obtain an IP address. You can, however, create a connection *from* a system that uses DHCP *to* a system that does not use DHCP. In the examples in this tutorial, `earth` can be a system that uses DHCP, and `jupiter` can be a system that does not use DHCP.

---

Before you begin, start the server on both systems:

1. Start the Enterprise Server on `earth`.
2. Start the Enterprise Server on `jupiter`.

## Creating Administered Objects for Multiple Systems

To run these clients, you must do the following:

- Create a new connection factory on both `earth` and `jupiter`
- Create a destination resource on both `earth` and `jupiter`

You do not have to install the tutorial examples on both systems, but you must be able to access the filesystem where it is installed. You may find it more convenient to install the tutorial examples on both systems if the two systems use different operating systems (for example, Windows and Solaris). Otherwise you will have to edit the file `tut-install/examples/bp-project/build.properties` and change the location of the `javaee.home` property each time you build or run a client on a different system.

To create a new connection factory on `jupiter`, perform these steps:

1. From a command shell on `jupiter`, go to the directory `tut-install/examples/jms/simple/producer/`.
2. Type the following command:

```
ant create-local-factory
```

The `create-local-factory` target, defined in the `build.xml` file for the `Producer` example, creates a connection factory named `jms/JupiterConnectionFactory`.

To create a new connection factory on `earth` that points to the connection factory on `jupiter`, perform these steps:

1. From a command shell on `earth`, go to the directory `tut-install/examples/jms/simple/producer/`.
2. Type the following command:

```
ant create-remote-factory -Dsys=remote-system-name
```

Replace `remote-system-name` with the actual name of the remote system.

The `create-remote-factory` target, defined in the `build.xml` file for the `Producer` example, also creates a connection factory named `jms/JupiterConnectionFactory`. In addition, it sets the `AddressList` property for this factory to the name of the remote system.

If you have already been working on either `earth` or `jupiter`, you have the queue and topic on one system. On the system that does not have the queue and topic, type the following command:

```
ant create-resources
```

When you run the clients, they will work as shown in [Figure 34-1](#). The client run on `earth` needs the queue on `earth` only in order that the resource injection will succeed. The connection, session, and message producer are all created on `jupiter` using the connection factory that points to `jupiter`. The messages sent from `earth` will be received on `jupiter`.

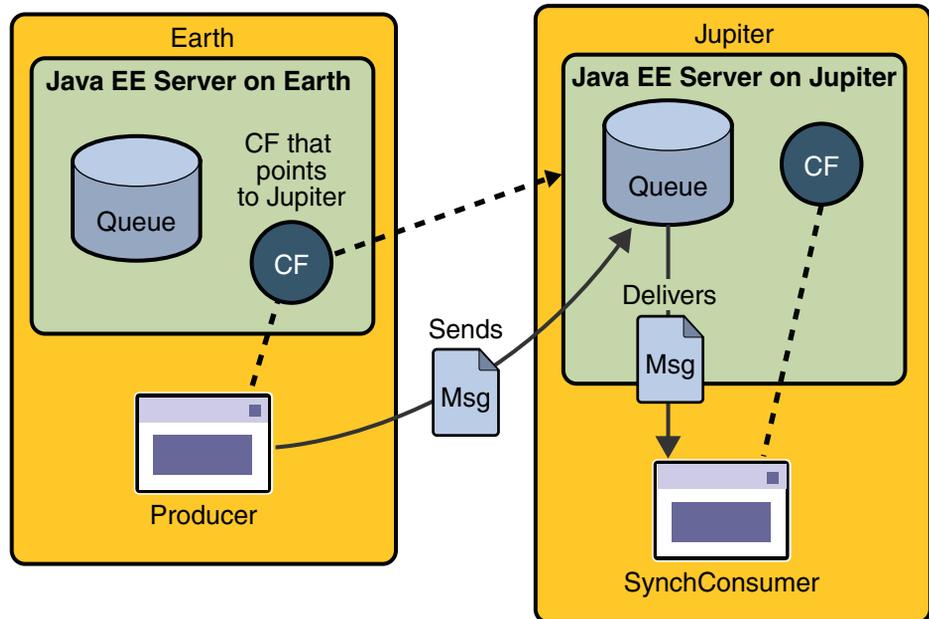


FIGURE 34-1 Sending Messages from One System to Another

## Changing the Default Host Name

By default, the default host name for the JMS service on the Enterprise Server is `localhost`. To access the JMS service from another system, however, you must change the host name. You can change it to either the actual host name or to `0.0.0.0`.

1. On `jupiter`, start the Administration Console by opening a browser at `http://localhost:4848/`.
2. In the navigation tree, expand the Configuration node.
3. Under the Configuration node, expand the Java Message Service node.
4. Under the Java Message Service node, expand the JMS Hosts node.
5. Under the JMS Hosts node, select `default_JMS_host`.

The Edit JMS Host page opens.

6. In the Host field, type the name of the system, or type `0.0.0.0`.
7. Click Save.
8. Restart the Enterprise Server.

To make this change using the `asadmin` command, use a command like one of the following:

```
asadmin set server-config.jms-service.jms-host.default_JMS_host.host="0.0.0.0"
```

```
asadmin set server-config.jms-service.jms-host.default_JMS_host.host="hostname"
```

## Editing, Recompiling, Repackaging, and Running the Clients

These steps assume that you have the tutorial installed on both of the two systems you are using and that you are able to access the file system of `jupiter` from `earth` or vice versa.

After you create the connection factories, edit the source files to specify the new connection factory. Then recompile, repackage, and run the clients. Perform the following steps:

1. On `earth`, open the following file in a text editor or in NetBeans IDE:

```
tut-install/examples/jms/simple/producer/src/java/Producer.java
```

2. Find the following line:

```
@Resource(mappedName="jms/ConnectionFactory")
```

3. Change the line to the following:

```
@Resource(mappedName="jms/JupiterConnectionFactory")
```

4. Recompile and repackage the `Producer` example on `earth`.

If you are using NetBeans IDE, right-click the `producer` project and choose `Clean and Build`.

If you are using Ant, type the following:

```
ant
```

5. On `jupiter`, open the following file in a text editor:

```
tut-install/examples/jms/simple/synchconsumer/src/java/SynchConsumer.java
```

6. Repeat steps 2 and 3.

7. Recompile and repackage the `SynchConsumer` example on `jupiter`.

If you are using NetBeans IDE, right-click the `synchconsumer` project and choose `Clean and Build`.

If you are using Ant, go to the `synchconsumer` directory and type:

```
ant
```

8. On `earth`, deploy and run `Producer`. If you are using NetBeans IDE on `earth`, perform these steps:

- a. Right-click the `producer` project and choose `Properties`.

- b. Select `Run` from the `Categories` tree.

- c. In the `Arguments` field, type the following:

```
queue 3
```

- d. Click `OK`.

- e. Right-click the project and choose `Run`.

If you are using `appclient`, perform these steps:

- a. On earth, from the producer directory, deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

- b. Type the following:

```
appclient -cclient client-jar/producerClient.jar queue 3
```

9. On jupiter, run SynchConsumer. If you are using NetBeans IDE on jupiter, perform these steps:

- a. Right-click the synchconsumer project and choose Properties.
- b. Select Run from the Categories tree.
- c. In the Arguments field, type the following:

```
queue
```

- d. Click OK.
- e. Right-click the project and choose Run.

If you are using appclient, perform these steps:

- a. From the synchconsumer directory, deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

- b. Type the following:

```
appclient -cclient client-jar/synchconsumerClient.jar queue
```

For examples showing how to deploy applications on two different systems, see [“An Application Example That Consumes Messages from a Remote Server” on page 695](#) and [“An Application Example That Deploys a Message-Driven Bean on Two Servers” on page 701](#).

## Deleting the Connection Factory and Stopping the Server

You will need the connection factory `jms/JupiterConnectionFactory` in [“An Application Example That Consumes Messages from a Remote Server” on page 695](#) and [“An Application Example That Deploys a Message-Driven Bean on Two Servers” on page 701](#). However, if you wish to delete it, go to the producer directory and type the following command:

```
ant delete-remote-factory
```

Remember to delete the connection factory on both systems.

You can also use Ant targets in the `producer/build.xml` file to delete the destinations and connection factories you created in “[Creating JMS Administered Objects for the Synchronous Receive Example](#)” on page 650. However, it is recommended that you keep them, because they will be used in most of the examples later in this chapter. After you have created them, they will be available whenever you restart the Enterprise Server.

To delete the class and JAR files for each client using NetBeans IDE, right-click each project and choose Clean.

To undeploy each client using Ant, go to the directory for the client and type the following:

```
ant undeploy
```

To delete the class and JAR files for each client using Ant, type the following:

```
ant clean
```

You can also stop the Enterprise Server, but you will need it to run the sample clients in the next section.

## Writing Robust JMS Applications

The following examples show how to use some of the more advanced features of the JMS API.

### A Message Acknowledgment Example

The `AckEquivExample.java` client shows how both of the following two scenarios ensure that a message will not be acknowledged until processing of it is complete:

- Using an asynchronous message consumer (a message listener) in an `AUTO_ACKNOWLEDGE` session
- Using a synchronous receiver in a `CLIENT_ACKNOWLEDGE` session

With a message listener, the automatic acknowledgment happens when the `onMessage` method returns (that is, after message processing has finished). With a synchronous receiver, the client acknowledges the message after processing is complete. If you use `AUTO_ACKNOWLEDGE` with a synchronous receive, the acknowledgment happens immediately after the receive call; if any subsequent processing steps fail, the message cannot be redelivered.

The example is in the following directory:

```
tut-install/examples/jms/advanced/ackequivexample/src/java/
```

The example contains an `AsynchSubscriber` class with a `TextListener` class, a `MultiplePublisher` class, a `SynchReceiver` class, a `SynchSender` class, a `main` method, and a method that runs the other classes' threads.

The example uses the following objects:

- `jms/ConnectionFactory`, `jms/Queue`, and `jms/Topic`: resources that you created in “Creating JMS Administered Objects for the Synchronous Receive Example” on page 650
- `jms/ControlQueue`: an additional queue
- `jms/DurableConnectionFactory`: a connection factory with a client ID (see “Creating Durable Subscriptions” on page 633, for more information)

To create the new queue and connection factory, you can use Ant targets defined in the file `tut-install/examples/jms/advanced/ackequivexample/build.xml`.

To run this example, follow these steps:

1. In a terminal window, go to the following directory:

```
tut-install/examples/jms/advanced/ackequivexample/
```

2. To create the objects needed in this example, type the following commands:

```
ant create-control-queue
ant create-durable-cf
```

3. To compile and package the client using NetBeans IDE, follow these steps:
  - a. In NetBeans IDE, choose Open Project from the File menu.
  - b. In the Open Project dialog, navigate to `tut-install/examples/jms/advanced/`.
  - c. Select the `ackequivexample` folder.
  - d. Select the Open as Main Project check box.
  - e. Click Open Project.
  - f. Right-click the project and choose Build.

To compile and package the client using Ant, type the following command:

```
ant
```

4. To run the client using NetBeans IDE, right-click the `ackequivexample` project and choose Run.

To run the client from the command line, follow these steps:

- a. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getClient
```

Ignore the message that states that the application is deployed at a URL.

- b. Type the following command:

```
appliant -cclient client-jar/ackequivexampleClient.jar
```

The client output looks something like this (along with some application client container output):

```
Queue name is jms/ControlQueue
Queue name is jms/Queue
Topic name is jms/Topic
Connection factory name is jms/DurableConnectionFactory
  SENDER: Created client-acknowledge session
  SENDER: Sending message: Here is a client-acknowledge message
  RECEIVER: Created client-acknowledge session
  RECEIVER: Processing message: Here is a client-acknowledge message
  RECEIVER: Now I'll acknowledge the message
SUBSCRIBER: Created auto-acknowledge session
SUBSCRIBER: Sending synchronize message to control queue
PUBLISHER: Created auto-acknowledge session
PUBLISHER: Receiving synchronize messages from control queue; count = 1
PUBLISHER: Received synchronize message; expect 0 more
PUBLISHER: Publishing message: Here is an auto-acknowledge message 1
PUBLISHER: Publishing message: Here is an auto-acknowledge message 2
SUBSCRIBER: Processing message: Here is an auto-acknowledge message 1
PUBLISHER: Publishing message: Here is an auto-acknowledge message 3
SUBSCRIBER: Processing message: Here is an auto-acknowledge message 2
SUBSCRIBER: Processing message: Here is an auto-acknowledge message 3
```

After you run the client, you can delete the destination resource `jms/ControlQueue`. Go to the directory `tut-install/examples/jms/advanced/ackequivexample/` and type the following command:

```
ant delete-control-queue
```

You will need the other resources for other examples.

To delete the class and JAR files for the client using NetBeans IDE, right-click the project and choose Clean.

To undeploy the client using Ant, type the following:

```
ant undeploy
```

To delete the class and JAR files for the client using Ant, type the following:

```
ant clean
```

## A Durable Subscription Example

The `DurableSubscriberExample.java` example shows how durable subscriptions work. It demonstrates that a durable subscription is active even when the subscriber is not active. The example contains a `DurableSubscriber` class, a `MultiplePublisher` class, a main method, and a method that instantiates the classes and calls their methods in sequence.

The example is in the following directory:

*tut-install/examples/jms/advanced/durablesubscriberexample/src/java/*

The example begins in the same way as any publish/subscribe client: The subscriber starts, the publisher publishes some messages, and the subscriber receives them. At this point, the subscriber closes itself. The publisher then publishes some messages while the subscriber is not active. The subscriber then restarts and receives the messages.

Before you run this example, compile and package the source file and create a connection factory that has a client ID. Perform the following steps:

1. To compile and package the client using NetBeans IDE, follow these steps:
  - a. In NetBeans IDE, choose Open Project from the File menu.
  - b. In the Open Project dialog, navigate to *tut-install/examples/jms/advanced/*.
  - c. Select the *durablesubscriberexample* folder.
  - d. Select the Open as Main Project check box.
  - e. Click Open Project.
  - f. Right-click the project and choose Build.

To compile and package the client using Ant, follow these steps:

- a. Go to the following directory:

```
tut-install/examples/jms/advanced/durablesubscriberexample/
```

- b. Type the following command:

```
ant
```

2. If you did not do so for “[A Message Acknowledgment Example](#)” on page 670, create a connection factory named `jms/DurableConnectionFactory`:

```
ant create-durable-cf
```

To run the client using NetBeans IDE, right-click the *durablesubscriberexample* project and choose Run.

To run the client from the command line, follow these steps:

1. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

2. Type the following command:

```
appclient -client client-jar/durablesubscriberexampleClient.jar
```

The output looks something like this (along with some application client container output):

```
Connection factory without client ID is jms/ConnectionFactory
Connection factory with client ID is jms/DurableConnectionFactory
Topic name is jms/Topic
```

```
Starting subscriber
PUBLISHER: Publishing message: Here is a message 1
SUBSCRIBER: Reading message: Here is a message 1
PUBLISHER: Publishing message: Here is a message 2
SUBSCRIBER: Reading message: Here is a message 2
PUBLISHER: Publishing message: Here is a message 3
SUBSCRIBER: Reading message: Here is a message 3
Closing subscriber
PUBLISHER: Publishing message: Here is a message 4
PUBLISHER: Publishing message: Here is a message 5
PUBLISHER: Publishing message: Here is a message 6
Starting subscriber
SUBSCRIBER: Reading message: Here is a message 4
SUBSCRIBER: Reading message: Here is a message 5
SUBSCRIBER: Reading message: Here is a message 6
Closing subscriber
Unsubscribing from durable subscription
```

After you run the client, you can delete the connection factory `jms/DurableConnectionFactory`. Go to the directory `tut-install/examples/jms/advanced/durablesubscriberexample/` and type the following command:

```
ant delete-durable-cf
```

To delete the class and JAR files for the client using NetBeans IDE, right-click the project and choose Clean.

To undeploy the client using Ant, type the following:

```
ant undeploy
```

To delete the class and JAR files for the client using Ant, type the following:

```
ant clean
```

## A Local Transaction Example

The `TransactedExample.java` example demonstrates the use of transactions in a JMS client application. The example is in the following directory:

```
tut-install/examples/jms/advanced/transactedexample/src/java/
```

This example shows how to use a queue and a topic in a single transaction as well as how to pass a session to a message listener's constructor function. The example represents a highly simplified e-commerce application in which the following things happen.

1. A retailer sends a `MapMessage` to the vendor order queue, ordering a quantity of computers, and waits for the vendor's reply:

```
producer = session.createProducer(vendorOrderQueue);
outMessage = session.createMapMessage();
outMessage.setString("Item", "Computer(s)");
outMessage.setInt("Quantity", quantity);
outMessage.setJMSReplyTo(retailerConfirmQueue);
producer.send(outMessage);
System.out.println("Retailer: ordered " + quantity + " computer(s)");
orderConfirmReceiver = session.createConsumer(retailerConfirmQueue);
connection.start();
```

2. The vendor receives the retailer's order message and sends an order message to the supplier order topic in one transaction. This JMS transaction uses a single session, so you can combine a receive from a queue with a send to a topic. Here is the code that uses the same session to create a consumer for a queue and a producer for a topic:

```
vendorOrderReceiver = session.createConsumer(vendorOrderQueue);
supplierOrderProducer = session.createProducer(supplierOrderTopic);
```

The following code receives the incoming message, sends an outgoing message, and commits the session. The message processing has been removed to keep the sequence simple:

```
inMessage = vendorOrderReceiver.receive();
// Process the incoming message and format the outgoing
// message
...
supplierOrderProducer.send(orderMessage);
...
session.commit();
```

3. Each supplier receives the order from the order topic, checks its inventory, and then sends the items ordered to the queue named in the order message's `JMSReplyTo` field. If it does not have enough in stock, the supplier sends what it has. The synchronous receive from the topic and the send to the queue take place in one JMS transaction.

```
receiver = session.createConsumer(orderTopic);
...
inMessage = receiver.receive();
if (inMessage instanceof MapMessage) {
    orderMessage = (MapMessage) inMessage;
}
// Process message
MessageProducer producer =
    session.createProducer((Queue) orderMessage.getJMSReplyTo());
```

```
outMessage = session.createMapMessage();
// Add content to message
producer.send(outMessage);
// Display message contentsession.commit();
```

4. The vendor receives the replies from the suppliers from its confirmation queue and updates the state of the order. Messages are processed by an asynchronous message listener; this step shows the use of JMS transactions with a message listener.

```
MapMessage component = (MapMessage) message;
...
orderNumber = component.getInt("VendorOrderNumber");
Order order = Order.getOrder(orderNumber).processSubOrder(component);
session.commit();
```

5. When all outstanding replies are processed for a given order, the vendor message listener sends a message notifying the retailer whether it can fulfill the order.

```
Queue replyQueue = (Queue) order.order.getJMSReplyTo();
MessageProducer producer = session.createProducer(replyQueue);
MapMessage retailerConfirmMessage = session.createMapMessage();
// Format the message
producer.send(producerConfirmMessage);
session.commit();
```

6. The retailer receives the message from the vendor:

```
inMessage = (MapMessage) orderConfirmReceiver.receive();
```

[Figure 34–2](#) illustrates these steps.

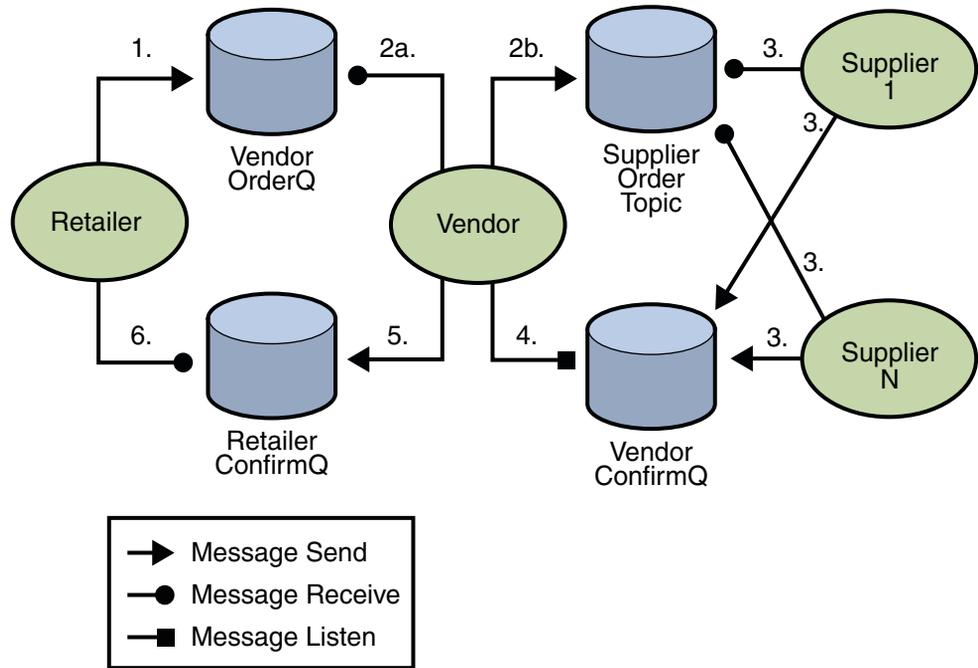


FIGURE 34-2 Transactions: JMS Client Example

The example contains five classes: `GenericSupplier`, `Order`, `Retailer`, `Vendor`, and `VendorMessageListener`. The example also contains a main method and a method that runs the threads of the `Retailer`, `Vendor`, and two supplier classes.

All the messages use the `MapMessage` message type. Synchronous receives are used for all message reception except for the case of the vendor processing the replies of the suppliers. These replies are processed asynchronously and demonstrate how to use transactions within a message listener.

At random intervals, the `Vendor` class throws an exception to simulate a database problem and cause a rollback.

All classes except `Retailer` use transacted sessions.

The example uses three queues named `jms/AQueue`, `jms/BQueue`, and `jms/CQueue`, and one topic named `jms/OTopic`.

Before you run the example, do the following:

1. In a terminal window, go to the following directory:
 

```
tut-install/examples/jms/advanced/transactedexample/
```
2. Create the necessary resources using the following command:

**ant create-resources**

This command creates three destination resources with the names `jms/AQueue`, `jms/BQueue`, and `jms/CQueue`, all of type `javax.jms.Queue`, and one destination resource with the name `jms/OTopic`, of type `javax.jms.Topic`.

3. To compile and package the client using NetBeans IDE, follow these steps:
  - a. In NetBeans IDE, choose Open Project from the File menu.
  - b. In the Open Project dialog, navigate to `tut-install/examples/jms/advanced/`.
  - c. Select the `transactedexample` folder.
  - d. Select the Open as Main Project check box.
  - e. Click Open Project.
  - f. Right-click the project and choose Build.

To compile and package the client using Ant, follow these steps:

- a. Go to the following directory:  
`tut-install/examples/jms/advanced/transactedexample/`
- b. Type the following command:

```
ant
```

To run the client using NetBeans IDE, follow these steps:

1. Right-click the `transactedexample` project and choose Properties.
2. Select Run from the Categories tree.
3. In the Arguments field, type a number that specifies the number of computers to order:  
`3`
4. Click OK.
5. Right-click the project and choose Run.

To run the client from the command line, follow these steps:

1. Deploy the client JAR file to the Enterprise Server, then retrieve the client stubs:

```
ant getclient
```

Ignore the message that states that the application is deployed at a URL.

2. Use a command like the following to run the client. The argument specifies the number of computers to order:

```
appclient -client client-jar/transactedexampleClient.jar 3
```

The output looks something like this (along with some application client container output):

```
Quantity to be ordered is 3
Retailer: ordered 3 computer(s)
Vendor: Retailer ordered 3 Computer(s)
```

```

Vendor: ordered 3 monitor(s) and hard drive(s)
Monitor Supplier: Vendor ordered 3 Monitor(s)
Monitor Supplier: sent 3 Monitor(s)
  Monitor Supplier: committed transaction
  Vendor: committed transaction 1
Hard Drive Supplier: Vendor ordered 3 Hard Drive(s)
Hard Drive Supplier: sent 1 Hard Drive(s)
Vendor: Completed processing for order 1
  Hard Drive Supplier: committed transaction
Vendor: unable to send 3 computer(s)
  Vendor: committed transaction 2
Retailer: Order not filled
Retailer: placing another order
Retailer: ordered 6 computer(s)
Vendor: JMSException occurred: javax.jms.JMSException:
Simulated database concurrent access exception
javax.jms.JMSException: Simulated database concurrent access exception
    at TransactedExample$Vendor.run(Unknown Source)
  Vendor: rolled back transaction 1
Vendor: Retailer ordered 6 Computer(s)
Vendor: ordered 6 monitor(s) and hard drive(s)
Monitor Supplier: Vendor ordered 6 Monitor(s)
Hard Drive Supplier: Vendor ordered 6 Hard Drive(s)
Monitor Supplier: sent 6 Monitor(s)
  Monitor Supplier: committed transaction
Hard Drive Supplier: sent 6 Hard Drive(s)
  Hard Drive Supplier: committed transaction
  Vendor: committed transaction 1
Vendor: Completed processing for order 2
Vendor: sent 6 computer(s)
Retailer: Order filled
  Vendor: committed transaction 2

```

After you run the client, you can delete the physical destinations and the destination resources. Go to the directory *tut-install/examples/jms/advanced/ttransactedexample/* and type the following command:

```
ant delete-resources
```

To undeploy the client using Ant, type the following:

```
ant undeploy
```

To delete the class and JAR files for the client using Ant, type the following:

```
ant clean
```

## An Application That Uses the JMS API with a Session Bean

This section explains how to write, compile, package, deploy, and run an application that uses the JMS API in conjunction with a session bean. The application contains the following components:

- An application client that invokes a session bean
- A session bean that publishes several messages to a topic
- A message-driven bean that receives and processes the messages using a durable topic subscriber and a message selector

The section covers the following topics:

- [“Writing the Application Components for the `clientsessionmdb` Example” on page 680](#)
- [“Creating Resources for the `clientsessionmdb` Example” on page 683](#)
- [“Building, Deploying, and Running the `clientsessionmdb` Example Using NetBeans IDE” on page 683](#)
- [“Building, Deploying, and Running the `clientsessionmdb` Example Using Ant” on page 685](#)

You will find the source files for this section in the directory `tut-install/examples/jms/clientsessionmdb/`. Path names in this section are relative to this directory.

## Writing the Application Components for the `clientsessionmdb` Example

This application demonstrates how to send messages from an enterprise bean (in this case, a session bean) rather than from an application client, as in the example in [Chapter 19, “A Message-Driven Bean Example.”](#) [Figure 34–3](#) illustrates the structure of this application.

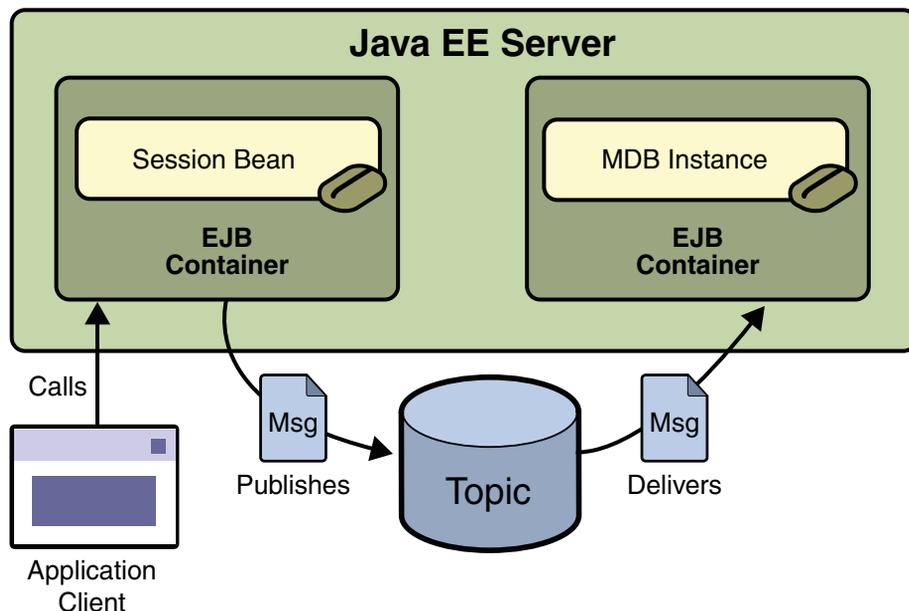


FIGURE 34-3 An Enterprise Bean Application: Client to Session Bean to Message-Driven Bean

The Publisher enterprise bean in this example is the enterprise-application equivalent of a wire-service news feed that categorizes news events into six news categories. The message-driven bean could represent a newsroom, where the sports desk, for example, would set up a subscription for all news events pertaining to sports.

The application client in the example injects the Publisher enterprise bean’s remote home interface and then calls the bean’s business method. The enterprise bean creates 18 text messages. For each message, it sets a `String` property randomly to one of six values representing the news categories and then publishes the message to a topic. The message-driven bean uses a message selector for the property to limit which of the published messages it receives.

Writing the components of the application involves the following:

- “Coding the Application Client: `MyAppClient.java`” on page 681
- “Coding the Publisher Session Bean” on page 682
- “Coding the Message-Driven Bean: `MessageBean.java`” on page 682

### Coding the Application Client: `MyAppClient.java`

The application client, `clientsessionmdb-app-client/src/java/MyAppClient.java`, performs no JMS API operations and so is simpler than the client in [Chapter 19, “A Message-Driven Bean Example.”](#) The client uses dependency injection to obtain the Publisher enterprise bean’s business interface:

```
@EJB(name="PublisherRemote")
static private PublisherRemote publisher;
```

The client then calls the bean's business method twice.

## Coding the Publisher Session Bean

The Publisher bean is a stateless session bean that has one business method. The Publisher bean uses a remote interface rather than a local interface because it is accessed from the application client.

The remote interface, `clientsessionmdb-ejb/src/java/sb/PublisherRemote.java`, declares a single business method, `publishNews`.

The bean class, `clientsessionmdb-ejb/src/java/sb/PublisherBean.java`, implements the `publishNews` method and its helper method `chooseType`. The bean class also injects `SessionContext`, `ConnectionFactory`, and `Topic` resources and implements `@PostConstruct` and `@PreDestroy` callback methods. The bean class begins as follows:

```
@Stateless
@Remote({PublisherRemote.class})
public class PublisherBean implements PublisherRemote {

    @Resource
    private SessionContext sc;

    @Resource(mappedName="jms/ConnectionFactory")
    private ConnectionFactory connectionFactory;

    @Resource(mappedName="jms/Topic")
    private Topic topic;
    ...
}
```

The `@PostConstruct` callback method of the bean class, `makeConnection`, creates the `Connection` used by the bean. The business method `publishNews` creates a `Session` and a `MessageProducer` and publishes the messages.

The `@PreDestroy` callback method, `endConnection`, deallocates the resources that were allocated by the `@PostConstruct` callback method. In this case, the method closes the `Connection`.

## Coding the Message-Driven Bean: `MessageBean.java`

The message-driven bean class, `clientsessionmdb-ejb/src/java/mdb/MessageBean.java`, is almost identical to the one in [Chapter 19, "A Message-Driven Bean Example."](#) However, the `@MessageDriven` annotation is different, because instead of a queue the bean is using a topic

with a durable subscription, and it is also using a message selector. Therefore, the annotation sets the activation config properties `messageSelector`, `subscriptionDurability`, `clientId`, and `subscriptionName`, as follows:

```
@MessageDriven(mappedName="jms/Topic",
activationConfig=
{ @ActivationConfigProperty(propertyName="messageSelector",
    propertyValue="NewsType = 'Sports' OR NewsType = 'Opinion'"),
  @ActivationConfigProperty(
    propertyName="subscriptionDurability",
    propertyValue="Durable"),
  @ActivationConfigProperty(propertyName="clientId",
    propertyValue="MyID"),
  @ActivationConfigProperty(propertyName="subscriptionName",
    propertyValue="MySub")
})
```

The JMS resource adapter uses these properties to create a connection factory for the message-driven bean that allows the bean to use a durable subscriber.

## Creating Resources for the `clientsessionmdb` Example

This example uses the topic named `jms/Topic` and the connection factory `jms/ConnectionFactory`, which you created in [“Creating JMS Administered Objects for the Synchronous Receive Example” on page 650](#). If you deleted the connection factory or topic, you can create them again using targets in the `build.xml` file for this example. Use the following commands to create the resources:

```
ant create-cf
ant create-topic
```

## Building, Deploying, and Running the `clientsessionmdb` Example Using NetBeans IDE

To build, deploy, and run the application using NetBeans IDE, do the following:

1. In NetBeans IDE, choose Open Project from the File menu.
2. In the Open Project dialog, navigate to `tut-install/examples/jms/`.
3. Select the `clientsessionmdb` folder.
4. Select the Open as Main Project check box and the Open Required Projects check box.
5. Click Open Project.

6. Right-click the `clientsessionmdb` project and choose Build.

This task creates the following:

- An application client JAR file that contains the client class file and the session bean's remote interface, along with a manifest file that specifies the main class and places the EJB JAR file in its classpath
- An EJB JAR file that contains both the session bean and the message-driven bean
- An application EAR file that contains the two JAR files

7. Right-click the project and choose Run.

This command deploys the project, returns a JAR file named `clientsessionmdbClient.jar`, and then executes it.

The output of the application client in the Output pane looks like this (preceded by application client container output):

```
To view the bean output,  
check <install_dir>/domains/domain1/logs/server.log.
```

The output from the enterprise beans appears in the server log (*domain-dir/logs/server.log*), wrapped in logging information. The Publisher session bean sends two sets of 18 messages numbered 0 through 17. Because of the message selector, the message-driven bean receives only the messages whose `NewsType` property is `Sports` or `Opinion`.

Undeploy the application after you finish running the client. To undeploy the application, follow these steps:

1. Click the Services tab.
2. Expand the Servers node.
3. Expand the GlassFish v3 Domain node.
4. Expand the Applications node.

You may need to right-click the node and choose Refresh to see the contents of the Applications node.

5. Right-click `clientsessionmdb` and choose Undeploy.

To remove the generated files, click the Projects tab, then right-click the `clientsessionmdb` project and choose Clean.

## Building, Deploying, and Running the `clientsessionmdb` Example Using Ant

To build the application using Ant, do the following:

1. Go to the following directory:

```
tut-install/examples/jms/clientsessionmdb/
```

2. To compile the source files and package the application, use the following command:

```
ant
```

The ant command creates the following:

- An application client JAR file that contains the client class file and the session bean's remote interface, along with a manifest file that specifies the main class and places the EJB JAR file in its classpath
- An EJB JAR file that contains both the session bean and the message-driven bean
- An application EAR file that contains the two JAR files

The `clientsessionmdb.ear` file is created in the `dist` directory.

To deploy the application and run the client, use the following command:

```
ant run
```

Ignore the message that states that the application is deployed at a URL.

The client displays these lines (preceded by application client container output):

```
To view the bean output,  
check <install_dir>/domains/domain1/logs/server.log.
```

The output from the enterprise beans appears in the server log file, wrapped in logging information. The Publisher session bean sends two sets of 18 messages numbered 0 through 17. Because of the message selector, the message-driven bean receives only the messages whose `NewsType` property is `Sports` or `Opinion`.

Undeploy the application after you finish running the client. Use the following command:

```
ant undeploy
```

To remove the generated files, use the following command in the `clientsessionmdb`, `clientsessionmdb-app-client`, and `clientsessionmdb-ejb` directories:

```
ant clean
```

## An Application That Uses the JMS API with an Entity

This section explains how to write, compile, package, deploy, and run an application that uses the JMS API with an entity. The application uses the following components:

- An application client that both sends and receives messages
- Two message-driven beans
- An entity class

This section covers the following topics:

- “[Overview of the `clientmdbentity` Example Application](#)” on page 686
- “[Writing the Application Components for the `clientmdbentity` Example](#)” on page 688
- “[Creating Resources for the `clientmdbentity` Example](#)” on page 690
- “[Building, Deploying, and Running the `clientmdbentity` Example Using NetBeans IDE](#)” on page 691
- “[Building, Deploying, and Running the `clientmdbentity` Example Using Ant](#)” on page 693

You will find the source files for this section in the directory `tut-install/examples/jms/clientmdbentity/`. Path names in this section are relative to this directory.

### Overview of the `clientmdbentity` Example Application

This application simulates, in a simplified way, the work flow of a company’s human resources (HR) department when it processes a new hire. This application also demonstrates how to use the Java EE platform to accomplish a task that many JMS applications need to perform.

A JMS client must often wait for several messages from various sources. It then uses the information in all these messages to assemble a message that it then sends to another destination. The common term for this process is *joining messages*. Such a task must be transactional, with all the receives and the send as a single transaction. If not all the messages are received successfully, the transaction can be rolled back. For an application client example that illustrates this task, see “[A Local Transaction Example](#)” on page 674.

A message-driven bean can process only one message at a time in a transaction. To provide the ability to join messages, an application can have the message-driven bean store the interim information in an entity. The entity can then determine whether all the information has been received; when it has, the entity can report this back to one of the message-driven beans, which then creates and sends the message to the other destination. After it has completed its task, the entity can be removed.

The basic steps of the application are as follows.

1. The HR department's application client generates an employee ID for each new hire and then publishes a message (M1) containing the new hire's name, employee ID, and position. The client then creates a temporary queue, `ReplyQueue`, with a message listener that waits for a reply to the message. (See [“Creating Temporary Destinations”](#) on page 632 for more information.)
2. Two message-driven beans process each message: One bean, `OfficeMDB`, assigns the new hire's office number, and the other bean, `EquipmentMDB`, assigns the new hire's equipment. The first bean to process the message creates and persists an entity named `SetupOffice`, then calls a business method of the entity to store the information it has generated. The second bean locates the existing entity and calls another business method to add its information.
3. When both the office and the equipment have been assigned, the entity business method returns a value of `true` to the message-driven bean that called the method. The message-driven bean then sends to the reply queue a message (M2) describing the assignments. Then it removes the entity. The application client's message listener retrieves the information.

[Figure 34-4](#) illustrates the structure of this application. Of course, an actual HR application would have more components; other beans could set up payroll and benefits records, schedule orientation, and so on.

[Figure 34-4](#) assumes that `OfficeMDB` is the first message-driven bean to consume the message from the client. `OfficeMDB` then creates and persists the `SetupOffice` entity and stores the office information. `EquipmentMDB` then finds the entity, stores the equipment information, and learns that the entity has completed its work. `EquipmentMDB` then sends the message to the reply queue and removes the entity.

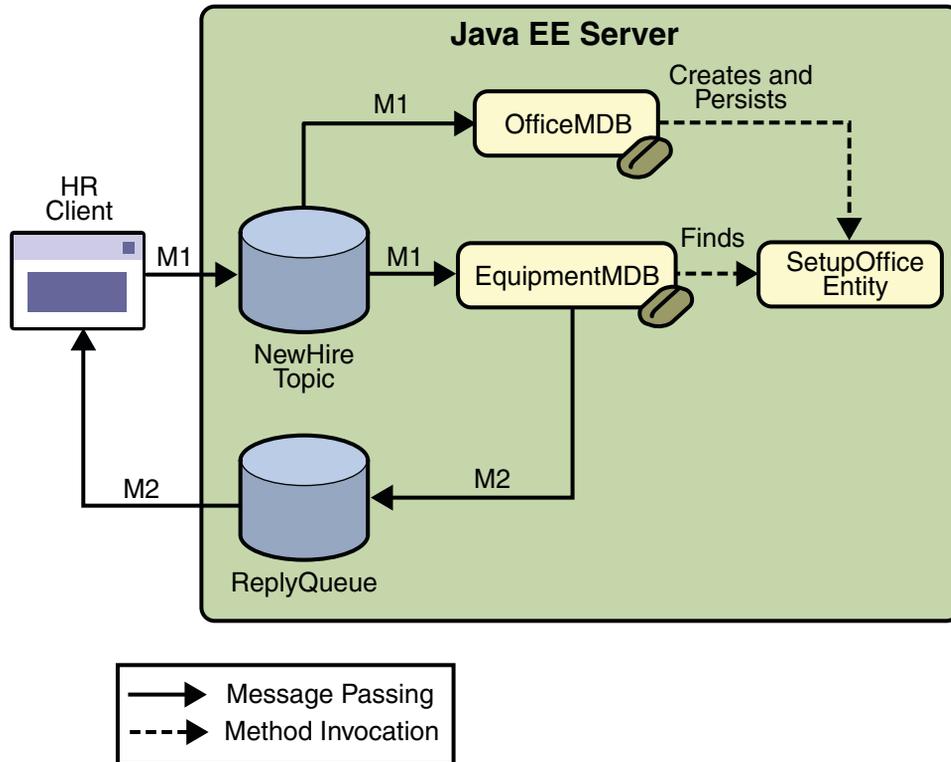


FIGURE 34-4 An Enterprise Bean Application: Client to Message-Driven Beans to Entity

## Writing the Application Components for the clientmdbentity Example

Writing the components of the application involves the following:

- “Coding the Application Client: `HumanResourceClient.java`” on page 688
- “Coding the Message-Driven Beans for the `clientmdbentity` Example” on page 689
- “Coding the Entity Class for the `clientmdbentity` Example” on page 689

### Coding the Application Client: `HumanResourceClient.java`

The application client, `clientmdbentity-app-client/src/java/HumanResourceClient.java`, performs the following steps:

1. Injects `ConnectionFactory` and `Topic` resources
2. Creates a `TemporaryQueue` to receive notification of processing that occurs, based on new-hire events it has published

3. Creates a `MessageConsumer` for the `TemporaryQueue`, sets the `MessageConsumer`'s message listener, and starts the connection
4. Creates a `MessageProducer` and a `MapMessage`
5. Creates five new employees with randomly generated names, positions, and ID numbers (in sequence) and publishes five messages containing this information

The message listener, `HRListener`, waits for messages that contain the assigned office and equipment for each employee. When a message arrives, the message listener displays the information received and determines whether all five messages have arrived. When they have, the message listener notifies the main method, which then exits.

## Coding the Message-Driven Beans for the `clientmdbentity` Example

This example uses two message-driven beans:

- `clientmdbentity-ejb/src/java/EquipmentMDB.java`
- `clientmdbentity-ejb/src/java/OfficeMDB.java`

The beans take the following steps:

1. They inject `MessageDrivenContext` and `ConnectionFactory` resources.
2. The `onMessage` method retrieves the information in the message. The `EquipmentMDB`'s `onMessage` method chooses equipment, based on the new hire's position; the `OfficeMDB`'s `onMessage` method randomly generates an office number.
3. After a slight delay to simulate real world processing hitches, the `onMessage` method calls a helper method, `compose`.
4. The `compose` method takes the following steps:
  - a. It either creates and persists the `SetupOffice` entity or finds it by primary key.
  - b. It uses the entity to store the equipment or the office information in the database, calling either the `doEquipmentList` or the `doOfficeNumber` business method.
  - c. If the business method returns `true`, meaning that all of the information has been stored, it creates a connection and a session, retrieves the reply destination information from the message, creates a `MessageProducer`, and sends a reply message that contains the information stored in the entity.
  - d. It removes the entity.

## Coding the Entity Class for the `clientmdbentity` Example

The `SetupOffice` class, `SetupOffice.java`, is an entity class. The entity and the message-driven beans are packaged together in an EJB JAR file. The entity class is declared as follows:

```
@Entity
public class SetupOffice implements Serializable {
```

The class contains a no-argument constructor and a constructor that takes two arguments, the employee ID and name. It also contains getter and setter methods for the employee ID, name, office number, and equipment list. The getter method for the employee ID has the `@Id` annotation to indicate that this field is the primary key:

```
@Id public String getEmployeeId() {
    return id;
}
```

The class also implements the two business methods, `doEquipmentList` and `doOfficeNumber`, and their helper method, `checkIfSetupComplete`.

The message-driven beans call the business methods and the getter methods.

The `persistence.xml` file for the entity specifies the most basic settings:

```
<?xml version="1.0" encoding="UTF-8"?>
<persistence version="2.0"
    xmlns="http://java.sun.com/xml/ns/persistence"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://java.sun.com/xml/ns/persistence http://java.sun.com/xml/ns/persistence/persistence_2_0.xsd">
  <persistence-unit name="clientmdbentity-ejbPU" transaction-type="JTA">
    <provider>org.eclipse.persistence.jpa.PersistenceProvider</provider>
    <jta-data-source>jdbc/___default</jta-data-source>
    <properties>
      <property name="eclipselink.ddl-generation" value="drop-and-create-tables"/>
    </properties>
  </persistence-unit>
</persistence>
```

## Creating Resources for the `clientmdbentity` Example

This example uses the connection factory `jms/ConnectionFactory` and the topic `jms/Topic`, both of which you used in [“An Application That Uses the JMS API with a Session Bean” on page 680](#). It also uses the JDBC resource named `jdbc/___default`, which is enabled by default when you start the Enterprise Server.

If you deleted the connection factory or topic, you can create them again using targets in the `build.xml` file for this example. Use the following commands to create the resources:

```
ant create-cf
ant create-topic
```

## Building, Deploying, and Running the `clientmdbentity` Example Using NetBeans IDE

To build, deploy, and run the application using NetBeans IDE, do the following:

1. In NetBeans IDE, choose Open Project from the File menu.
2. In the Open Project dialog, navigate to `tut-install/examples/jms/`.
3. Select the `clientmdbentity` folder.
4. Select the Open as Main Project check box and the Open Required Projects check box.
5. Click Open Project.
6. Right-click the `clientmdbentity` project and choose Build.

This task creates the following:

- An application client JAR file that contains the client class and listener class files, along with a manifest file that specifies the main class
  - An EJB JAR file that contains the message-driven beans and the entity class, along with the `persistence.xml` file
  - An application EAR file that contains the two JAR files along with an `application.xml` file
7. Right-click the project and choose Run.

This command deploys the project, returns a client JAR file named `clientmdbentityClient.jar` and then executes it.

---

**Note** – Because NetBeans IDE starts the database at the same time it starts the Enterprise Server, you do not need to start the database explicitly.

---

The output of the application client in the Output pane looks something like this:

```
PUBLISHER: Setting hire ID to 25, name Gertrude Bourbon, position Senior Programmer
PUBLISHER: Setting hire ID to 26, name Jack Verdon, position Manager
PUBLISHER: Setting hire ID to 27, name Fred Tudor, position Manager
PUBLISHER: Setting hire ID to 28, name Fred Martin, position Programmer
PUBLISHER: Setting hire ID to 29, name Mary Stuart, position Manager
Waiting for 5 message(s)
New hire event processed:
  Employee ID: 25
  Name: Gertrude Bourbon
  Equipment: Laptop
  Office number: 183
Waiting for 4 message(s)
New hire event processed:
```

```
Employee ID: 26
Name: Jack Verdon
Equipment: Pager
Office number: 20
Waiting for 3 message(s)
New hire event processed:
Employee ID: 27
Name: Fred Tudor
Equipment: Pager
Office number: 51
Waiting for 2 message(s)
New hire event processed:
Employee ID: 28
Name: Fred Martin
Equipment: Desktop System
Office number: 141
Waiting for 1 message(s)
New hire event processed:
Employee ID: 29
Name: Mary Stuart
Equipment: Pager
Office number: 238
```

The output from the message-driven beans and the entity class appears in the server log, wrapped in logging information.

For each employee, the application first creates the entity and then finds it. You may see runtime errors in the server log, and transaction rollbacks may occur. The errors occur if both of the message-driven beans discover at the same time that the entity does not yet exist, so they both try to create it. The first attempt succeeds, but the second fails because the bean already exists. After the rollback, the second message-driven bean tries again and succeeds in finding the entity. Container-managed transactions allow the application to run correctly, in spite of these errors, with no special programming.

You can run the application client repeatedly.

Undeploy the application after you finish running the client. To undeploy the application, follow these steps:

1. Click the Services tab.
2. Expand the Servers node.
3. Expand the GlassFish v3 Domain node.
4. Expand or refresh the Applications node.
5. Right-click `clientmdbentity` and choose Undeploy.

To remove the generated files, right-click the `clientmdbentity`, `clientmdbentity-app-client`, and `clientmdbentity-ejb` projects in turn and choose Clean.

## Building, Deploying, and Running the `clientmdbentity` Example Using Ant

To create and package the application using Ant, perform these steps:

1. Start the Enterprise Server, if it is not already running.
2. Go to the following directory:

```
tut-install/examples/jms/clientmdbentity/
```

3. To compile the source files and package the application, use the following command:

```
ant
```

The ant command creates the following:

- An application client JAR file that contains the client class and listener class files, along with a manifest file that specifies the main class
- An EJB JAR file that contains the message-driven beans and the entity class, along with the `persistence.xml` file
- An application EAR file that contains the two JAR files along with an `application.xml` file

To deploy the application and run the client, use the following command:

```
ant run
```

This command starts the database server if it is not already running, then deploys and runs the application.

Ignore the message that states that the application is deployed at a URL.

The output in the terminal window looks something like this (preceded by application client container output):

```
running application client container.  
PUBLISHER: Setting hire ID to 25, name Gertrude Bourbon, position Senior Programmer  
PUBLISHER: Setting hire ID to 26, name Jack Verdon, position Manager  
PUBLISHER: Setting hire ID to 27, name Fred Tudor, position Manager  
PUBLISHER: Setting hire ID to 28, name Fred Martin, position Programmer  
PUBLISHER: Setting hire ID to 29, name Mary Stuart, position Manager  
Waiting for 5 message(s)  
New hire event processed:  
  Employee ID: 25  
  Name: Gertrude Bourbon  
  Equipment: Laptop  
  Office number: 183  
Waiting for 4 message(s)  
New hire event processed:
```

```
Employee ID: 26
Name: Jack Verdon
Equipment: Pager
Office number: 20
Waiting for 3 message(s)
New hire event processed:
Employee ID: 27
Name: Fred Tudor
Equipment: Pager
Office number: 51
Waiting for 2 message(s)
New hire event processed:
Employee ID: 28
Name: Fred Martin
Equipment: Desktop System
Office number: 141
Waiting for 1 message(s)
New hire event processed:
Employee ID: 29
Name: Mary Stuart
Equipment: Pager
Office number: 238
```

The output from the message-driven beans and the entity class appears in the server log, wrapped in logging information.

For each employee, the application first creates the entity and then finds it. You may see runtime errors in the server log, and transaction rollbacks may occur. The errors occur if both of the message-driven beans discover at the same time that the entity does not yet exist, so they both try to create it. The first attempt succeeds, but the second fails because the bean already exists. After the rollback, the second message-driven bean tries again and succeeds in finding the entity. Container-managed transactions allow the application to run correctly, in spite of these errors, with no special programming.

Undeploy the application after you finish running the client:

**ant undeploy**

To remove the generated files, use the following command in the `clientmdbentity`, `clientmdbentity-app-client`, and `clientmdbentity-ejb` directories:

**ant clean**

# An Application Example That Consumes Messages from a Remote Server

This section and the following section explain how to write, compile, package, deploy, and run a pair of Java EE modules that run on two Java EE servers and that use the JMS API to interchange messages with each other. It is a common practice to deploy different components of an enterprise application on different systems within a company, and these examples illustrate on a small scale how to do this for an application that uses the JMS API.

However, the two examples work in slightly different ways. In this first example, the deployment information for a message-driven bean specifies the remote server from which it will *consume* messages. In the next example, the same message-driven bean is deployed on two different servers, so it is the client module that specifies the servers (one local, one remote) to which it is *sending* messages.

This first example divides the example in [Chapter 19, “A Message-Driven Bean Example,”](#) into two modules: one containing the application client, and the other containing the message-driven bean.

This section covers the following topics:

- [“Overview of the consumeremote Example Modules” on page 695](#)
- [“Writing the Module Components for the consumeremote Example” on page 696](#)
- [“Creating Resources for the consumeremote Example” on page 697](#)
- [“Using Two Application Servers for the consumeremote Example” on page 697](#)
- [“Building, Deploying, and Running the consumeremoteModules Using NetBeans IDE” on page 697](#)
- [“Building, Deploying, and Running the consumeremote Modules Using Ant” on page 699](#)

You will find the source files for this section in `tut-install/examples/jms/consumeremote/`. Path names in this section are relative to this directory.

## Overview of the consumeremote Example Modules

Except for the fact that it is packaged as two separate modules, this example is very similar to the one in [Chapter 19, “A Message-Driven Bean Example”](#):

- One module contains the application client, which runs on the remote system and sends three messages to a queue.
- The other module contains the message-driven bean, which is deployed on the local server and consumes the messages from the queue on the remote server.

The basic steps of the modules are as follows.

1. The administrator starts two Java EE servers, one on each system.

2. On the local server, the administrator deploys the message-driven bean module, which specifies the remote server where the client is deployed.
3. On the remote server, the administrator places the client JAR file.
4. The client module sends three messages to a queue.
5. The message-driven bean consumes the messages.

Figure 34–5 illustrates the structure of this application. You can see that it is almost identical to Figure 19–1 except that there are two Java EE servers. The queue used is the one on the remote server; the queue must also exist on the local server for resource injection to succeed.

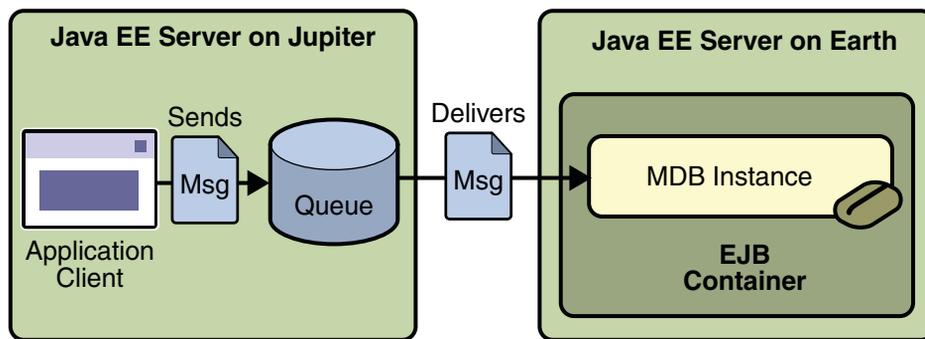


FIGURE 34–5 A Java EE Application That Consumes Messages from a Remote Server

## Writing the Module Components for the consumerremote Example

Writing the components of the modules involves

- Coding the application client
- Coding the message-driven bean

The application client, `jupiterclient/src/java/SimpleClient.java`, is almost identical to the one in “The `simplemessage` Application Client” on page 342.

Similarly, the message-driven bean, `earthmdb/src/java/MessageBean.java`, is almost identical to the one in “The Message-Driven Bean Class” on page 343. The only significant difference is that the activation config properties include one property that specifies the name of the remote system. You need to edit the source file to specify the name of your system.

## Creating Resources for the `consumerremote` Example

The application client can use any connection factory that exists on the remote server; it uses `jms/ConnectionFactory`. Both components use the queue named `jms/Queue`, which you created in “[Creating JMS Administered Objects for the Synchronous Receive Example](#)” on page 650. The message-driven bean does not need a previously created connection factory; the resource adapter creates one for it.

## Using Two Application Servers for the `consumerremote` Example

As in “[Running JMS Clients on Multiple Systems](#)” on page 664, the two servers are named `earth` and `jupiter`.

The Enterprise Server must be running on both systems.

Before you can run the example, you must change the default name of the JMS host on `jupiter`, as described in “[Changing the Default Host Name](#)” on page 667. If you have already performed this task, you do not have to repeat it.

Which system you use to package and deploy the modules and which system you use to run the client depend on your network configuration (which file system you can access remotely). These instructions assume that you can access the file system of `jupiter` from `earth` but cannot access the file system of `earth` from `jupiter`. (You can use the same systems for `jupiter` and `earth` that you used in “[Running JMS Clients on Multiple Systems](#)” on page 664.)

You can package both modules on `earth` and deploy the message-driven bean there. The only action you perform on `jupiter` is running the client module.

## Building, Deploying, and Running the `consumerremote` Modules Using NetBeans IDE

To edit the message-driven bean source file and package the modules using NetBeans IDE, perform these steps:

1. In NetBeans IDE, choose **Open Project** from the **File** menu.
2. In the **Open Project** dialog, navigate to `tut-install/examples/jms/consumerremote/`.
3. Select the `earthhmdb` folder.
4. Select the **Open as Main Project** check box.
5. Click **Open Project**.
6. Edit the `MessageBean.java` file as follows:

- a. Expand the `earthmdb`, `Source Packages`, and `mdb` nodes, then double-click `MessageBean.java`.
- b. Find the following line within the `@MessageBean` annotation:

```
@ActivationConfigProperty(propertyName = "addressList",  
    propertyValue = "remotesystem"),
```

- c. Replace `remotesystem` with the name of your remote system.
7. Right-click the `earthmdb` project and choose `Build`.
8. Choose `Open Project` from the `File` menu.
9. Select the `jupiterclient` folder.
10. Select the `Open as Main Project` check box.
11. Click `Open Project`.
12. Right-click the `jupiterclient` project and choose `Build`.

This target creates a JAR file that contains the client class file and a manifest file.

To deploy the `earthmdb` module and run the application client, perform these steps:

1. Right-click the `earthmdb` project and choose `Deploy`.
2. Copy the `jupiterclient` module to the remote system (`jupiter`). The module is `tut-install/examples/jms/consumerremote/jupiterclient/dist/jupiterclient.jar`. You can use either the file system user interface on your system or the command line.
3. Go to the directory on the remote system where you copied the client JAR file.
4. Use the following command:

```
asadmin deploy --retrieve . jupiterclient.jar
```

This command deploys the client JAR file and retrieves the client stubs in a file named `jupiterclientClient.jar`

5. Use the following command:

```
appclient -client jupiterclientClient.jar
```

On `jupiter`, the output of the `appclient` command looks like this (preceded by application client container output):

```
Sending message: This is message 1 from jupiter  
Sending message: This is message 2 from jupiter  
Sending message: This is message 3 from jupiter
```

On `earth`, the output in the server log looks something like this (preceded by logging information):

```
MESSAGE BEAN: Message received: This is message 1 from jupiter
MESSAGE BEAN: Message received: This is message 2 from jupiter
MESSAGE BEAN: Message received: This is message 3 from jupiter
```

Undeploy the message-driven bean after you finish running the client. To undeploy the `earthmdb` module, perform these steps on earth:

1. Click the Services tab.
2. Expand the Servers node.
3. Expand the GlassFish v3 Domain node.
4. Expand or refresh the Applications node.
5. Right-click `earthmdb` and choose Undeploy.

To undeploy the `jupiterclient` module, use the following command on jupiter:

```
asadmin undeploy jupiterclient
```

To remove the generated files, follow these steps:

1. On earth, right-click the `earthmdb` project and choose Clean.
2. Right-click the `jupiterclient` project and choose Clean.

You can also delete the `jupiterclient.jar` file from the remote filesystem, along with the `jupiterclientClient.jar` file and the directory `jupiterclientClient` that were created when you retrieved the client stubs.

## Building, Deploying, and Running the consumerremote Modules Using Ant

To edit the message-driven bean source file and package and deploy the `earthmdb` module using Ant, perform these steps:

1. Open the file  
`tut-install/examples/jms/consumerremote/earthmdb/src/java/mdb/MessageBean.java` in an editor.
2. Find the following line within the `@MessageBean` annotation:

```
    @ActivationConfigProperty(propertyName = "addressList",
                             propertyValue = "remotesystem"),
```

3. Replace `remotesystem` with the name of your remote system, then save and close the file.
4. Go to the following directory:

```
tut-install/examples/jms/consumerremote/earthmdb/
```

5. Type the following command:

```
ant
```

This command creates a JAR file that contains the bean class file.

6. Type the following command:

```
ant deploy
```

7. Go to the `jupiterclient` directory:

```
cd ../jupiterclient
```

8. Type the following command:

```
ant
```

This target creates a JAR file that contains the client class file and a manifest file.

To copy the `jupiterclient` module to the remote system, perform these steps:

1. Change to the directory `jupiterclient/dist`:

```
cd ../jupiterclient/dist
```

2. Type a command like the following:

```
cp jupiterclient.jar F:/
```

That is, copy the client JAR file to a location on the remote filesystem.

To run the client, perform the following steps:

1. Go to the directory on the remote system (`jupiter`) where you copied the client JAR file.
2. Use the following command:

```
asadmin deploy --retrieve . jupiterclient.jar
```

This command deploys the client JAR file and retrieves the client stubs in a file named `jupiterclientClient.jar`

3. Use the following command:

```
appclient -client jupiterclientClient.jar
```

On `jupiter`, the output of the `appclient` command looks like this:

```
Sending message: This is message 1 from jupiter  
Sending message: This is message 2 from jupiter  
Sending message: This is message 3 from jupiter
```

On `earth`, the output in the server log looks something like this (wrapped in logging information):

```
MESSAGE BEAN: Message received: This is message 1 from jupiter  
MESSAGE BEAN: Message received: This is message 2 from jupiter  
MESSAGE BEAN: Message received: This is message 3 from jupiter
```

Undeploy the message-driven bean after you finish running the client. To undeploy the `earthmdb` module, perform these steps:

1. Change to the directory `earthmdb`.
2. Type the following command:

```
ant undeploy
```

You can also delete the `jupiterclient.jar` file from the remote filesystem, along with the `jupiterclientClient.jar` file and the directory `jupiterclientClient` that were created when you retrieved the client stubs.

To remove the generated files, use the following command in both the `earthmdb` and `jupiterclient` directories:

```
ant clean
```

## An Application Example That Deploys a Message-Driven Bean on Two Servers

This section, like the preceding one, explains how to write, compile, package, deploy, and run a pair of Java EE modules that use the JMS API and run on two Java EE servers. The modules are slightly more complex than the ones in the first example.

The modules use the following components:

- An application client that is deployed on the local server. It uses two connection factories, one ordinary one and one that is configured to communicate with the remote server, to create two publishers and two subscribers and to publish and to consume messages.
- A message-driven bean that is deployed twice: once on the local server, and once on the remote one. It processes the messages and sends replies.

In this section, the term *local server* means the server on which both the application client and the message-driven bean are deployed (`earth` in the preceding example). The term *remote server* means the server on which only the message-driven bean is deployed (`jupiter` in the preceding example).

The section covers the following topics:

- [“Overview of the sendremote Example Modules” on page 702](#)
- [“Writing the Module Components for the sendremote Example” on page 703](#)
- [“Creating Resources for the sendremote Example” on page 704](#)
- [“Using Two Application Servers for the sendremote Example” on page 705](#)
- [“Building, Deploying, and Running the sendremote Modules Using NetBeans IDE” on page 705](#)

- [“Building, Deploying, and Running the sendremote Modules Using Ant” on page 708](#)

You will find the source files for this section in `tut-install/examples/jms/sendremote/`. Path names in this section are relative to this directory.

## Overview of the sendremote Example Modules

This pair of modules is somewhat similar to the modules in [“An Application Example That Consumes Messages from a Remote Server” on page 695](#) in that the only components are a client and a message-driven bean. However, the modules here use these components in more complex ways. One module consists of the application client. The other module contains only the message-driven bean and is deployed twice, once on each server.

The basic steps of the modules are as follows.

1. You start two Java EE servers, one on each system.
2. On the local server (earth), you create two connection factories: one local and one that communicates with the remote server (jupiter). On the remote server, you create a connection factory that has the same name as the one that communicates with the remote server.
3. The application client looks up the two connection factories (the local one and the one that communicates with the remote server) to create two connections, sessions, publishers, and subscribers. The subscribers use a message listener.
4. Each publisher publishes five messages.
5. Each of the local and the remote message-driven beans receives five messages and sends replies.
6. The client’s message listener consumes the replies.

[Figure 34–6](#) illustrates the structure of this application. M1 represents the first message sent using the local connection factory, and RM1 represents the first reply message sent by the local MDB. M2 represents the first message sent using the remote connection factory, and RM2 represents the first reply message sent by the remote MDB.

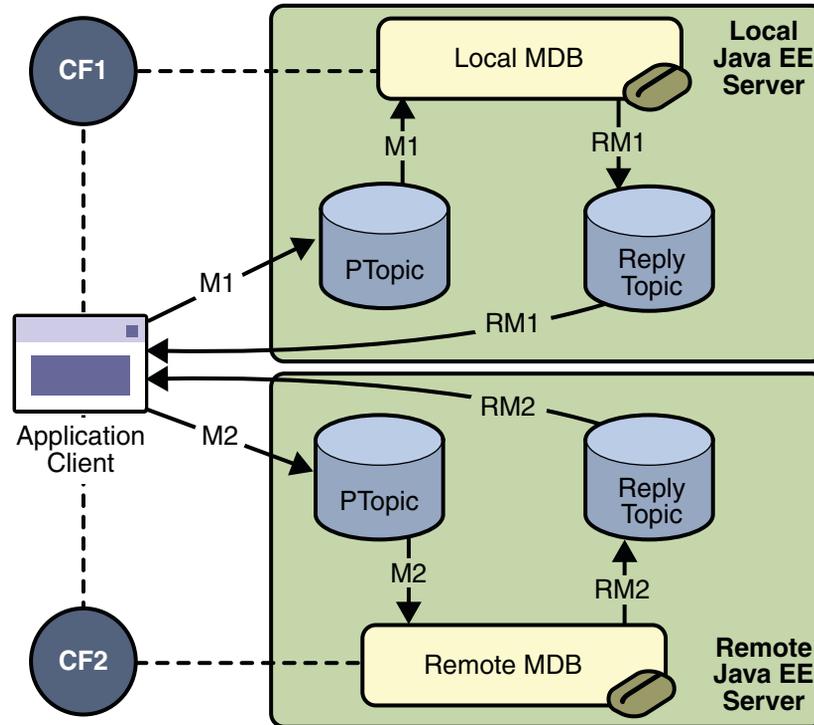


FIGURE 34-6 A Java EE Application That Sends Messages to Two Servers

## Writing the Module Components for the send remote Example

Writing the components of the modules involves two tasks:

- “Coding the Application Client: `MultiAppServerClient.java`” on page 703
- “Coding the Message-Driven Bean: `ReplyMsgBean.java`” on page 704

### Coding the Application Client: `MultiAppServerClient.java`

The application client class, `multiclient/src/java/MultiAppServerClient.java`, does the following.

1. It injects resources for two connection factories and a topic.
2. For each connection factory, it creates a connection, a publisher session, a publisher, a subscriber session, a subscriber, and a temporary topic for replies.
3. Each subscriber sets its message listener, `ReplyListener`, and starts the connection.
4. Each publisher publishes five messages and creates a list of the messages the listener should expect.

5. When each reply arrives, the message listener displays its contents and removes it from the list of expected messages.
6. When all the messages have arrived, the client exits.

## Coding the Message-Driven Bean: ReplyMsgBean.java

The message-driven bean class, `replybean/src/ReplyMsgBean.java`, does the following:

1. Uses the `@MessageDriven` annotation:

```
@MessageDriven(mappedName="jms/Topic")
```
2. Injects resources for the `MessageDrivenContext` and for a connection factory. It does not need a destination resource because it uses the value of the incoming message's `JMSReplyTo` header as the destination.
3. Uses a `@PostConstruct` callback method to create the connection, and a `@PreDestroy` callback method to close the connection.

The `onMessage` method of the message-driven bean class does the following:

1. Casts the incoming message to a `TextMessage` and displays the text
2. Creates a connection, a session, and a publisher for the reply message
3. Publishes the message to the reply topic
4. Closes the connection

On both servers, the bean will consume messages from the topic `jms/Topic`.

## Creating Resources for the `send remote` Example

This example uses the connection factory named `jms/ConnectionFactory` and the topic named `jms/Topic`. These objects must exist on both the local and the remote servers.

This example uses an additional connection factory, `jms/JupiterConnectionFactory`, which communicates with the remote system; you created it in [“Creating Administered Objects for Multiple Systems” on page 665](#). This connection factory must exist on the local server.

The `build.xml` file for the `multiclient` module contains targets that you can use to create these resources if you deleted them previously.

## Using Two Application Servers for the `sendremote` Example

If you are using NetBeans IDE, you need to add the remote server in order to deploy the message-driven bean there. To do so, perform these steps:

1. In NetBeans IDE, click the Runtime tab.
2. Right-click the Servers node and choose Add Server. In the Add Server Instance dialog, perform these steps:
  - a. Select GlassFish v3 (the default) from the Server list.
  - b. In the Name field, specify a name slightly different from that of the local server, such as `GlassFish v3 (2)`.
  - c. Click Next.
  - d. For the Server Location, you can either browse to the location of the Enterprise Server on the remote system or, if that location is not visible from the local system, use the default location on the local system.
  - e. Click Next.
  - f. Select the Register Remote Domain radio button.
  - g. In the Host Name field, type the name of the remote system.
  - h. Click Finish.

Before you can run the example, you must change the default name of the JMS host on `jupiter`, as described in [“Changing the Default Host Name” on page 667](#). If you have already performed this task, you do not have to repeat it.

## Building, Deploying, and Running the `sendremote` Modules Using NetBeans IDE

To package the modules using NetBeans IDE, perform these steps:

1. In NetBeans IDE, choose Open Project from the File menu.
2. In the Open Project dialog, navigate to `tut-install/examples/jms/sendremote/`.
3. Select the `replybean` folder.
4. Select the Open as Main Project check box.
5. Click Open Project.
6. Right-click the `replybean` project and choose Build.  
This command creates a JAR file that contains the bean class file.
7. Choose Open Project from the File menu.

8. Select the `multiclient` folder.
9. Select the Open as Main Project check box.
10. Click Open Project.
11. Right-click the `multiclient` project and choose Build.

This command creates a JAR file that contains the client class file and a manifest file.

To deploy the `multiclient` module on the local server, perform these steps:

1. Right-click the `multiclient` project and choose Properties.
2. Select Run from the Categories tree.
3. From the Server list, select GlassFish v3 Domain (the local server).
4. Click OK.
5. Right-click the `multiclient` project and choose Deploy.

To deploy the `replybean` module on the local and remote servers, perform these steps:

1. Right-click the `replybean` project and choose Properties.
2. Select Run from the Categories tree.
3. From the Server list, select GlassFish v3 Domain (the local server).
4. Click OK.
5. Right-click the `replybean` project and choose Deploy.
6. Right-click the `replybean` project again and choose Properties.
7. Select Run from the Categories tree.
8. From the Server list, select GlassFish v3 (2) (the remote server).
9. Click OK.
10. Right-click the `replybean` project and choose Deploy.

You can use the Services tab to verify that `multiclient` is deployed as an App Client Module on the local server and that `replybean` is deployed as an EJB Module on both servers.

To run the application client, right-click the `multiclient` project and choose Run Project.

This command returns a JAR file named `multiclientClient.jar` and then executes it.

On the local system, the output of the `appclient` command looks something like this:

```
running application client container.  
Sent message: text: id=1 to local app server  
Sent message: text: id=2 to remote app server  
ReplyListener: Received message: id=1, text=ReplyMsgBean processed message: text: id=1 to local  
app server  
Sent message: text: id=3 to local app server  
ReplyListener: Received message: id=3, text=ReplyMsgBean processed message: text: id=3 to local  
app server  
ReplyListener: Received message: id=2, text=ReplyMsgBean processed message: text: id=2 to remote  
app server
```

```
Sent message: text: id=4 to remote app server
ReplyListener: Received message: id=4, text=ReplyMsgBean processed message: text: id=4 to remote
app server
Sent message: text: id=5 to local app server
ReplyListener: Received message: id=5, text=ReplyMsgBean processed message: text: id=5 to local
app server
Sent message: text: id=6 to remote app server
ReplyListener: Received message: id=6, text=ReplyMsgBean processed message: text: id=6 to remote
app server
Sent message: text: id=7 to local app server
ReplyListener: Received message: id=7, text=ReplyMsgBean processed message: text: id=7 to local
app server
Sent message: text: id=8 to remote app server
ReplyListener: Received message: id=8, text=ReplyMsgBean processed message: text: id=8 to remote
app server
Sent message: text: id=9 to local app server
ReplyListener: Received message: id=9, text=ReplyMsgBean processed message: text: id=9 to local
app server
Sent message: text: id=10 to remote app server
ReplyListener: Received message: id=10, text=ReplyMsgBean processed message: text: id=10 to remote
app server
Waiting for 0 message(s) from local app server
Waiting for 0 message(s) from remote app server
Finished
Closing connection 1
Closing connection 2
```

On the local system, where the message-driven bean receives the odd-numbered messages, the output in the server log looks like this (wrapped in logging information):

```
ReplyMsgBean: Received message: text: id=1 to local app server
ReplyMsgBean: Received message: text: id=3 to local app server
ReplyMsgBean: Received message: text: id=5 to local app server
ReplyMsgBean: Received message: text: id=7 to local app server
ReplyMsgBean: Received message: text: id=9 to local app server
```

On the remote system, where the bean receives the even-numbered messages, the output in the server log looks like this (wrapped in logging information):

```
ReplyMsgBean: Received message: text: id=2 to remote app server
ReplyMsgBean: Received message: text: id=4 to remote app server
ReplyMsgBean: Received message: text: id=6 to remote app server
ReplyMsgBean: Received message: text: id=8 to remote app server
ReplyMsgBean: Received message: text: id=10 to remote app server
```

Undeploy the modules after you finish running the client. To undeploy the modules, perform these steps:

1. Click the Services tab.
2. Expand the Servers node.
3. Expand the GlassFish v3 Domain node (the local system).
4. Expand or refresh the Applications node.
5. Right-click `replybean` and choose Undeploy.
6. Right-click `multiclient` and choose Undeploy.
7. Expand the GlassFish v3 (2) node (the remote system).
8. Expand or refresh the Applications node.
9. Right-click `replybean` and choose Undeploy.

To remove the generated files, follow these steps:

1. Right-click the `replybean` project and choose Clean.
2. Right-click the `multiclient` project and choose Clean.

## Building, Deploying, and Running the `sendremote` Modules Using Ant

To package the modules, perform these steps:

1. Go to the following directory:

```
tut-install/examples/jms/sendremote/multiclient/
```

2. Type the following command:

```
ant
```

This command creates a JAR file that contains the client class file and a manifest file.

3. Change to the directory `replybean`:

```
cd ../replybean
```

4. Type the following command:

```
ant
```

This command creates a JAR file that contains the bean class file.

To deploy the `replybean` module on the local and remote servers, perform the following steps:

1. Verify that you are still in the directory `replybean`.
2. Type the following command:

```
ant deploy
```

Ignore the message that states that the application is deployed at a URL.

3. Type the following command:

```
ant deploy-remote -Dsys=remote-system-name
```

Replace *remote-system-name* with the actual name of the remote system.

To deploy and run the client, perform these steps:

1. Change to the directory `multiclient`:

```
cd ../multiclient
```

2. Type the following command:

```
ant run
```

On the local system, the output looks something like this:

```
running application client container.  
Sent message: text: id=1 to local app server  
Sent message: text: id=2 to remote app server  
ReplyListener: Received message: id=1, text=ReplyMsgBean processed message: text: id=1 to local  
app server  
Sent message: text: id=3 to local app server  
ReplyListener: Received message: id=3, text=ReplyMsgBean processed message: text: id=3 to local  
app server  
ReplyListener: Received message: id=2, text=ReplyMsgBean processed message: text: id=2 to remote  
app server  
Sent message: text: id=4 to remote app server  
ReplyListener: Received message: id=4, text=ReplyMsgBean processed message: text: id=4 to remote  
app server  
Sent message: text: id=5 to local app server  
ReplyListener: Received message: id=5, text=ReplyMsgBean processed message: text: id=5 to local  
app server  
Sent message: text: id=6 to remote app server  
ReplyListener: Received message: id=6, text=ReplyMsgBean processed message: text: id=6 to remote  
app server  
Sent message: text: id=7 to local app server  
ReplyListener: Received message: id=7, text=ReplyMsgBean processed message: text: id=7 to local  
app server  
Sent message: text: id=8 to remote app server  
ReplyListener: Received message: id=8, text=ReplyMsgBean processed message: text: id=8 to remote  
app server  
Sent message: text: id=9 to local app server  
ReplyListener: Received message: id=9, text=ReplyMsgBean processed message: text: id=9 to local  
app server  
Sent message: text: id=10 to remote app server  
ReplyListener: Received message: id=10, text=ReplyMsgBean processed message: text: id=10 to remote  
app server  
Waiting for 0 message(s) from local app server  
Waiting for 0 message(s) from remote app server
```

Finished

Closing connection 1

Closing connection 2

On the local system, where the message-driven bean receives the odd-numbered messages, the output in the server log looks like this (wrapped in logging information):

```
ReplyMsgBean: Received message: text: id=1 to local app server
ReplyMsgBean: Received message: text: id=3 to local app server
ReplyMsgBean: Received message: text: id=5 to local app server
ReplyMsgBean: Received message: text: id=7 to local app server
ReplyMsgBean: Received message: text: id=9 to local app server
```

On the remote system, where the bean receives the even-numbered messages, the output in the server log looks like this (wrapped in logging information):

```
ReplyMsgBean: Received message: text: id=2 to remote app server
ReplyMsgBean: Received message: text: id=4 to remote app server
ReplyMsgBean: Received message: text: id=6 to remote app server
ReplyMsgBean: Received message: text: id=8 to remote app server
ReplyMsgBean: Received message: text: id=10 to remote app server
```

Undeploy the modules after you finish running the client. To undeploy the `multiclient` module, perform these steps:

1. Verify that you are still in the directory `multiclient`.
2. Type the following command:

```
ant undeploy
```

To undeploy the `replybean` module, perform these steps:

1. Change to the directory `replybean`:

```
cd ../replybean
```

2. Type the following command:

```
ant undeploy
```

3. Type the following command:

```
ant undeploy-remote -Dsys=remote-system-name
```

Replace *remote-system-name* with the actual name of the remote system.

To remove the generated files, use the following command in both the `replybean` and `multiclient` directories:

```
ant clean
```

# Index

---

## Numbers and Symbols

@ApplicationScoped annotation, 356-358  
@Consumes, 257  
@ConversationScoped annotation, 356-358  
@DELETE, 248, 254  
@Dependent annotation, 356-358  
@DiscriminatorColumn annotation, 393-394  
@DiscriminatorValue annotation, 393-394  
@Embeddable annotation, 390-391  
@EmbeddedId annotation, 385  
@Entity annotation, 381  
@GET, 248, 254  
@Id annotation, 385  
@IdClass annotation, 385  
@Inject annotation, 356  
@Local annotation, 289, 308  
@ManyToMany annotation, 387, 388  
@ManyToOne annotation, 387  
@MessageDriven annotation, 682-683  
@Named annotation, 358  
@NamedQuery annotation, 424  
@OneToMany annotation, 387, 388, 389  
@OneToOne annotation, 387, 388, 389  
@Path, 248, 251  
@PathParam, 259  
@PersistenceContext annotation, 396  
@PersistenceUnit annotation, 397  
@POST, 248, 254  
@PostActivate annotation, 309, 310-311  
@PostConstruct annotation, 297-300  
@PostConstruct annotation, 309, 310-311

@PostConstruct method, session beans using  
JMS, 682  
@PreDestroy annotation, 297-300  
@PreDestroy annotation, 309, 310-311  
@PreDestroy method, session beans using JMS, 682  
@PrePassivate annotation, 309, 310-311  
@Produces, 257  
@Produces annotation, 360  
@PUT, 248, 254  
@Qualifier annotation, 355-356  
@QueryParam, 259  
@Remote annotation, 289, 308  
@Remove annotation, 297  
@Remove annotation, 309, 312  
@RequestScoped annotation, 356-358  
@Resource annotation, 595-598  
JMS resources, 342, 620  
@SessionScoped annotation, 356-358  
@Stateful annotation, 308  
@Timeout annotation, 331  
@Timeout method, 334  
@Transient annotation, 383  
@WebMethod annotation, 311

## A

abstract schemas  
    defined, 423  
    types, 423  
access control, 478  
acknowledge method, 630

- acknowledging messages., *See* message acknowledgment
- action events, 140
  - ActionEvent class
    - and UICommand component, 140
  - addListener attribute, 168
    - and backing bean methods, 168
    - and UICommand component, 140
    - referencing methods that handle action events, 169
  - ActionListener class, 166
  - addListener tag, 157, 166
  - referencing methods that handle action events, 169, 185
  - writing a backing-bean method to handle action events, 185
- Admin Console, 58
- administered objects, JMS, 619-621
  - See also* connection factories, destinations
  - definition, 615
- Administration Console, starting, 64
- annotations
  - DeclareRoles, 508-512
  - DenyAll, 509
  - JAX-RS, 250-251
  - Jersey, 248, 250-251
  - PermitAll, 509
  - RolesAllowed, 508
  - security, 483-484, 504
- Ant tool, 60-61
- appClient tool, 58
- applet containers, 43
- applets, 37, 39
- application client containers, 43
- application client examples, JMS, 646-670
- application clients, 38
  - examples, 342
  - securing, 521-522
- application clients, JMS
  - deploying and running, 651-655
  - packaging, 650-651, 656-657
  - running, 657-660
  - running on multiple systems, 664-670
- applications
  - creating
    - JAX-RS, 263-275
  - deploying
    - JAX-RS, 263-275
  - running
    - JAX-RS, 263-275
  - security, 480
- asadmin tool, 58
- asynchronous message consumption, 618
  - See also* message-driven beans
  - JMS client example, 655-660
- attributes referencing backing bean methods, 168
  - action attribute, 168
    - and backing bean methods, 168
    - and navigation, 169
  - addListener attribute, 168, 169
  - validator attribute, 168, 169
  - valueChangeListener attribute, 168, 170
- audit modules, pluggable, 485
- auditing, 478
- auth-constraint, 534-535
- auth-constraint element, 492
- authenticate method, 549-551
- authenticating users, 537-544
- authentication, 478, 486, 494
  - basic, 538
    - example, 562-565
  - certificate-based mutual, 542
  - digest, 544
  - form-based
    - example, 566-571
  - mutual, 542-543
  - web resources
    - form-based, 539
    - HTTP basic, 538, 566-571
- authentication mechanism, EJB, 512
- authorization, 478, 486
- authorization constraint, 492, 534-535
- authorization providers, pluggable, 485
- AUTO\_ACKNOWLEDGE mode, 630

**B**

backing bean methods, 168, 183  
 attributes referencing  
   *See* attributes referencing backing bean methods  
 referencing  
   *See* referencing backing bean methods  
 writing  
   *See* writing backing bean methods

backing bean properties, 160, 171, 173  
 bound to component instances, 181-182  
 properties for `UISelectItems` composed of  
   `SelectItem` instances, 180  
`UIData` properties, 176-177  
`UIInput` and `UIOutput` properties, 175  
`UISelectBoolean` properties, 177  
`UISelectItems` properties, 179-180  
`UISelectMany` properties, 177-178  
`UISelectOne` properties, 178-179

backing beans, 171-174  
 method binding  
   *See* method binding  
 methods  
   *See* backing bean methods  
 properties  
   *See* backing bean properties

basic authentication, 538-539  
 EJB, 512

bean-managed transactions, 643  
   *See* transactions, bean-managed

beans, defined for CDI, 353

beans.xml file, 361

BLOBs, *See* persistence, BLOBs

`BufferedReader` class, 218

business logic, 281

business methods, 291  
 client calls, 311  
 exceptions, 312  
 locating, 303  
 requirements, 311  
 transactions, 585, 587, 589, 590

`BytesMessage` interface, 626

**C**

`CallbackHandler` interface, 521

capture-schema tool, 58

certificate authority, 497

certificates, 479  
 digital, 481, 497-499  
   managing, 497  
 server  
   generating, 498-499  
   using for authentication, 491

`CLIENT_ACKNOWLEDGE` mode, 630

client ID, for durable subscriptions, 633

client-side, 262-263

clients  
 authenticating, 541-543  
 securing, 521-522

CLOBs, *See* persistence, CLOBs

collections  
 persistence, 383-385, 469

`commit` method, 587

`commit` method (JMS), 635-637

commits, *See* transactions, commits

Common Client Interface, Connector  
 Architecture, 608-609

component binding, 173, 174  
 binding attribute  
   value expressions, 173

component classes  
`SelectItem` class, 148, 179  
`SelectItemGroup` class, 179

component-managed sign-on, 523

component properties, *See* backing bean properties

component tags  
`panelGrid` tag, 142  
`selectItems` tag, 179

composite components, facelets, 109-112

concurrent access, 581

confidentiality, 494

configuring beans, 192-202

configuring JavaServer Faces applications  
 Application class, 190  
 application configuration resource files, 189-191  
   `commandButton` tag, 140  
   conversion model, 161

- configuring JavaServer Faces applications (*Continued*)
  - configuring beans
    - See configuring beans
  - configuring navigation rules
    - See configuring navigation rules
  - faces-config.xml files, 205
  - including the classes, pages, and other resources, 212
  - javax.faces.application.CONFIG\_FILES context parameter, 190
  - registering messages
    - See registering messages
  - specifying a path to an application configuration resource file, 210-211
  - specifying where UI component state is saved, 211
- configuring navigation rules, 204-207
  - from-action element, 205
  - from-view-id element, 205
  - navigation-case element, 205, 206
  - navigation-rule element, 205
  - to-view-id element, 205
- connection factories, JMS
  - creating, 345-346, 650
  - injecting resources, 342, 620
  - introduction, 620
  - specifying for remote servers, 665-666
- Connection interface, 591
- Connection interface (java.sql), 587
- Connection interface (JMS), 621
- connection pooling, 594
- ConnectionFactory interface (JMS), 620
- connections
  - secure, 494
  - securing, 494-499
- connections, JMS
  - introduction, 621
  - managing in enterprise bean applications, 639
- connectors, *See* Java EE Connector architecture
- constraint
  - authorization, 492
  - security, 492
  - user data, 492
- container-managed sign-on, 523
- container-managed transactions, *See* transactions, container-managed
- containers, 41-43
  - See also applet containers
  - See also application client containers
  - See also EJB containers
  - See also web containers
  - configurable services, 42
  - nonconfigurable services, 42
  - securing, 482-484
  - security, 474-478
  - services, 42
  - trust between, 520
- context roots, 77
- Contexts and Dependency Injection for the Java EE Platform (CDI), 351-361
  - beans, 353
  - configuring applications, 361
  - examples, 363-377
  - Facelets pages, 359
  - injectable objects, 354
  - injecting beans, 356
  - managed beans, 353-354
  - overview, 352
  - qualifiers, 355-356
  - scopes, 356-358
  - setter and getter methods, 358-359
  - unified expression language, 358
- conversion model
  - converter attribute, 160-161
    - text components, 135
  - Converter implementations, 159
  - converter tags
    - See converter tags
  - converterId attribute, 160
  - converters
    - See converters
  - javax.faces.convert package, 159
- Converter implementation classes
  - BigDecimalConverter class, 159
  - BigIntegerConverter class, 159
  - BooleanConverter class, 159
  - ByteConverter class, 159
  - CharacterConverter class, 159

Converter implementation classes (*Continued*)

- DateTimeConverter, 160
- DateTimeConverter class, 159, 161
- DoubleConverter class, 159
- EnumConverter class, 160
- FloatConverter class, 160
- IntegerConverter class, 160
- LongConverter class, 160
- NumberConverter class, 160, 161, 163
- ShortConverter class, 160

## converter tags

- convertDateTime tag, 161
- convertDateTime tag attributes, 162-163
- converter tag, 161
- convertNumber tag, 161, 163
- convertNumber tag attributes, 164
- parseLocale attribute, 162

## converters, 90

## cookie parameters, 259

## core tags, convertNumber tag, 163

## createBrowser method, 660

## createTimer method, 331

## creating applications, JAX-RS, 263-275

## credential, 489

## Criteria API, 457-469

- creating queries, 461-462
- expressions, 464-465, 465-467
- query execution, 468-469
- query results, 464-467, 467-468

## cryptography, public key, 497

## custom UI components, specifying where state is saved, 211

## custom validators

- validate method, 185
- Validator implementation
  - backing bean methods, 183

**D**

## data, encryption, 541

## data integrity, 478, 581, 582

## data sources, 594

## databases

- See also* persistence

databases (*Continued*)

- clients, 281
- connections, 312, 589
- data recovery, 581
- EIS tier, 35
- message-driven beans and, 285
- multiple, 588, 590
- transactions
  - See* transactions

## DataSource interface, 594

## debugging, Java EE applications, 66-67

## declarative security, 474, 504, 529

## DeclareRoles annotation, 508-512

## delivery modes, JMS, 631

- JMSDeliveryMode message header field, 625

## DeliveryMode interface, 631

## DenyAll annotation, 509

## Dependency Injection for Java (JSR-330), 351-361

## deployer roles, 49

## deploying, without NetBeans IDE, 266-267

## deploying applications, JAX-RS, 263-275

## deployment, 303-304

## deployment descriptor, 530-533

- auth-constraint element, 492
- security-constraint element, 492
- specifying SSL, 495
- transport-guarantee element, 492
- user-data-constraint element, 492
- web applications, 530-548
- web-resource-collection element, 492

## deployment descriptors, 474, 482-483, 504, 529

- ejb-jar.xml file, 483

- portable, 46

- runtime, 46

- security-role-mapping element, 493

- security-role-ref element, 554

- web application, 73, 75, 207

- runtime, 75

- web services, 483

- web.xml file, 483

## Destination interface, 620-621

## destinations, JMS

- See also* queues, temporary destinations, topics
- creating, 345-346, 650

- destinations, JMS (*Continued*)
  - injecting resources, 342, 620
  - introduction, 620-621
  - JMSDestination message header field, 625
  - temporary, 632-633, 688-689, 703-704
- destroy method, 228
- development roles, 47-49
  - application assemblers, 49
  - application client developers, 48
  - application deployers and administrators, 49
  - enterprise bean developers, 48
  - Java EE product providers, 48
  - tool providers, 48
  - web component developers, 48
- digest authentication, 544
- digital signature, 497
- DNS, 55
- doFilter method, 220, 221, 224
- doGet method, 217
- domains, 63
- doPost method, 217
- downloading, Enterprise Server, 60
- DUPS\_OK\_ACKNOWLEDGE mode, 630
- durable subscriptions, JMS, 633-635
  - examples, 672-674, 680-685
- E**
- EAR files, 46
- EIS tier, 41
  - security, 523-526
- EJB, security, 504-521
- EJB containers, 43
  - container-managed transactions, 582
  - message-driven beans, 640-641
  - onMessage method, invoking, 344-345
  - services, 281, 504-521
- EJB JAR files, 294
  - portability, 294
- EJBContext interface, 586, 587, 589
- embeddable classes, *See* persistence: embeddable classes
- end-to-end security, 481-482
- enterprise bean applications
  - JMS examples, 680-685, 686-694
  - enterprise beans, 40, 50
    - See also* Java EE components
    - accessing, 286
    - business methods
      - See* business methods
    - classes, 294
    - compiling, 303-304
    - contents, 294-296
    - defined, 281
    - deployment, 294
    - distribution, 288
    - exceptions, 339
    - getCallerPrincipal method, 515-516
    - implementor of business logic, 40
    - interfaces, 286-293, 294
    - isCallerInRole method, 515-516
    - life cycles, 297-300
    - local access, 289-291
    - message-driven beans.
      - See* message-driven beans
    - packaging, 303-304
    - performance, 289
    - persistence
      - See* persistence
    - programmatic security, 515-516
    - protecting, 504-521
    - remote access, 291-292
    - securing, 503-526
    - session beans
      - See* session beans
    - timer service, 328-339
    - types, 282
    - web services, 282, 292-293, 324-327
- Enterprise Information Systems, *See* EIS tier
- Enterprise Server
  - adding users to, 490-491
  - downloading, 60
  - enabling debugging, 67
  - installation tips, 60
  - securing, 484-485
  - server logs, 66-67
  - SSL connectors, 494
  - starting, 63
  - stopping, 63

Enterprise Server (*Continued*)

- tools, 57
- entities
  - abstract, 391
  - abstract schema names, 426
  - cascading operations, 389
    - orphans, 389-390
  - collections, 439-440
  - entity manager, 396-400
  - finding, 398, 412
  - inheritance, 391-395, 416-417
  - life cycle, 398-400
  - managing, 395-401
  - overview, 381-395
  - persistent fields, 382-385
  - persistent properties, 382-385
  - persisting, 398-399
  - primary keys, 385-387
  - querying, 401-402
  - relationships, 412
  - removing, 399, 413
  - requirements, 381-382
  - superclasses, 392
  - synchronizing, 400
- entity providers, 255
- entity relationships
  - bidirectional, 388
  - many-to-many, 387
  - many-to-one, 387
  - multiplicity, 387
  - one-to-many, 387
  - one-to-one, 387
  - query language, 388
  - unidirectional, 388
- equals method, 386
- event and listener model
  - ActionEvent class, 136, 138
  - event listeners
    - JavaServer Faces UI, 90
  - Listener class, 183
  - value-change events
    - See value-change events
  - ValueChangeEvent class, 170
- examples, 59-67

examples (*Continued*)

- building, 65
- CDI, 363-377
- classpath, 304
- directory structure, 65
- JMS
  - asynchronous message consumption, 655-660
  - browsing messages in a queue, 660-664
  - durable subscriptions, 672-674
  - enterprise bean examples, 680-685, 686-694
  - Java EE examples, 695-701, 701-710
  - local transactions, 674-679
  - message acknowledgment, 670-672
  - synchronous message consumption, 646-655
- primary keys, 386
- query language, 426-431
- required software, 59-62
- security, 474-477
  - basic authentication, 562-565
  - form-based authentication, 566-571
- servlets, 302
- session beans, 302
- simple servlets, 74
- timer service, 335-338
- web clients, 302
- web services, 238, 324
- exceptions
  - business methods, 312
  - enterprise beans, 339
  - JMS, 628
  - mapping to error screens, 85
  - rolling back transactions, 339, 586
  - transactions, 584, 585
- expiration of JMS messages, 632
  - JMSExpiration message header field, 625
- expressions, lvalue expressions, 173

**F**

- facelets, 97-112
  - composite components, 109-112
  - developing simple applications, 99-106
  - resources, 112
  - templating, 107-109

- filter chains, 220, 224
- Filter interface, 220
- filters, 220
  - defining, 220
  - mapping to web components, 222
  - mapping to web resources, 222, 223, 224
  - overriding request methods, 221
  - overriding response methods, 221
  - response wrappers, 221
- foreign keys, 405
- form parameters, 259
- forward method, 225

## G

- garbage collection, 300
- GenericServlet interface, 213
- getCallerPrincipal method, 515-516
- getCallerPrincipal method, 515
- getConnection method, 594
- getRemoteUser, 549-551
- getRemoteUser method, 551
- getRequestDispatcher method, 224
- getRollbackOnly method, 643
- getServletContext method, 226
- getSession method, 226
- getUserPrincipal, 549-551
- getUserPrincipal method, 551
- groups, 488-489
  - managing, 489-491

## H

- hashCode method, 386
- header parameters, 259
- helper classes, 294, 312
- HTTP, 237
  - over SSL, 541
- HTTP methods, 254
- HTTP request URLs, 218
  - query strings, 219
  - request paths, 218
- HTTP requests, 218

- HTTP requests (*Continued*)
  - See also* requests
- HTTP responses, 220
  - See also* responses
  - status codes, 85
    - mapping to error screens, 85
- HTTPS, 480, 496, 497, 535-536
- HttpServletRequest interface, 213
- HttpServletRequest interface, 551
- HttpServletRequest interface, 218
- HttpServletResponse interface, 220
- HttpSession interface, 226

## I

- identification, 478, 486
- include method, 225
- init method, 217
- InitialContext interface, 55
- initializing properties with the managed-property element
  - initializing Array and List properties, 199
  - initializing managed-bean properties, 199-201
  - initializing Map properties, 197-198
  - initializing maps and lists, 201
  - referencing an initialization parameter, 196-197
- injectable objects, 354
- integrity, 494
  - of data, 478
- internationalizing JavaServer Faces applications
  - basename, 203
  - FacesContext.getLocale method, 162
  - loadBundle tag, 158
  - using localized static data and messages, 145-146
- invalidate method, 227
- isCallerInRole method, 515-516
- isCallerInRole method, 516
- isUserInRole, 549-551
- isUserInRole method, 551

## J

- JAAS, 57, 479, 521-522

- 
- JAAS (*Continued*)
    - login modules, 522
  - JACC, 485
  - JAF, 56
  - JAR files
    - See also* EJB JAR files
    - javaee.jar, 304
    - query language, 438
  - JAR signatures, 479
  - Java API for XML Binding, *See* JAXB
  - Java API for XML Processing, *See* JAXP
  - Java API for XML Registries, *See* JAXR
  - Java API for XML Web Services, *See* JAX-WS
  - Java Authentication and Authorization Service, 479
    - See* JAAS
  - Java BluePrints, 65
  - Java Cryptography Extension (JCE), 479
  - Java DB database, 58
    - starting, 64-65
    - stopping, 64
  - Java EE 6 platform, APIs, 50-54
  - Java EE applications, 35
    - debugging, 66-67
    - deploying, 303-304
    - iterative development, 305
    - JMS examples, 695-701, 701-710
    - running on more than one system, 695-701, 701-710
    - tiers, 35
  - Java EE clients, 37-38
    - application clients, 38
      - See also* application clients
    - web clients, 71-88
      - See also* web clients
  - Java EE components, 37
    - types, 37
  - Java EE Connector architecture, 53
  - Java EE modules, 46, 47
    - application client modules, 47
    - EJB modules, 47, 294
    - resource adapter modules, 47
    - web modules
      - See* web modules
  - Java EE platform, 35
  - Java EE platform (*Continued*)
    - JMS and, 613-614
  - Java EE security model, 42
  - Java EE servers, 43
  - Java EE transaction model, 42
  - Java Generic Security Services, 479
  - Java GSS-API, 479
  - Java Message Service (JMS) API
    - message-driven beans.
      - See* message-driven beans
  - Java Naming and Directory Interface, *See* JNDI
  - Java Persistence API query language, *See* query language
  - Java Persistence Criteria API, *See* Criteria API
  - Java Secure Sockets Extension, 479
  - Java Servlet technology, 50
    - See also* servlets
  - Java Transaction API, *See* JTA
  - JavaBeans Activation Framework, *See* JAF
  - JavaBeans components, 38
    - in WAR files, 75
  - JavaMail API, 53
  - JavaServer Faces, 51
  - JavaServer Faces application development roles
    - application architects
      - responsibilities, 189
    - application developers, 176
      - responsibilities, 171
    - page authors
      - ActionListener registration, 166
      - responsibilities, 127
  - JavaServer Faces applications, queuing messages, 186
  - JavaServer Faces core tag library, 127
    - action attribute, 140
    - actionListener tag, 157, 166
    - attribute tag, 157
    - convertDateTime tag, 157, 161
    - convertDateTime tag attributes, 162-163
    - converter tag, 157, 161
    - converterId attribute, 160
    - convertNumber tag, 157, 161, 163
    - convertNumber tag attributes, 164
    - facet tag, 143, 157, 158
    - jsf\_core TLD, 158
    - loadBundle tag, 157, 158

JavaServer Faces core tag library (*Continued*)

- param tag, 139, 157, 158
  - parseLocale attribute, 162
  - selectItem tag, 130, 145, 146, 158
  - selectitem tag, 147, 157
  - selectItems tag, 130, 145, 146, 158
  - selectitems tag, 147, 157
  - type attribute, 165
  - validateDoubleRange tag, 158, 166
  - validateLength tag, 158, 166
  - validateLongRange tag, 158, 166, 167
  - validator tag, 158
  - validator tags
    - See validator tags
  - valueChangeListener tag, 157, 165
- JavaServer Faces standard components, UISelectItems component, 179
- JavaServer Faces standard HTML render kit tag library, 127
- UI component tags
    - See UI component tags
- JavaServer Faces standard UI components
- UIColumn component, 141
  - UICommand component, 140, 166
  - UIData component, 141, 176
  - UIForm component, 133
  - UIGraphic component, 141
  - UIInput component, 135, 136
  - UIOutput component, 133, 135, 136
  - UIPanel component, 141
  - UISelectBoolean component, 177
  - UISelectItem component, 178, 179
  - UISelectItems component, 147, 178
  - UISelectMany component, 147, 157, 177
  - UISelectOne component, 147, 157
  - UISelectOne properties, 178
- JavaServer Faces tag libraries
- JavaServer Faces core tag library, 157
    - See JavaServer Faces core tag library
  - JavaServer Faces standard HTML tag library
    - See JavaServer Faces standard HTML render kit tag library
  - taglib directives, 128
- JavaServer Faces technology, 89-96

JavaServer Faces technology (*Continued*)

- advantages of, 91
  - FacesContext class
    - Validator interface, 185
  - FacesServlet class, 208
  - UI component tags
    - See UI component tags
- JavaServer Pages Standard Tag Library, *See* JSTL
- javax.servlet.http package, 213
- javax.servlet package, 213
- JAX-RS, 247
- APIs, 248
  - described, 233-235
- JAX-WS, 57
- defined, 237
  - described, 233-235
  - service endpoint interfaces, 238
  - specification, 245
- JAXB, 56
- JAXP, 56
- JAXR, 54
- JCE, 479
- JDBC API, 55, 594
- Jersey, 247
- APIs, 248
  - other info sources, 275-276
- JMS
- achieving reliability and performance, 628-637
  - application client examples, 646-670
  - architecture, 615
  - basic concepts, 614-618
  - definition, 612
  - enterprise bean examples, 680-685, 686-694
  - examples, 645-710
  - introduction, 611-614
  - Java EE examples, 341-348, 695-701, 701-710
  - Java EE platform, 613-614, 638-644
  - messaging domains, 616-617
  - programming model, 618-628
- JMS API, 52
- JMS API, *See* Java Message Service (JMS) API
- JMSCorrelationID message header field, 626
- JMSDeliveryMode message header field, 625
- JMSDestination message header field, 625

- 
- JMSException class, 628
  - JMSExpiration message header field, 625
  - JMSMessageID message header field, 626
  - JMSPriority message header field, 625
  - JMSRedelivered message header field, 626
  - JMSReplyTo message header field, 626
  - JMSTimestamp message header field, 626
  - JMSType message header field, 626
  - JNDI, 55, 593
    - data source naming subcontexts, 55
    - enterprise bean naming subcontexts, 55
    - environment naming contexts, 55
    - jms naming subcontext, 620
    - namespace for JMS administered objects, 619-621
    - naming and directory services, 55
    - naming contexts, 55
    - naming environments, 55
    - naming subcontexts, 55
  - JSR-299, *See* Contexts and Dependency Injection for the Java EE Platform (CDI)
  - JSR-311, 247
  - JSSE, 479
  - JSTL, 51
  - JTA, 52
    - See also* transactions, JTA
  - JTS API, 588
- K**
- Kerberos, 479
  - Kerberos tickets, 479
  - key pairs, 497
  - keystores, 479, 497-499
    - managing, 497
  - keytool utility, 497
- L**
- LDAP, 55
  - listener classes, 214
    - defining, 214
  - listener interfaces, 214
  - listeners
    - HTTP, 484
    - IIOP, 484
  - local interfaces, defined, 289
  - local transactions, JMS, 635-637
  - login, configuring, 537-544
  - login configuration, 537-544
  - login method, 549-551
  - login modules, 521-522
  - logout method, 549-551
- M**
- managed bean creation facility, 192
    - initializing properties with managed-property elements, 195-201
  - managed bean declarations
    - key-class element, 198
    - list-entries element, 195
    - managed-bean element, 193-195, 200
    - managed-bean-name element, 194
    - managed-bean-scope element, 194
    - managed-property element, 195-201
    - map-entries element, 196, 197
    - map-entry element, 197
    - null-value elements, 196
    - value element, 196
  - managed beans, defined for CDI, 353-354
  - Managed Beans specification, 351-361
  - MapMessage interface, 626
  - matrix parameters, 259
  - message acknowledgment, JMS
    - bean-managed transactions, 644
    - introduction, 629-630
    - message-driven beans, 640
  - message bodies, JMS, 626-627
  - message consumers, JMS, 623-625
  - message consumption, JMS
    - asynchronous, 618, 655-660
    - introduction, 618
    - synchronous, 618, 646-655
  - message-driven beans, 50, 285-286
    - accessing, 285
    - coding, 343-345, 682-683, 689, 704

message-driven beans (*Continued*)

- defined, 285
  - examples, 341-348, 680-685, 686-694, 695-701, 701-710
  - garbage collection, 300
  - introduction, 640-641
  - onMessage method, 286, 344-345
  - requirements, 343-345
  - transactions, 582, 587, 588
- message headers, JMS, 625-626
- message IDs, JMSMessageID message header field, 626
- Message interface, 627
- message listeners, JMS, 285
- message listeners, JMS
- examples, 656, 688-689, 703-704
  - introduction, 623-624
- message producers, JMS, 622-623
- message properties, JMS, 626
- message security, 485, 529, 555
- message selectors, JMS, introduction, 624-625
- MessageBodyReader, 255
- MessageBodyWriter, 255
- MessageConsumer interface, 623-625
- MessageListener interface, 623-624
- MessageProducer interface, 622-623
- messages
- integrity, 541
  - MessageFormat pattern, 138, 157
  - outputFormat tag, 138
  - param tag, 139
  - parameter substitution tags
    - See JavaServer Faces core tag library
    - param tag, 157
  - queueing messages, 186, 202
  - securing, 481-482
  - security, 485
- messages, JMS
- body formats, 626-627
  - browsing, 627
  - definition, 615
  - delivery modes, 631
  - expiration, 632
  - headers, 625-626
  - introduction, 625-627

messages, JMS (*Continued*)

- persistence, 631
  - priority levels, 631
  - properties, 626
- messaging, definition, 611-612
- messaging domains, JMS, 616-617
- common interfaces, 617
  - point-to-point, 616
  - publish/subscribe, 617
- metadata annotations, 603-604, 604-608
- security, 483-484
- Metamodel API, 457-459
- using, 459-460
- method binding, 136
- method-binding expressions, 136, 206
  - method expressions, 168
- method expressions, 114
- method permissions, 507
- annotations, 508-512

**N**

## navigation model

- action attribute, 168, 169
    - and backing bean methods, 168
    - and UICommand component, 140
  - action methods, 183, 204
  - ActionEvent class, 169
  - configuring navigation rules, 204-207
  - logical outcome, 183, 204
    - commandButton tag, 140
    - referencing backing bean methods, 169
  - navigation rules, 140, 204
  - NavigationHandler class, 184
  - referencing methods that perform navigation, 169, 183
    - writing a backing bean method to perform navigation processing, 183-184
- NDS, 55
- NetBeans IDE, 62
- NIS, 55
- NON\_PERSISTENT delivery mode, 631
- non-repudiation, 478

**O**

ObjectMessage interface, 626  
 objects, administered (JMS), 619-621  
 onMessage method  
   introduction, 623-624  
   message-driven beans, 286, 344-345, 640  
 overview, further topics, 262-263

**P**

package-appclient tool, 58  
 parameters, extracting, 259  
 path, templates, 251  
 path parameters, 259  
 permissions, policy, 485  
 PermitAll annotation, 509  
 persistence  
   BLOBs, 410-411  
   cascade operations, 410  
   CLOBs, 410-411  
   collections, 383-385  
     maps, 384-385  
   configuration, 401  
   context, 396  
   embeddable classes, 390-391  
   entities, 381-395  
   JMS messages, 631  
   many-to-many, 415-416  
   one-to-many, 405  
   one-to-one, 404-405  
   overview, 381-402  
   persistence units, 400-401  
   persistent fields, 382  
   primary keys, 385-387  
     compound, 406-409  
   properties, 383  
   queries, 381-402, 413, 424-425  
     *See also* query language  
     creating, 461-462  
     Criteria, 457-469  
     executing, 468-469  
     expressions, 464-465, 465-467  
     joins, 463  
     parameters, 424-425, 425  
   persistence, queries (*Continued*)  
     results, 464-467, 467-468  
     type-safe, 457-469  
   query language, 388  
   relationships, 403-405  
   scope, 400-401  
   self-referential relationships, 404  
   session beans, 283  
   temporal types, 411  
 persistence units  
   query language, 423, 438  
 PERSISTENT delivery mode, 631  
 persistent entities, JMS example, 686-694  
 pluggable audit modules, 485  
 pluggable authorization providers, 485  
 point-to-point messaging domain, 616  
   *See also* queues  
 policy files, 479  
 prerequisites, 23  
 primary keys, 405  
   compound, 406-409  
   defined, 385-387  
   examples, 386  
 principal, 489  
   default, 520-521  
 PrintWriter class, 219  
 priority levels, for messages, 631  
   JMSPriority message header field, 625  
 producer methods, 360  
 programmatic login, 485  
 programmatic security, 474, 484, 505, 529  
 programming model, JMS, 618-628  
 properties., *See* message properties, JMS  
 providers, JMS, 615  
 proxies, 237  
 public key certificates, 541  
 public key cryptography, 497  
 publish/subscribe messaging domain  
   *See also* topics  
   durable subscriptions, 633-635  
   introduction, 617

**Q**

qualifiers, using, 355-356

Quality of Service (QOS), 478

query language

ABS function, 449

abstract schemas, 423, 426, 439

ALL expression, 447

ANY expression, 447

arithmetic functions, 448-449

ASC keyword, 454

AVG function, 452

BETWEEN expression, 430-431, 444-445

boolean literals, 443

boolean logic, 450

collection member expressions, 439, 446-447

collections, 439-440, 446

compared to SQL, 428, 438, 441

comparison operators, 431, 444

CONCAT function, 448

conditional expressions, 429, 442, 443, 451

constructors, 453-454

COUNT function, 452

DELETE expression, 431

DELETE statement, 426

DESC keyword, 454

DISTINCT keyword, 427

domain of query, 423, 437, 438

duplicate values, 427

enum literals, 443

equality, 451

ESCAPE clause, 445

examples, 426-431

EXISTS expression, 447

FETCH JOIN operator, 440

FROM clause, 426, 437-441

grammar, 432-455

GROUP BY clause, 426, 454-455

HAVING clause, 426, 455

identification variables, 426, 437, 438-441

identifiers, 437-438

IN operator, 440, 445

INNER JOIN operator, 440

input parameters, 429, 443

IS EMPTY expression, 430

query language (*Continued*)

IS FALSE operator, 451

IS NULL expression, 430

IS TRUE operator, 451

JOIN statement, 428, 440-441

LEFT JOIN operator, 440

LEFT OUTER JOIN operator, 440

LENGTH function, 448

LIKE expression, 430, 445-446

literals, 442-443

LOCATE function, 448

LOWER function, 448

MAX function, 452

MEMBER expression, 446-447

MIN function, 452

MOD function, 449

multiple declarations, 438

multiple relationships, 429

named parameters, 427, 443

navigation, 428-429, 429, 439, 442

negation, 451

NOT operator, 451

null values, 446, 450

numeric comparisons, 451

numeric literals, 443

operator precedence, 443-444

operators, 443-444

ORDER BY clause, 426, 454

parameters, 427

parentheses, 443

path expressions, 423, 441-442

positional parameters, 443

range variables, 439

relationship fields, 424

relationships, 423, 428, 429

return types, 452

scope, 423

SELECT clause, 426, 452-454

setNamedParameter method, 427

SIZE function, 449

SQRT function, 449

state fields, 424

string comparison, 451

string functions, 448-449

- query language (*Continued*)
    - string literals, 442
    - subqueries, 447
    - SUBSTRING function, 448
    - SUM function, 453
    - syntax, 432-455
    - TRIM function, 448
    - types, 441-442, 451
    - UPDATE expression, 426, 431
    - UPPER function, 448
    - WHERE clause, 426, 442-452
    - wildcards, 445
  - query parameters, 259
  - Queue interface, 620-621
  - QueueBrowser interface, 627
    - JMS client example, 660-664
  - queues
    - browsing, 627, 660-664
    - creating, 620-621, 650
    - injecting resources, 342
    - introduction, 620-621
    - temporary, 632-633, 688-689
- R**
- realms, 486, 487-488
    - admin-realm, 488
    - certificate, 488
      - adding users, 491
    - configuring, 485
    - file, 488
  - recover method, 630
  - redelivery of messages, 629, 630
    - JMSRedelivered message header field, 626
  - referencing backing bean methods, 168-170
    - for handling action events, 169, 185
    - for handling value-change events, 170
    - for performing navigation, 169, 183
    - for performing validation, 169-170, 185
  - registering messages, 202
    - message-bundle element, 202
  - relationship fields, query language, 424
  - relationships
    - direction, 388
    - relationships (*Continued*)
      - unidirectional, 405
  - reliability, JMS
    - advanced mechanisms, 633-637
    - basic mechanisms, 629-633
    - durable subscriptions, 633-635
    - local transactions, 635-637
    - message acknowledgment, 629-630
    - message expiration, 632
    - message persistence, 631
    - message priority levels, 631
    - temporary destinations, 632-633
  - remote interfaces, defined, 291
  - Remote Method Invocation (RMI), and
    - messaging, 611-612
  - request method designator, 248, 254
  - request/reply mechanism
    - JMSCorrelationID message header field, 626
    - JMSReplyTo message header field, 626
    - temporary destinations and, 632
  - RequestDispatcher interface, 224
  - requests, 218
    - See also* HTTP requests
    - customizing, 221
    - getting information from, 218
  - Required transaction attribute, 644
  - resource adapter, security, 524-525
  - resource adapters, 53, 598-603
  - resource class, 248
  - resource injection, 595-598
  - resource method, 248
  - resources, 593-609
    - See also* data sources
    - facelets, 112
    - JMS, 639
  - ResponseBuilder, 255
  - responses, 219
    - See also* HTTP responses
    - buffering output, 219
    - customizing, 221
    - setting headers, 217
  - RESTful web services, 247
  - roles, 489
    - application, 493

roles (*Continued*)

- declaring, 546-547
  - development
    - See development roles
  - mapping
    - to groups, 547-548
  - mapping to groups, 493
  - mapping to users, 493
  - referencing, 508-512
  - security, 491-493, 507, 508-512, 544-548
- RolesAllowed annotation, 508
- rollback method, 587, 588, 589
- rollback method (JMS), 635-637
- rollbacks, *See* transactions, rollbacks
- run-as identity, 518-520
- RunAs annotation, 518-520
- running applications, JAX-RS, 263-275

**S**

- SAAJ, 56
- SASL, 479
- schema, deployment descriptors, 482-483
- schemagen tool, 58
- scopes, using, 356-358
- secure connections, 494-499
- Secure Socket Layer (SSL), 494-499
- security, 262-263
  - advanced concepts, 573
  - annotations, 483-484, 504
    - web applications, 529
  - application, 480
    - characteristics of, 478
  - application client tier
    - callback handlers, 521-522
  - callback handlers, 521
  - clients, 521-522
  - constraints, 533-536
  - container, 474-478
  - container trust, 520
  - containers, 482-484
  - context
    - enterprise beans, 515-516
  - declarative, 474, 482-483, 504, 529

security (*Continued*)

- default principal, 520-521
  - deploying enterprise beans, 520-521
  - EIS applications, 523-526
    - component-managed sign-on, 523-524
    - container-managed sign-on, 523
  - end-to-end, 481-482
  - enterprise beans, 504-521
  - example, 474-477
  - functions, 477-478
  - groups, 488-489
  - implementation mechanisms, 479-482
  - introduction, 473-501
  - JAAS login modules, 522
  - Java EE
    - mechanisms, 480-482
  - Java SE, 479
  - login forms, 521
  - login modules, 521-522
  - mechanisms, 477-478
  - message, 529, 555
  - message-layer, 481-482
  - method permissions, 507
    - annotations, 508-512
  - policy domain, 489
  - programmatic, 474, 484, 505, 529, 548-554
  - programmatic login, 485
  - propagating identity, 518-520
  - realms, 487-488
  - resource adapter, 524-525
  - role names, 508-512, 546-547
  - roles, 489, 491-493, 507, 544-548
  - run-as identity, 518-520
  - single sign-on, 485
  - specifying run-as identity, 518-520
  - transport-layer, 480-481, 494-499
  - users, 488, 546
  - web applications, 527
    - deployment descriptor, 530-548
    - overview, 528
  - web components, 527
  - web services, 527
- security constraint, 492
- security-constraint element, 492

- security constraints, 533-536
  - multiple, 536
- security domain, 489
- security identity
  - propagating, 518-520
  - specific identity, 519
- security-role-mapping element, 493
- security-role-ref element, 554
- security roles, 491-493, 507, 547-548
- send method, 622-623
- server, authentication, 541
- servers, certificates, 497-499
- servers, Java EE
  - deploying on more than one, 695-701, 701-710
  - running JMS clients on more than one, 664-670
- Servlet interface, 213
- ServletContext interface, 226
- ServletInputStream class, 218
- ServletOutputStream class, 219
- ServletRequest interface, 218
- ServletResponse interface, 219
- servlets, 213
  - binary data
    - reading, 218
    - writing, 219
  - character data
    - reading, 218
    - writing, 219
  - compiling, 303-304
  - examples, 74, 302
  - finalization, 228
  - initialization, 217
    - failure, 217
  - life cycle, 214-215
  - life-cycle events
    - handling, 214
  - packaging, 303-304
  - service methods, 217
    - notifying, 229
    - programming long running, 230
    - tracking service requests, 228
- session beans, 50, 283-285
  - activation, 297
  - clients, 283
  - session beans (*Continued*)
    - databases, 587
    - examples, 302, 680-685
    - passivation, 297
    - requirements, 308
    - stateful, 283, 284
    - stateless, 283, 284
    - transactions, 582, 587, 588
    - web services, 293, 325-326
- Session interface, 621-622
- sessions, 226
  - associating attributes, 226
  - associating with user, 227
  - invalidating, 227
  - notifying objects associated with, 227
- sessions, JMS
  - introduction, 621-622
  - managing in enterprise bean applications, 639
- setRollbackOnly method, 643
- sign-on
  - component-managed, 523
  - container-managed, 523
- Simple Authentication and Security Layer, 479
- single sign-on, 485
- SingleThreadModel interface, 216
- SOAP, 233-235, 237, 245
- SOAP messages, 45
  - securing, 481-482
- SOAP with Attachments API for Java, *See* SAAJ
- SQL, 55, 428, 438, 441
- SQL92, 450
- SSL, 480, 494-499, 535-536, 541
  - connector, 494
  - connectors
    - Enterprise Server, 494
    - specifying, 492
    - tips, 496
    - verifying support, 496
- SSL HTTPS Connector, configuring, 494
- SSO, 485
- standard converters
  - Converter implementation classes, 159
  - converter tags, 157, 158, 161
  - NumberConverter class, 160

standard converters (*Continued*)

using, 159

## standard validators

using, 166

validator tags

*See* validator tags

state fields, query language, 424

StreamMessage interface, 626

subscription names, for durable subscribers, 633

substitution parameters, defining, *See* messages, param tag

synchronous message consumption, 618

JMS client example, 646-655

**T**

templating, facelets, 107-109

temporary JMS destinations, 632-633

examples, 688-689, 703-704

testing, without NetBeans IDE, 266-267

TextMessage interface, 626

Thawte certificate authority, 497

timer service, 328-339

cancelling timers, 334

creating timers, 331-333

examples, 335-338

exceptions, 334

getInfo method, 334

getNextTimeout method, 334

getTimeRemaining method, 334

getting information, 334-335

saving timers, 334

transactions, 335

timestamps, for messages, JMSTimestamp message

header field, 626

Topic interface, 620-621

## topics

creating, 620-621, 650

durable subscriptions, 633-635

introduction, 620-621

temporary, 632-633, 703-704

transactions, 581-591

application-managed, 396-398

attributes, 583-586

transactions (*Continued*)

bean-managed, 587-589, 643

boundaries, 582, 587

business methods

*See* business methods, transactions

commits, 582, 587

container-managed, 582-587, 589, 642

default transaction demarcation, 582

defined, 582

distributed, JMS, 642-644

examples, 674-679

exceptions

*See* exceptions transactions

JDBC, 590

JMS and enterprise bean applications, 639

JTA, 588

local, JMS, 635-637

managers, 585, 588, 590

message-driven beans, 286

*See* message-driven beans, transactions

nested, 582, 588

Required attribute, 644

rollbacks, 582, 586-587, 587, 588

scope, 583

session beans

*See* session beans, transactions

timeouts, 589

timer service, 335

web components, 591

transport guarantee, 535-536

transport-guarantee, 535-536

transport-guarantee element, 492

transport-layer security, 480-481, 494-499

truststores, 497-499

managing, 497

**U**

UI component behavioral interfaces

ActionSource interface

action and ActionListener attributes, 168

UI component classes

SelectedItem class, 147

UISelectBoolean class, 144

---

 UI component classes (*Continued*)

- UISelectItem class, 147

- UISelectItems class, 147

- UISelectMany class, 146

- UISelectOne class, 144

## UI component renderers

- Table renderer, 141

- Text renderer, 135, 136

## UI component tag attributes, 131

- action attribute, 183

- actionListener attribute, 168, 185
  - and backing bean methods, 168
  - and UICommand component, 140

- alt attribute, 141

- basename attribute, 203

- binding attribute, 131, 133

- value expressions, 173

- columns attribute, 143

- converter attribute, 160-161

- text components, 135

- first attribute, 141

- for attribute, 138, 152

- id attribute, 131

- immediate attribute, 131

- redisplay attribute, 137

- rendered attribute, 131, 132

- rows attribute, 141

- style attribute, 131, 132, 141, 152

- styleClass attribute, 131, 132

- validator attribute, 185

- text components, 136

- value attribute, 131, 133

- binding to a backing-bean property, 174

- commandButton tag, 140

- graphicImage tag, 141

- outputFormat tag, 139

- outputLabel tag, 138

- selectItems tag, 147

- valueChangeListener attribute, 136, 170, 186

- var attribute

- graphicImage tag, 141

- registering static localized text, 203

## UI component tags, 131, 174

 UI component tags (*Continued*)

- attributes

- See UI component tag attributes

- column tag, 128

- commandButton tag, 129, 140

- commandLink tag, 129, 140

- dataTable tag, 129, 141, 176

- form tag, 129, 134

- graphicImage tag, 129

- inputHidden tag, 129, 135

- inputSecret tag, 129, 135, 137

- inputText tag, 129, 135

- text components, 135

- text fields, 136

- inputTextarea tag, 129, 135

- message tag, 129, 152

- messages tag, 129, 152

- outputFormat tag, 129, 138, 140

- outputLabel tag, 129, 136, 137

- outputLink tag, 129, 136, 138

- outputMessage tag, 136

- outputText tag, 130, 136, 138, 177

- text fields, 136

- panelGrid tag, 130, 141

- panelGroup tag, 130, 141

- selectBooleanCheckbox tag, 130, 144, 177

- selectManyCheckbox tag, 130, 146, 177

- selectManyListbox tag, 130, 146

- selectManyMenu tag, 130

- selectOneListbox tag, 130, 144

- selectOneMenu tag, 130, 145, 178, 179

- selectOneRadio tag, 130, 144

## UI components

- buttons, 129

- check boxes, 130

- combo boxes, 130

- data grids, 129

- hidden fields, 129

- hyperlinks, 129

- labels, 129, 130

- list boxes, 130

- password fields, 129

- radio buttons, 130

- table columns, 128

UI components (*Continued*)

- tables, 130
- text areas, 129
- text fields, 129
- UnavailableException class, 217
- unified expression language, 113-126, 173
  - deferred evaluation expressions, 114
  - expression examples, 125
  - immediate evaluation expressions, 114
  - literal expressions, 119, 123
  - literals, 118
  - lvalue expressions, 114, 116
  - managed beans, 358
  - method expressions, 120
  - operators, 124
  - reserved words, 125
  - rvalue expressions, 114, 116
  - type conversion during expression evaluation, 120
  - value expressions, 114, 116
- URI path templates, 251
- user data constraint, 492, 535-536
- user-data-constraint, 535-536
- user-data-constraint element, 492
- users, 488, 546
  - adding to Enterprise Server, 490-491
  - managing, 489-491
- UserTransaction interface, 587, 588, 589, 591
  - message-driven beans, 643
- utility classes, 294

**V**

## validation model

- referencing a method that performs validation, 169-170
- validator attribute, 168, 185
  - and backing bean methods, 168
  - referencing backing bean methods, 169
  - text components, 136
- Validator class, 183
- Validator interface, 185-186
- validators
  - See validators

validation model (*Continued*)

- writing a backing bean method to perform validation, 185-186
- Validator implementation classes, 166-167
  - DoubleRangeValidator class, 158, 166
  - LengthValidator class, 158, 166
  - LongRangeValidator class, 158, 166, 167
- validator tags, 158
  - validateDoubleRange tag, 166
  - validateLength tag, 166
  - validateLongRange tag, 166, 167
- validators, 90
  - custom validators, 158
- value binding, 175
  - a component instance to a bean property
    - See component binding
  - acceptable types of component values, 175
  - value attribute
    - binding to a backing-bean property, 174
    - commandButton tag, 140
    - graphicImage tag, 141
    - outputFormat tag, 139
    - outputLabel tag, 138
    - selectItems tag, 147
  - value expressions, 176
- value-change events, 165
  - processValueChangeEvent (ValueChangeEvent) method, 186
  - referencing methods that handle value-change events, 170
  - type attribute, 165
  - ValueChangeEvent class, 165
  - valueChangeListener attribute, 136, 168, 186
  - ValueChangeListener class, 165, 186
  - valueChangeListener tag, 157, 165
  - writing a backing bean method to handle value-change events, 186-187
- value expressions, 173
  - ValueExpression class, 174
- VeriSign certificate authority, 497

**W**

- W3C, 237, 245

- WAR files, JavaBeans components in, 75
  - web applications, 75
    - configuring, 73, 81
    - maintaining state across requests, 226
    - presentation-oriented, 71
      - securing, 527
    - securing, 527
    - security
      - overview, 528
    - service oriented, 71
    - service-oriented
      - securing, 527
    - specifying initialization parameters, 83-84
    - specifying welcome files, 83
  - web beans, *See* Contexts and Dependency Injection for the Java EE Platform (CDI)
  - web clients, 37, 71-88
    - examples, 302
  - web components, 39, 71
    - See also* Java EE components
    - applets bundled with, 39
    - concurrent access to shared resources, 216
    - forwarding to other web components, 225
    - including other web resources, 225
    - invoking other web resources, 224
  - Web components, JMS and, 644
  - web components
    - mapping exceptions to error screens, 85
    - mapping filters to, 222
    - scope objects, 216
    - securing, 527
    - sharing information, 216
    - specifying aliases, 81
    - specifying initialization parameters, 84
    - transactions, 591
    - types, 39
    - utility classes bundled with, 39
    - web context, 226
  - web containers, 43
    - loading and initializing servlets, 214
    - mapping URLs to web components, 81
  - web modules, 47, 75
    - deploying, 77
      - packaged, 78
    - web modules (*Continued*)
      - dynamic reloading, 80
      - undeploying, 80
      - updating, 79
      - viewing deployed, 79
  - web resource collection, 492, 534
  - web-resource-collection, 534
  - web-resource-collection element, 492
  - web resources, 75
    - mapping filters to, 222, 223, 224
    - unprotected, 534
  - web services, 44
    - described, 233
    - EJB.
      - See* enterprise beans, web services
    - endpoint implementation classes, 325
    - examples, 238, 324
    - JAX-RS compared to JAX-WS, 233-235
    - securing, 527
  - web.xml file, 563-564
  - web.xml file, security elements, 530-533
  - work flows, 284
  - writing backing bean methods, 183-187
    - for handling action events, 185
    - for handling value-change events, 186-187
    - for performing navigation, 183-184
    - for performing validation, 185-186
  - writing backing-bean methods, for performing validation, 136
  - WSDL, 45, 233-235, 237, 245
  - wsgen tool, 58, 239
  - wsimport tool, 58
- X**
- xjc tool, 58
  - XML, 44, 237

