


SCWCD

SCE 310-081

EXAM STUDY KIT

SECOND EDITION



JAVA WEB

COMPONENT

DEVELOPER

CERTIFICATION

Hanumant Deshmukh
Jignesh Malavia
Matthew Scarpino

 **MANNING**

Praise for the First Edition

“Written in a very easy-to-read, conversational tone and is an excellent resource for someone who’s familiar with Java but not with Servlets and JSPs or even for someone familiar with them, but who needs to brush up on some of the details for the exam ... The bundled CD is chock-full of excellent resources ... I will definitely use this book as a resource even after the exam.”

— *JavaRanch.com*

“If you want to buy just one book for the SCWCD exam, then this is the book to buy. The book is well-written and should act as a good reference for you.”

— *JavaPrepare.com*

“An excellent study guide highly recommended not only for SCWCD exam takers, but for anyone intending to put their exam credentials to good use ... a solid reference for dedicated programmers.”

— *Internet Bookwatch*

Five stars! “Well written and well organized by folks who create testing software and mock exams. The Java source code examples are concise and illustrate the point well ... The Bottom Line: A terrific study guide for the new Sun Certified Web Component Developer Certification (SCWCD).”

— *Focus on Java at About.com*

“Certainly recommended for the web component developer examination ... extremely well organized and goes through each and every objective explaining the concepts in a lucid manner ... this book avoids the hassles of going through any API’s or specs because of its thorough coverage.

“... the discussion is thorough and not intimidating to a novice and even a beginner of web programming can digest the material easily. Overall I strongly recommend this book as a study guide for the examination and also as a general reference for JSP technology.”

— *Austin JUG*

“Like other Manning titles I’ve reviewed, this title is very dense with little fluff ... indispensable if you are studying to earn this certification or just getting your feet wet in the web tier of Java technology ... the perfect reference for the experienced developer who needs to learn the salient features of JSP/servlet technology quickly and without a lot of introductory ‘this is web programming’ fluff ... it is a very thorough Servlet/JSP/Tag Library reference and developer guide.”

— *DiverseBooks.com*

“!!!! Exceptional!”

— *Today’s Books*

SCWCD
Exam Study Kit
Second Edition

JAVA WEB COMPONENT DEVELOPER CERTIFICATION

MATTHEW SCARPINO (Second Edition author)
HANUMANT DESHMUKH
JIGNESH MALAVIA
with Jacquelyn Carter



MANNING

Greenwich
(74° w. long.)

For online information and ordering of this and other Manning books, please go to www.manning.com. The publisher offers discounts on this book when ordered in quantity. For more information, please contact:

Special Sales Department
Manning Publications Co.
209 Bruce Park Avenue
Greenwich, CT 06830

Fax: (203) 661-9018
email: orders@manning.com


©2005 by Manning Publications Co. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in the book, and Manning Publications was aware of a trademark claim, the designations have been printed in initial caps or all caps.

The authors and publisher have taken care in the preparation of this book, but make no express or implied warranty of any kind and assume no responsibility for errors or omissions. The authors and publisher assume no liability for losses or damages in connection with or resulting from the use of information or programs in the book and the accompanying downloads.

☉ Recognizing the importance of preserving what has been written, it is Manning's policy to have the books we publish printed on acid-free paper, and we exert our best efforts to that end.

 Manning Publications Co.
209 Bruce Park Avenue
Greenwich, CT 06830

Copyeditor: Liz Welch
Typesetter: D. Dalinnik
Cover designer: Leslie Haimes

ISBN 1-932394-38-9

Printed in the United States of America

1 2 3 4 5 6 7 8 9 10 – VHG – 09 08 07 06 05

brief contents

Part 1 Getting started 1

- 1 Understanding Java servlets 3*
- 2 Understanding JavaServer Pages 14*
- 3 Web application and HTTP basics 21*

Part 2 Servlets 29

- 4 The servlet model 31*
- 5 Structure and deployment 67*
- 6 The servlet container model 83*
- 7 Using filters 97*
- 8 Session management 119*
- 9 Developing secure web applications 139*

Part 3 JavaServer Pages and design patterns 163

- 10 The JSP technology model—the basics 165*
- 11 The JSP technology model—advanced topics 188*
- 12 Reusable web components 219*

13	<i>Creating JSPs with the Expression Language (EL)</i>	236
14	<i>Using JavaBeans</i>	251
15	<i>Using custom tags</i>	285
16	<i>Developing “Classic” custom tag libraries</i>	309
17	<i>Developing “Simple” custom tag libraries</i>	352
18	<i>Design patterns</i>	376

Appendices

A	<i>Installing Tomcat 5.0.25</i>	403
B	<i>A sample web.xml file</i>	408
C	<i>Review Q & A</i>	412
D	<i>Exam Quick Prep</i>	475

contents

<i>preface to the second edition</i>	<i>xv</i>
<i>preface to the first edition</i>	<i>xvii</i>
<i>acknowledgments</i>	<i>xviii</i>
<i>about the Sun certification exams</i>	<i>xix</i>
<i>about this book</i>	<i>xxii</i>
<i>about the authors</i>	<i>xxv</i>
<i>about the cover illustration</i>	<i>xxvi</i>

Part 1 Getting started 1

<i>1 Understanding Java servlets</i>	<i>3</i>
1.1 What is a servlet?	4
Server responsibilities	4 ♦ Server extensions 5
1.2 What is a servlet container?	5
The big picture	5 ♦ Understanding servlet containers 5
Using Tomcat	8
1.3 Hello World servlet	8
The code	8 ♦ Compilation 9 ♦ Deployment 9
Execution	10
1.4 The relationship between a servlet container and the Servlet API	10
The javax.servlet package	10 ♦ The javax.servlet.http package 11 ♦ Advantages and disadvantages of the Servlet API 12
1.5 Summary	13
<i>2 Understanding JavaServer Pages</i>	<i>14</i>
2.1 What is a JSP page?	15
Server-side includes	15

2.2	Hello User	15
	The HTML code	16 ♦ The servlet code 16
	The JSP code	17
2.3	Servlet or JSP?	17
2.4	JSP architecture models	18
	The Model 1 architecture	18 ♦ The Model 2 architecture 18
2.5	A note about JSP syntax	19
2.6	Summary	20
3	<i>Web application and HTTP basics</i>	21
3.1	What is a web application?	22
	Active and passive resources	22 ♦ Web applications and the web application server 22
3.2	Understanding the HTTP protocol	23
	HTTP basics	24 ♦ The structure of an HTTP request 24
	The structure of an HTTP response	26
3.3	Summary	27

Part 2 Servlets 29

4	<i>The servlet model</i>	31
4.1	Sending requests: Web browsers and HTTP methods	32
	Comparing HTTP methods	33
4.2	Handling HTTP requests in an HttpServlet	35
4.3	Analyzing the request	36
	Understanding ServletRequest	37 ♦ Understanding HttpServletResponse 37
4.4	Sending the response	40
	Understanding ServletResponse	40 ♦ Understanding HttpServletResponse 43
4.5	Servlet life cycle	45
	Loading and instantiating a servlet	46 ♦ Initializing a servlet 46
	Servicing client requests	47 ♦ Destroying a servlet 48
	Unloading a servlet	48 ♦ Servlet state transition from the servlet container's perspective 48
4.6	ServletConfig: a closer look	50
	ServletConfig methods	50 ♦ Example: a servlet and its deployment descriptor 50
4.7	ServletContext: a closer look	53

4.8	Beyond servlet basics	54
	Sharing the data (attribute scopes)	55 ♦ Coordinating servlets using RequestDispatcher
	Accessing request-scoped attributes with RequestDispatcher	58 ♦ Putting it all together: A simple banking application
	A simple banking application	59
4.9	Summary	63
4.10	Review questions	63
5	<i>Structure and deployment</i>	67
5.1	Directory structure of a web application	68
	Understanding the document root directory	68 ♦ Understanding the WEB-INF directory
	The web archive (WAR) file	70 ♦ Resource files and HTML access
	The default web application	70 ♦ The default web application
	71	
5.2	The deployment descriptor: an overview	71
	Example: A simple deployment descriptor	72 ♦ Using the <servlet> element
	Using the <servlet-mapping> element	73 ♦ Mapping a URL to a servlet
	76	
5.3	Summary	80
5.4	Review questions	80
6	<i>The servlet container model</i>	83
6.1	Initializing ServletContext	84
6.2	Adding and listening to scope attributes	85
	Adding and removing scope attributes	85 ♦ Listening to attribute events
	86	
6.3	Servlet life-cycle events and listeners	88
	javax.servlet.ServletContextListener	88
	javax.servlet.Http.HttpSessionListener	89
	javax.servlet.Http.HttpServletRequestListener	89
6.4	Adding listeners in the deployment descriptor	90
6.5	Web applications in a distributed environment	92
	Behavior of a ServletContext	92 ♦ Behavior of an HttpSession
	93	
6.6	Summary	94
6.7	Review questions	94
7	<i>Using filters</i>	97
7.1	What is a filter?	98
	How filtering works	99 ♦ Uses of filters
	99 ♦ The Hello World filter	100

7.2	The Filter API	102
	The Filter interface	103 ♦ The FilterConfig interface 105
	The FilterChain interface	105 ♦ The request and response wrapper classes 106
7.3	Configuring a filter	106
	The <filter> element	106 ♦ The <filter-mapping> element 107
	Configuring a filter chain	107
7.4	Advanced features	110
	Using the request and response wrappers	110 ♦ Important points to remember about filters 116 ♦ Using filters with MVC 116
7.5	Summary	117
7.6	Review questions	117
8	<i>Session management</i>	119
8.1	Understanding state and sessions	120
8.2	Using HttpSession	121
	Working with an HttpSession	122 ♦ Handling session events with listener interfaces 124 ♦ Invalidating a Session 130
8.3	Understanding session timeout	130
8.4	Implementing session support	131
	Supporting sessions using cookies	132 ♦ Supporting sessions using URL rewriting 133
8.5	Summary	136
8.6	Review questions	136
9	<i>Developing secure web applications</i>	139
9.1	Basic concepts	140
	Authentication	140 ♦ Authorization 140
	Data integrity	141 ♦ Confidentiality or data privacy 141
	Auditing	141 ♦ Malicious code 141 ♦ Web site attacks 141
9.2	Understanding authentication mechanisms	142
	HTTP Basic authentication	143 ♦ HTTP Digest authentication 145 ♦ HTTPS Client authentication 145
	FORM-based authentication	146 ♦ Defining authentication mechanisms for web applications 146
9.3	Securing web applications declaratively	149
	display-name	149 ♦ web-resource-collection 149
	auth-constraint	150 ♦ user-data-constraint 151
	Putting it all together	152
9.4	Securing web applications programmatically	156
9.5	Summary	158
9.6	Review questions	159

Part 3 JavaServer Pages and design patterns 163

10 The JSP technology model—the basics 165

- 10.1 SP syntax elements 166
 - Directives 167 ♦ Declarations 168 ♦ Scriptlets 169
 - Expressions 170 ♦ Actions 171 ♦ Comments 172
- 10.2 The JSP page life cycle 173
 - JSP pages are servlets 174 ♦ Understanding translation units 174 ♦ JSP life-cycle phases 175 ♦ JSP life-cycle example 178
- 10.3 Understanding JSP page directive attributes 181
 - The import attribute 182 ♦ The session attribute 182
 - The errorPage and isErrorPage attributes 182 ♦ The language and extends attributes 184 ♦ The buffer and autoFlush attributes 184 ♦ The info attribute 185 ♦ The contentType and pageEncoding attributes 185
- 10.4 Summary 186
- 10.5 Review questions 186

11 The JSP technology model—advanced topics 188

- 11.1 Understanding the translation process 189
 - Using scripting elements 189 ♦ Using conditional and iterative statements 191 ♦ Using request-time attribute expressions 194
 - Using escape sequences 194
- 11.2 Understanding JSP implicit variables and JSP implicit objects 198
 - application 200 ♦ session 201 ♦ request and response 202
 - page 202 ♦ pageContext 202 ♦ out 203 ♦ config 204
 - exception 206
- 11.3 Understanding JSP page scopes 207
 - Application scope 207 ♦ Session scope 207
 - Request scope 208 ♦ Page scope 209
- 11.4 JSP pages as XML documents 211
 - The root element 212 ♦ Directives and scripting elements 213
 - Text, comments, and actions 214
- 11.5 Summary 215
- 11.6 Review questions 216

12 Reusable web components 219

- 12.1 Static inclusion 220
 - Accessing variables from the included page 221 ♦ Implications of static inclusion 222

12.2	Dynamic inclusion	223
	Using jsp:include	223 ♦ Using jsp:forward
	Using jsp:forward	225
	Passing parameters to dynamically included components	226
	Sharing objects with dynamically included components	228
12.3	Summary	232
12.4	Review questions	232
13	<i>Creating JSPs with the Expression Language (EL)</i>	236
13.1	Understanding the Expression Language	237
	EL expressions and JSP script expressions	237 ♦ Using implicit variables in EL expressions
	Using implicit variables in EL expressions	238
13.2	Using EL operators	241
	EL operators for property and collection access	241
	EL arithmetic operators	242 ♦ EL relational and logical operators
	EL relational and logical operators	243
13.3	Incorporating functions with EL	244
	Creating the static methods	244 ♦ Creating a tag library descriptor (TLD)
	Creating a tag library descriptor (TLD)	245 ♦ Modifying the deployment descriptor
	Modifying the deployment descriptor	246 ♦ Accessing EL functions within a JSP
	Accessing EL functions within a JSP	247
13.4	Summary	249
13.5	Review questions	249
14	<i>Using JavaBeans</i>	251
14.1	JavaBeans: a brief overview	252
	JavaBeans from the JSP perspective	252 ♦ The JavaBean advantage
	The JavaBean advantage	253 ♦ Serialized JavaBeans
	Serialized JavaBeans	255
14.2	Using JavaBeans with JSP actions	258
	Declaring JavaBeans using <jsp:useBean>	258 ♦ Mutating properties using <jsp:setProperty>
	Mutating properties using <jsp:setProperty>	266 ♦ Accessing properties using <jsp:getProperty>
	Accessing properties using <jsp:getProperty>	269
14.3	JavaBeans in servlets	271
14.4	Accessing JavaBeans from scripting elements	274
14.5	More about properties in JavaBeans	276
	Using nonstring data type properties	276 ♦ Using indexed properties
	Using indexed properties	278
14.6	Summary	280
14.7	Review questions	281
15	<i>Using custom tags</i>	285
15.1	Getting started	286
	New terms	286 ♦ Understanding tag libraries
	Understanding tag libraries	287

- 15.2 Informing the JSP engine about a custom tag library 288
 - Location of a TLD file 289 ♦ Associating URIs with TLD file locations 290 ♦ Understanding explicit mapping 290
 - Resolving URIs to TLD file locations 291 ♦ Understanding the prefix 293
 - 15.3 Using custom tags in JSP pages 293
 - Empty tags 294 ♦ Tags with attributes 295 ♦ Tags with JSP code 296 ♦ Tags with nested custom tags 297
 - 15.4 Using the JSP Standard Tag Library (JSTL) 298
 - Acquiring and installing the JSTL 298 ♦ General purpose JSTL tags: `<c:catch>` and `<c:out>` 299 ♦ Variable support JSTL tags: `<c:set>` and `<c:remove>` 300 ♦ Flow control JSTL: `<c:if>`, `<c:choose>`, `<c:forEach>`, and `<c:forEachTokens>` 301
 - 15.5 Summary 305
 - 15.6 Review questions 305
- 16 *Developing “Classic” custom tag libraries* 309
- 16.1 Understanding the tag library descriptor 310
 - The `<taglib>` element 311 ♦ The `<tag>` element 313
 - The `<attribute>` element 314 ♦ The `<body-content>` element 316
 - 16.2 The Tag Extension API 318
 - 16.3 Implementing the Tag interface 320
 - Understanding the methods of the Tag interface 321
 - An empty tag that prints HTML text 324 ♦ An empty tag that accepts an attribute 326 ♦ A nonempty tag that includes its body content 328
 - 16.4 Implementing the IterationTag interface 329
 - Understanding the IterationTag methods 329 ♦ A simple iterative tag 330
 - 16.5 Implementing the BodyTag interface 333
 - Understanding the methods of BodyTag 334 ♦ A tag that processes its body 335
 - 16.6 Extending TagSupport and BodyTagSupport 338
 - The TagSupport class 338 ♦ The BodyTagSupport class 339
 - Accessing implicit objects 339 ♦ Writing cooperative tags 343
 - 16.7 What’s more? 347
 - 16.8 Summary 348
 - 16.9 Review questions 349

<i>17</i>	<i>Developing “Simple” custom tag libraries</i>	<i>352</i>
17.1	Understanding SimpleTags	353
	A brief example	353 ♦ Exploring SimpleTag and SimpleTagSupport
		354
17.2	Incorporating SimpleTags in JSPs	357
	Coding empty SimpleTags	357 ♦ Adding dynamic attributes to SimpleTags
		359 ♦ Processing body content inside SimpleTags
		362
17.3	Creating Java-free libraries with tag files	364
	Introducing tag files	364 ♦ Tag files and TLDs
		365
	Controlling tag processing with tag file directives	366
	Processing fragments and body content with tag file actions	368
17.4	Summary	371
17.5	Review questions	372

18 Design patterns 376

18.1	Design patterns: a brief history	377
	The civil engineering patterns	377 ♦ The Gang of Four patterns
		377 ♦ The distributed design patterns
		379
	The J2EE patterns	379
18.2	Patterns for the SCWCD exam	382
	The pattern template	382 ♦ The Intercepting Filter
		385
	Model-View-Controller (MVC)	386 ♦ Front Controller
		389
	Service Locator	391 ♦ Business Delegate
		393
	Transfer Object	397
18.3	Summary	400
18.4	Review questions	401

Appendices

<i>A</i>	<i>Installing Tomcat 5.0.25</i>	<i>403</i>
<i>B</i>	<i>A sample web.xml file</i>	<i>408</i>
<i>C</i>	<i>Review Q & A</i>	<i>412</i>
<i>D</i>	<i>Exam Quick Prep</i>	<i>475</i>
	<i>index</i>	<i>523</i>

preface to the second edition

When I first considered taking the Sun Certified Web Component Developer (SCWCD) exam, I thought it was going to be a breeze. After all, I'd deployed some servlets and I had a solid working knowledge of JavaServer Pages (JSPs). But before I registered, I figured a few simulation questions couldn't hurt. What an eye-opener! The questions seemed better suited to Trivial Pursuit than a software exam. How could these sadists ask for every Java exception, interface method, and XML element? Do I look like a Javadoc?

I bought a few books covering the exam, but Manning's *SCWCD Exam Study Kit* stood out from the rest. With its in-depth explanations, multiple helpful appendices, and powerful simulation software, it became apparent that this was something special. Building this immense course must have been a labor of love, and the authors' dedication shone on every page. It goes without saying that I passed the exam with flying colors.

When Manning approached me to assist in creating a second edition for the new 310–081 exam, I was honored and nervous. Hanumant and Jignesh had set the standard for clarity and precise technical understanding, and it would take no small effort to maintain their degree of merit. But, after passing the new exam, I looked forward to presenting Sun's new features for simplifying web development, including the Expression Language, the JSP Standard Tag Library, and SimpleTag development. This new edition covers these topics and more, holding as closely as possible to the quality of its predecessor.

MATTHEW SCARPINO

preface to the first edition

We first started thinking about writing this book when we were preparing to take the Sun Certified Web Component Developer (SCWCD) exam. We had difficulty finding any books that thoroughly covered the objectives published by Sun. The idea continued to percolate during the time we were developing JWebPlus, our exam simulator for the SCWCD. With its successful release, we finally turned our attention to putting our combined knowledge and experience into this book.

We have been interacting with Java Certification aspirants for a long time. Through our discussion forums and our exam simulators, JWebPlus and JQPlus (for SCJP—Sun Certified Java Programmer), we have helped people gain the skills they need. Our goal in this book is to leverage that experience and help you feel confident about taking the exam. This book and the accompanying CD will prepare you to do so; they are all you need to pass with flying colors. Of course, you'll still have to write a lot of code yourself!

HANUMANT DESHMUKH
JIGNESH MALAVIA

acknowledgments

No book gets published without the hard work of a lot of people. We are very grateful to...

Michael Curwen, who tech-proofed all the chapters in the second edition and added material where appropriate. His detailed knowledge of J2EE ensured that the material in this book was presented clearly and accurately.

Our reviewers, who provided valuable feedback and comments: Rob Abbe, Phil Hanna, William Lopez, and Muhammad Ashikuzzaman.

Our publisher, Marjan Bace for his guidance and encouragement, and the entire publishing team at Manning: Liz Welch for her incredible patience in copyediting, Karen Tegtmeyer for setting up the reviews, Susan Forsyth for proofreading, Denis Dalinnik for typesetting the manuscript, and Mary Piergies for managing the production process. Also the terrific crew in the back office who printed the book and brought it to the market in record time.

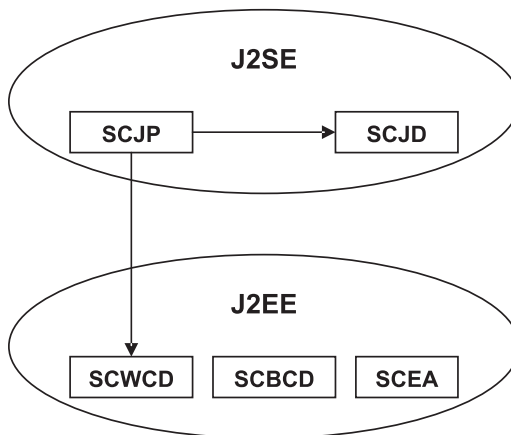
Finally, our kudos to Jackie Carter. She took great care with the “presentation logic” throughout the book and put in an incredible amount of effort to format and polish every chapter. She made sure that the concepts were explained in a clear and professional manner. We cannot thank her enough for all the hard work she put in to help us shape a better book.

about the Sun certification exams

The Java platform comes in three flavors: Standard Edition, Enterprise Edition, and Micro Edition. The figure below shows the certification exams that Sun offers for the first two editions.

The Standard Edition (J2SE) is the basis of the Java platform and is used in the development of Java applets and applications. The standard library includes important packages, such as `java.io`, `java.net`, `java.rmi`, and `javax.swing`. Sun offers two certifications for this platform: the Java Programmer (SCJP) certification and the Java Developer (SCJD) certification. While the Java Programmer certification process consists of only one multiple-choice exam covering the basics of the Java language, the Java Developer certification requires you to develop a simple but nontrivial client server application using the `java.net`, `java.rmi`, and `javax.swing` packages, followed by an essay-type exam on the application.

The Enterprise Edition (J2EE) builds on the Standard Edition and includes a number of technologies, such as Enterprise JavaBeans (EJB), Servlet, and JavaServer Pages, used for building enterprise-class server-side applications. Sun offers three certifications for this platform: the Web Component Developer (SCWCD) certification, the Business Component Developer (SCBCD) certification, and the Enterprise Architect (SCEA) certification. The SCWCD certification process is designed for programmers developing



A roadmap for Sun's certifications in the J2SE and the J2EE platforms. SCJP certification is required before taking the SCWCD exam.

web applications using Servlet and JSP technology and consists of one multiple-choice exam. You must be a Sun Certified Java Programmer (SCJP) before you can take this exam. The Business Component Developer certification is for developers creating applications with Enterprise JavaBeans (EJBs) and EJB containers. The Enterprise Architect certification is designed for senior developers who are using the whole gamut of J2EE technologies to design enterprise-class applications. The certification process consists of one multiple-choice exam and one architecture and design project, followed by an essay-type exam on the project.

The Micro Edition (J2ME) is an optimized Java runtime environment meant for use in consumer electronic products, such as cell phones and pagers.

Preparing for the SCWCD exam

We believe that studying for a test is very different than just learning a technology. Of course, you also learn the technology when you study for the test. But when you take the exam, you have to show that you understand what the examiner expects you to know about the technology. And that's what makes studying for a test a different ball game altogether. It is not surprising that even people with many years of experience sometimes fail the tests. In this book, we'll teach you the technology while training you for the test.

Here are the things that you will need:

- *A copy of the exam objectives.* It is very important to take a look at the objectives before you start a chapter and after you finish it. It helps to keep you focused. For your convenience, we have included the relevant exam objectives at the beginning of each chapter, as well as in appendix D.
- *A Servlet engine that implements the Servlet 2.4 and JSP 2.0 specifications.* You will need it because we'll do some coding exercises to illustrate the concepts. In this book, we have decided to use Tomcat 5.0.25 because it is now the official reference implementation for the JSP/Servlet technology and it conforms to the specifications. In addition, it is free and easy to install and run. Appendix A explains where to get Tomcat 5.0.25 and how to install it. If you are clueless about what Tomcat is, don't worry. Chapters 1 and 2 will bring you up to speed.
- *A copy of the Servlet 2.4 and JSP 2.0 specifications.* The specifications are the best source of information on this technology. Don't get scared; unlike the Java Language specs, these specs are readable and easy to understand. You can download the specs for Servlet 2.4 from <http://www.jcp.org/aboutJava/communityprocess/final/jsr154/> and for JSP 2.0 from <http://jcp.org/aboutJava/communityprocess/final/jsr152/>.
- *The JWebPlus exam simulator.* We've developed this exam simulator to help you judge your level of preparedness. It not only includes detailed explanations of the questions but also explains why a certain option is right or wrong. You can download an abbreviated version of this tool from www.manning.com/deshmukh2. You can buy the full version at www.enthuware.com.

Taking the SCWCD exam

Exam code: 310–081

Cost: \$150

Number of questions: 69 multiple-choice questions

The questions tell you the number of correct answers. You may also get questions that ask you to match options on the left side with options on the right side, or that ask you to drag and drop options to the correct place. In general, many exam takers have reported that questions on this test are easier than the ones on the Sun Certified Java Programmer's exam. The exam starts with a survey that asks you questions about your level and experience with Servlet/JSP technology, but these questions are not a part of the actual exam.

At the time of this writing, the duration of the test was 135 minutes. But Sun has changed the duration for the SCJP exam a couple of times, so they could change the duration of this test as well. Please verify it before you take the exam. You can get the latest information about the exam from <http://suned.sun.com>.

Here's how to register and what to expect:

- First, purchase an exam voucher from your local Sun Educational Services office. In the United States, you can purchase an exam voucher by visiting the Sun web site, at www.sun.com/training/catalog/courses/CX-310-081.xml. If you reside outside the United States, you should contact your local Sun Educational Services office. You'll be given a voucher number.
- Tests are conducted by Prometric all across the world. You have to contact them to schedule the test. Please visit the Prometric web site at www.2test.com for information about testing centers. Before you schedule the test, check out the testing center where you plan to take the exam. Make sure you feel comfortable with the environment there. Believe us, you do not want to take the test at a noisy place. Once you finalize the center, you can schedule the test.
- You should reach the testing center at least 15 minutes before the test, and don't forget to take two forms of ID. One of the IDs should have your photograph on it.
- After you finish the test, the screen will tell you whether or not you passed. You will need a score of 62% in order to pass (43 correct answers out of 69 questions). You will receive a printed copy of the detailed results.

Best of luck!

about this book

This book is built around the objectives that Sun has published for the updated SCWCD exam. If you know everything that is covered by the objectives, you will pass the exam. The chapters in the book examine each objective in detail and explain everything you need to understand about web component development.

Who is this book for?

This book is for Java programmers who want to prepare for the SCWCD exam, which focuses on the Servlet and JavaServer Pages technologies. This book will also be very useful for beginners since we have explained the concepts using simple examples. The text will bring you up to speed even if you are totally new to these technologies. Even expert Servlet/JSP programmers should read the book to ensure that they do not overlook any exam objectives. However, since this book is a study guide, we do not try to cover advanced tricks and techniques for expert Servlet/JSP developers.

How this book is organized

This book has three parts:

Part	Topic	Chapters
1	The basics of web component development	1 through 3
2	The Servlet technology	4 through 9
3	The JavaServerPages (JSP) technology and design patterns	10 through 18

For those of you new to web component development, we've included one introductory chapter each on Servlets and JavaServer Pages. The objectives of chapters 1 and 2 are to make you comfortable with this technology. They won't make you an expert, but they'll teach you enough so that you can understand the rest of the book. If you already have experience with the Servlet and JavaServerPages technologies, you can skip these two chapters. Since in practice servlets are written for HTTP, we have also included a brief discussion of the HTTP protocol and the basics of web applications in chapter 3. You should read this chapter even if you know the HTTP protocol.

Chapters 4 through 18 cover the exam objectives. Some chapters start with basic concepts that do not necessarily correspond to exam objectives but that are very important in order to understand the remaining sections. In the chapters, we illustrate the concepts with simple test programs. You should try to write and run the programs, and we encourage you to modify them and try out similar examples. From our experience, we've seen that people tend to understand and remember the concepts a lot better if they actually put them in code and see them in action.

There are four appendices. Appendix A will help you set up Tomcat. Appendix B contains a sample web.xml file that illustrates the use of various deployment descriptor tags. Appendix C contains the answers to each chapter's review questions. In appendix D, you will find the Quick Prep, a summary of key concepts and helpful tips that you can review as part of your last-minute exam preparations.

How each chapter is organized

After the introductory chapters in part 1, each chapter begins with a list of the exam objectives that are discussed within it, along with the chapter sections in which each objective is addressed. In some of the chapters, the order of the objectives departs slightly from the original Sun numbering to better correspond to the way the topics within the chapters have been organized.

As you read through the chapters, you will encounter Quizlets about the material you have just read. Try to answer the Quizlet without looking at the answer; if you are correct, you can feel confident that you have understood the concepts.

At the end of each chapter, you will find review questions that will help you to evaluate your ability to answer the exam questions related to the objectives for the chapter. The answers to these questions are in appendix C.

Code conventions

Italic typeface is used to introduce new terms.

Courier typeface is used to denote code samples, as well as elements and attributes, method names, classes, interfaces, and other identifiers.

Bold courier is used to denote important parts of the code samples.

Code annotations accompany many segments of code.

Line continuations are indented.

Downloads

Source code for all the programming examples in this book is available for download from the publisher's web site, www.manning.com/deshmukh2. Any corrections to code will be updated on an ongoing basis.

Also available for download is the abbreviated version of the JWebPlus exam simulator which contains a practice exam. Please go to www.manning.com/deshmukh2 to download the exam simulator and follow the instructions that accompany the file.

System requirements for JWebPlus are:

- OS: Win 98, NT, 2000, XP, Must have IE 5.0 or later version.
- Processor (Min Speed): AMD/Intel Pentium (500MHz)
- Min RAM: 128MB
- HDD space: 2 MB

Author Online

Purchase of the *SCWCD Exam Study Kit Second Edition* includes free access to a private web forum run by Manning Publications, where you can make comments about the book, ask technical questions, and receive help from the authors and from other users. To access the forum and subscribe to it, point your web browser to www.manning.com/deshmukh2. This page provides information on how to get on the forum once you are registered, what kind of help is available, and the rules of conduct on the forum.

Manning's commitment to our readers is to provide a venue where a meaningful dialogue between individual readers and between readers and the authors can take place. It is not a commitment to any specific amount of participation on the part of the authors, whose contribution to the AO remains voluntary (and unpaid). We suggest you try asking the authors some challenging questions lest their interest stray!

The Author Online forum and the archives of previous discussions will be accessible from the publisher's web site as long as the book is in print.

You can also reach the authors through their web site at www.jdiscuss.com, where they maintain forums for the discussion of Java topics, especially those related to the Sun exams. Additionally, the web site contains material that you will find useful in your preparation for the exam, such as information about books, tutorials, free and commercial practice exams, and study notes. The site will continue to be updated with exciting new resources as they become available.

about the authors

HANUMANT DESHMUKH is the president and founder of Enthware.com Pvt. Ltd. He also manages www.jdiscuss.com, a free site designed for Java certification aspirants. He has been working in the information technology industry for over eight years, mainly consulting for projects with the Distributed Object Oriented System using J2EE technologies. Hanumant also designs and develops the Java certification software for his company. The exam simulators from Enthware.com, JQPlus (for SCJP) and JWeb-Plus (for SCWCD), are well known and respected in the Java community.

JIGNESH MALAVIA is a senior technical architect at SourceCode, Inc. in New York. For over eight years, he has been involved in the design and development of various types of systems, from language interpreters to business applications. Teaching is one of his passions, and he has taught courses on Java and web development, as well as C, C++, and Unix, at various locations, including the Narsee Monjee Institute of Management Science (NMIMS), Mumbai. He has been actively involved with Enthware projects and currently provides online guidance to candidates preparing for Sun certification exams.

MATTHEW SCARPINO is a Sun Certified Web Component Developer and has developed a number of web sites for business. He has worked with Java for over six years, with particular emphasis on the Eclipse IDE. He has been recently involved in designing with Eclipse's Rich Client Platform seeks to extend these applications across a network.

JACQUELYN CARTER is an editor and technical writer who also has many years' experience providing information technology solutions for organizations in both the business and nonprofit worlds.

about the cover illustration

The figure on the cover of *SCWCD Exam Study Kit Second Edition* is taken from a Spanish compendium of regional dress customs first published in Madrid in 1799. The book's title page states:

Coleccion general de los Trages que usan actualmente todas las Naciones del Mundo desubierto, dibujados y grabados con la mayor exactitud por R.M.V.A.R. Obra muy util y en special para los que tienen la del viajero universal

which we translate, as literally as possible, thus:

General collection of costumes currently used in the nations of the known world, designed and printed with great exactitude by R.M.V.A.R. This work is very useful especially for those who hold themselves to be universal travelers

Although nothing is known of the designers, engravers, and workers who colored this illustration by hand, the “exactitude” of their execution is evident in this drawing which is just one of many figures in this colorful collection. Their diversity speaks vividly of the uniqueness and individuality of the world's towns and regions just 200 years ago. This was a time when the dress codes of two regions separated by a few dozen miles identified people uniquely as belonging to one or the other. The collection brings to life a sense of isolation and distance of that period—and of every other historic period except our own hyperkinetic present.

Dress codes have changed since then and the diversity by region, so rich at the time, has faded away. It is now often hard to tell the inhabitant of one continent from another. Perhaps, trying to view it optimistically, we have traded a cultural and visual diversity for a more varied personal life. Or a more varied and interesting intellectual and technical life.

We at Manning celebrate the inventiveness, the initiative, and, yes, the fun of the computer business with book covers based on the rich diversity of regional life of two centuries ago, brought back to life by the pictures from this collection.

P A R T

1

Getting started

Part 1 is intended for readers who are new to web component development. We introduce you to the concepts you'll need to understand before you begin the chapters that focus on the exam objectives. Our topics here include the Servlet and JSP technologies, web applications, and the HTTP protocol.



CHAPTER 1

Understanding Java servlets

- 1.1 What is a servlet? 4
- 1.2 What is a servlet container? 5
- 1.3 Hello World servlet 8
- 1.4 The relationship between a servlet container and the Servlet API 10
- 1.5 Summary 13

INTRODUCTION

The goal of this book is to explain how you can use J2EE to create these dynamic web components. We'll do this by discussing servlets and JavaServer Pages (JSPs) in great technical depth. We'll present the theory behind these concepts, and then supplement the theory with practical code. Then, by using Tomcat or a similar web server, you can construct your own code to cement the material in your mind.

1.1 **WHAT IS A SERVLET?**

As is apparent from its name, a servlet is a server-side entity. But what exactly does it mean? Is it a new design pattern for writing servers? Is it a new Java class? Or is it a new technology? The answer to all these questions is yes, albeit in different contexts. To understand any new concept, it is important to know the reasons behind its conception. So, let's start by having a look at the tasks a server needs to do.

1.1.1 **Server responsibilities**

Every server that provides services to remote clients has two main responsibilities. The first is to handle client requests; the second is to create a response to be sent back. The first task involves programming at the socket level, extracting information from request messages, and implementing client-server protocols, such as FTP and HTTP. The second task, creating the response, varies from service to service. For example, in the case of FTP servers that serve file transfer requests, response creation is as simple as locating a file on the local machine. On the other hand, HTTP servers that host full-fledged web applications are required to be more sophisticated in the way they generate output. They have to create the response dynamically, which may involve complicated tasks, such as retrieving data from the database, applying business rules, and presenting the output in the formats desired by different clients.

One way to write a simple server that serves only static data would be to code everything in a single executable program. This single program would take care of all the different chores, such as managing the network, implementing protocols, locating data, and replying. However, for HTTP servers that serve syndicated data, we require a highly flexible and extensible design. Application logic keeps changing, clients need personalized views of information, and business partners need customized processing rules. We cannot write a single program that handles all these tasks. Furthermore, what if a new functionality has to be added? What if the data format changes? Modifying the source files (especially after the developer has left!) to add new code is surely the last thing we want to do.

Well, there is a better design for these kinds of servers: divide the code into two executable parts—one that handles the network and one that provides the application logic—and let the two executables have a standard interface between them. This kind of separation makes it possible to modify the code in the application logic without affecting the network module, as long as we follow the rules of the interface. Traditionally, people have implemented this design for HTTP servers using Common Gateway Interface (CGI). On one side of this interface is the main web server, and on the other side are the CGI scripts. The web server acts as the network communications module and manages the clients, while the CGI scripts act as data processing modules and deliver the output. They follow the rules of the “common gateway interface” to pass data between them.

1.1.2 Server extensions

Although CGI provides a modular design, it has several shortcomings. The main issue for high-traffic web sites is scalability. Each new request invocation involves the creation and destruction of new processes to run the CGI scripts. This is highly inefficient, especially if the scripts perform initialization routines, such as connecting to a database. Moreover, they use file input/output (I/O) as a means of communication with the server, causing a significant increase in the overall response time.

A better way is to have the server support separate executable modules that can be loaded into its memory and initialized only once—when the server starts up. Each request can then be served by the already in-memory and ready-to-serve copy of the modules. Fortunately, most of the industrial-strength servers have been supporting such modules for a long time, and they have made the out-of-memory CGI scripts obsolete. These separate executable modules are known as *server extensions*. On platforms other than Java, server extensions are written using native-language APIs provided by the server vendors. For example, Netscape Server provides the Netscape Server Application Programming Interface (NSAPI), and Microsoft's Internet Information Server (IIS) provides the Internet Server Application Programming Interface (ISAPI). In Java, server extensions are written using the Servlet API,¹ and the server extension modules are called *servlets*.

1.2 WHAT IS A SERVLET CONTAINER?

A web server uses a separate module to load and run servlets. This specialized module, which is dedicated to servlet management, is called a *servlet container*, or *servlet engine*.

1.2.1 The big picture

Figure 1.1 shows how different components fit into the big picture. HTML files are stored in the file system, servlets run within a servlet container, and business data is in the database.

The browser sends requests to the web server. If the target is an HTML file, the server handles it directly. If the target is a servlet, the server delegates the request to the servlet container, which in turn forwards it to the servlet. The servlet uses the file-system and database to generate dynamic output.

1.2.2 Understanding servlet containers

Conceptually, a servlet container is a part of the web server, even though it may run in a separate process. In this respect, servlet containers are classified into the following three types:

¹ An overview of the Servlet API is given in section 1.4. The details of the different elements of this API are explained in chapters 4 through 9.




Figure 1.1 The big picture: all the components of a web-based application

- *Standalone*—Servlet containers of this type are typically Java-based web servers where the two modules—the main web server and the servlet container—are integral parts of a single program (figure 1.2).




Figure 1.2
A standalone
servlet container

Tomcat (we'll learn about Tomcat shortly) running all by itself is an example of this type of servlet container. We run Tomcat as we would any normal Java program inside a Java Virtual Machine (JVM). It contains handlers for static content, like HTML files, and handlers for running servlets and JSP pages.

- *In-process*—Here, the main web server and the servlet container are different programs, but the container runs within the address space of the main server as a plug-in (figure 1.3).

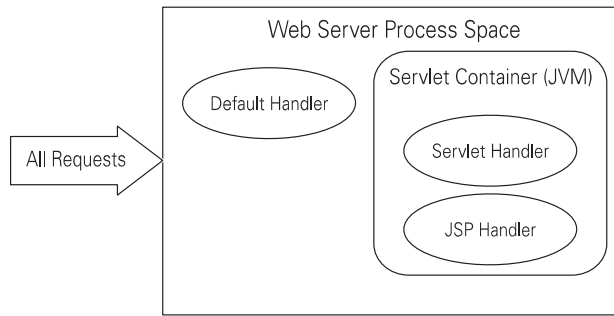


Figure 1.3
An in-process
servlet container

An example of this type is Tomcat running inside Apache Web Server. Apache loads a JVM that runs Tomcat. In this case, the web server handles the static content by itself, and Tomcat handles the servlets and JSP pages.

- *Out-of-process*—Like in-process servers, the main web server and the servlet container are different programs. However, with out-of-process, the web server runs in one process while the servlet container runs in a separate process (figure 1.4). To communicate with the servlet container, the web server uses a plug-in, which is usually provided by the servlet container vendor.

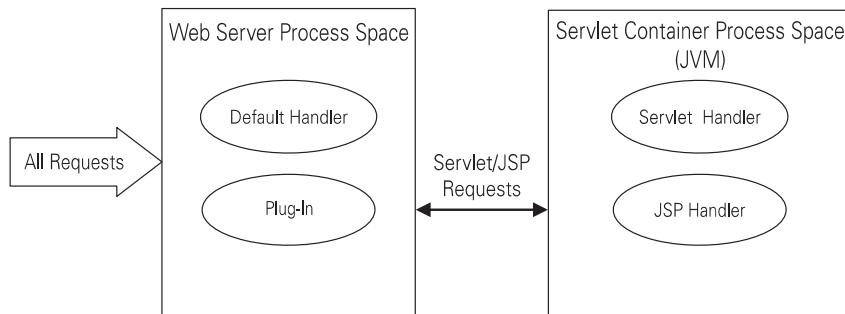


Figure 1.4 An out-of-process servlet container

An example of this type is Tomcat running as a separate process configured to receive requests from Apache Web Server. Apache loads the `mod_jk` plug-in to communicate with Tomcat.

Each of these types has its advantages, limitations, and applicability. We will not discuss these details, since they are beyond the scope of this book.

Many servlet containers are available on the market—Tomcat (Apache), Resin (Caucho Technology), JRun (Macromedia), WebLogic (BEA), and WebSphere (IBM), just to name a few. Some of these, like WebLogic and WebSphere, are much more than just servlet containers. They also provide support for Enterprise JavaBeans (EJB), Java Message Service (JMS), and other J2EE technologies.

1.2.3 Using Tomcat

Tomcat is a servlet container developed under the Jakarta project at the Apache Software Foundation (ASF). You can get a wealth of information about Tomcat from <http://jakarta.apache.org/tomcat>. We have decided to use Tomcat version 5.0.25 for the examples in this book because of the following reasons:

- It is free.
- It implements the latest Servlet 2.4 and JSP 2.0 specifications, which is what we need for the exam.
- It has the capability of running as a web server by itself (Standalone mode). There is no need for a separate web server.

We have given installation instructions for Tomcat in appendix A. In the discussions of the examples throughout the book, we have assumed that the Tomcat installation directory is `c:\jakarta-tomcat-5.0.25`. Note that once you have installed Tomcat, you must set the `CATALINA_HOME`, `JAVA_HOME`, and `CLASSPATH` variables, as described in appendix A.

1.3 HELLO WORLD SERVLET

In this section, we will look at the four basic steps—coding, compiling, deploying, and running—required to develop and run the customary Hello World servlet,² which prints Hello World! in the browser window. By the way, do you know who started the trend of writing “Hello World!” as an introductory program?³

1.3.1 The code

Listing 1.1 contains the code for `HelloWorldServlet.java`.

Listing 1.1 `HelloWorldServlet.java`

```
import java.io.*;
import javax.servlet.*;
import javax.servlet.http.*;
public class HelloWorldServlet extends HttpServlet
{
    public void service(HttpServletRequest request,
                        HttpServletResponse response)
        throws ServletException,
            IOException
    {
        PrintWriter pw = response.getWriter();
        pw.println("<html>");
        pw.println("<head>");
    }
}
```

² The details of the code will become clear as we move through the chapters.

³ Kernighan, Brian and Ritchie, Dennis. *The C Programming Language*. Prentice-Hall, 1988.

```

        pw.println("</head>");
        pw.println("<body>");
        pw.println("<h3>Hello World!</h3>");
        pw.println("</body>");
        pw.println("</html>");
    }
}

```

1.3.2 Compilation

Note the import statements in listing 1.1. They import the classes from the `javax.servlet` and `javax.servlet.http` packages. In Tomcat, they are provided as part of the `servlet-api.jar` file, which is in the directory `c:\jakarta-tomcat-5.0.25\common\lib\`. To compile the program in listing 1.1, include the JAR file in the classpath, as directed in appendix A. We will explain the details of these packages in section 1.4.

1.3.3 Deployment

Deployment is a two-step process. (We'll discuss the deployment structure in chapter 5.) First, we put the resources into the required directory. Then, we inform Tomcat about our servlet by editing the `web.xml` file:

- 1 Copy the `HelloWorldServlet.class` file to the directory

```
c:\jakarta-tomcat-5.0.25\webapps\chapter01\WEB-INF\classes
```

- 2 Create a text file named `web.xml` in the `c:\jakarta-tomcat-5.0.25\web-apps\chapter01\WEB-INF` directory. Write the following lines in the file:

```

<?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_4.xsd"
  version="2.4">
  <servlet>
    <servlet-name>HelloWorldServlet</servlet-name>
    <servlet-class>HelloWorldServlet</servlet-class>
  </servlet>
</web-app>

```

You can also copy the `chapter01` directory directly from the Manning web site to your `c:\jakarta-tomcat-5.0.25\webapps` directory. This will provide all the files you need to run the example.

1.3.4 Execution

Start Tomcat with a shortcut or with the DOS prompt (c:\jakarta-tomcat-5.0.25\bin\startup.bat). Open a browser window and go to the URL `http://localhost/chapter01/servlet/HelloWorldServlet`.

Hello World! should appear in the browser window.

1.4 THE RELATIONSHIP BETWEEN A SERVLET CONTAINER AND THE SERVLET API

Sun's Servlet specification provides a standard and a platform-independent framework for communication between servlets and their containers. This framework is made up of a set of Java interfaces and classes. These interfaces and classes are collectively called the *Servlet Application Programming Interfaces*, or the *Servlet API*. Simply put, we develop servlets using this API, which is implemented by the servlet container (see figure 1.5). The Servlet API is all we as servlet developers need to know. Since all the servlet containers must provide this API, the servlets are truly platform- and servlet container-independent. Essentially, understanding the rules of this API and the functionality that it provides is what servlet programming is all about!

The Servlet API is divided into two packages: `javax.servlet` and `javax.servlet.http`. We will discuss these packages in more detail as we progress through the book, but for now, let's take a quick look at them.

1.4.1 The `javax.servlet` package

This package contains the generic servlet interfaces and classes that are independent of any protocol.

The `javax.servlet.Servlet` interface

This is the central interface in the Servlet API. Every servlet class must directly or indirectly implement this interface. It has five methods, as shown in table 1.1.

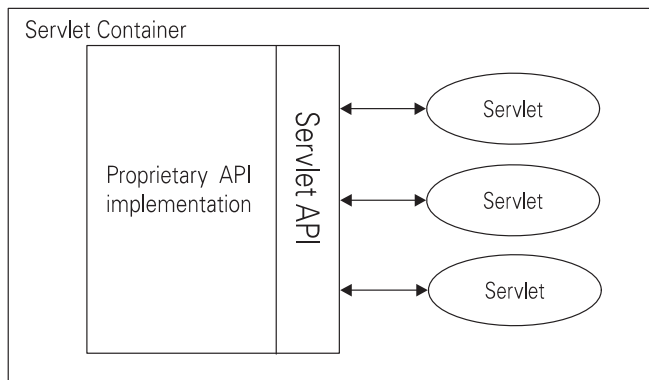


Figure 1.5
Servlets interact with the servlet container through the Servlet API.

Table 1.1 Methods of the `javax.servlet.Servlet` interface

Method	Description
<code>init()</code>	This method is called by the servlet container to indicate to the servlet that it must initialize itself and get ready for service. The container passes an object of type <code>ServletConfig</code> as a parameter.
<code>service()</code>	This method is called by the servlet container for each request from the client to allow the servlet to respond to the request.
<code>destroy()</code>	This method is called by the servlet container to indicate to the servlet that it must clean up itself, release any required resources, and get ready to go out of service.
<code>getServletConfig()</code>	Returns information about the servlet, such as a parameter to the <code>init()</code> method.
<code>getServletInfo()</code>	The implementation class must return information about the servlet, such as the author, the version, and copyright information.

The `service()` method handles requests and creates responses. The servlet container automatically calls this method when it gets any request for this servlet. The complete signature of this method is

```
public void service (ServletRequest, ServletResponse)
    throws ServletException, java.io.IOException;
```

The `javax.servlet.GenericServlet` class

The `GenericServlet` class implements the `Servlet` interface. It is an abstract class that provides implementation for all the methods except the `service()` method of the `Servlet` interface. It also adds a few methods to support logging. We can extend this class and implement the `service()` method to write any kind of servlet.

The `javax.servlet.ServletRequest` interface

The `ServletRequest` interface provides a generic view of the request that was sent by a client. It defines methods that extract information from the request.

The `javax.servlet.ServletResponse` interface

The `ServletResponse` interface provides a generic way of sending responses. It defines methods that assist in sending a proper response to the client.

1.4.2 The `javax.servlet.http` package

This package provides the basic functionality required for HTTP servlets. Interfaces and classes in this package extend the corresponding interfaces and classes of the `javax.servlet` package to build support for the HTTP protocol.

The javax.servlet.http.HttpServlet class

`HttpServlet` is an abstract class that extends `GenericServlet`. It adds a new `service()` method with this signature:

```
protected void service (HttpServletRequest, HttpServletResponse)
    throws ServletException, java.io.IOException;
```

In the Hello World example, we extended our servlet class from this class and we overrode the `service()` method.

The javax.servlet.http.HttpServletRequest interface

The `HttpServletRequest` interface extends `ServletRequest` and provides an HTTP-specific view of the request. It defines methods that extract information, such as HTTP headers and cookies, from the request.

The javax.servlet.http.HttpServletResponse interface

The `HttpServletResponse` interface extends `ServletResponse` and provides an HTTP-specific way of sending responses. It defines methods that assist in setting information, such as HTTP headers and cookies, into the response.

1.4.3 Advantages and disadvantages of the Servlet API

The advantages of the Servlet API are as follows:

- *Flexibility*—Each time we need to add new functionality to the server, all we have to do is write a new servlet specific to that set of requirements and plug it into the server, without modifying the server itself.
- *Separation of responsibilities*—The main server now only needs to worry about the network connections and communications part. The job of interpreting requests and creating appropriate responses is delegated to the servlets.
- *It's Java*—Java programmers don't need to learn a new scripting language. Also, they can use all the object-oriented features provided by Java.
- *Portability*—We can develop and test a servlet in one container and deploy it in another. Unlike proprietary solutions, the Servlet API is independent of web servers and servlet containers. We can “write once, run anywhere,” as long as the containers support the standard Servlet API.

One obvious limitation, or rather restriction, of the Servlet API is one that is common to all kinds of frameworks: you have to stick to the rules set forth by the framework. This means we have to follow certain conventions to make the servlet container happy.

Another disadvantage involves the containers available in the market and not the Servlet API itself. Theoretically, using the API, you can write servlets for almost any kind of protocol, including FTP, SMTP, or even proprietary protocols. Nevertheless, it would not be fair to expect the servlet container providers to build support for all

of them. As of now, the Servlet specification mandates support only for HTTP through the `javax.servlet.http` package.

1.5 **SUMMARY**

In this chapter, we learned about the basics of servlets and the servlet container, and how they provide extensions to a server's functionality. We also ran a sample `HelloWorld` servlet that displayed a line of text in the browser window. Finally, we looked at the Servlet API and its classes and interfaces.

Armed with this knowledge, we can now answer the question “What is a servlet?” from several different perspectives. Conceptually, a servlet is a piece of code that can be

- Plugged into an existing server to extend the server functionality
- Used to generate the desired output dynamically

For a servlet container, a servlet is

- A Java class like any other normal Java class
- A class that implements the `javax.servlet.Servlet` interface

For a web component developer, a servlet, or specifically an HTTP servlet, is a class that

- Extends `javax.servlet.http.HttpServlet`
- Resides in a servlet container (such as Tomcat or JRun)
- Serves HTTP requests



CHAPTER 2

Understanding JavaServer Pages

- | | | | |
|-----|------------------------|-----|----------------------------|
| 2.1 | What is a JSP page? 15 | 2.4 | JSP architecture models 18 |
| 2.2 | Hello User 15 | 2.5 | A note about JSP syntax 19 |
| 2.3 | Servlet or JSP? 17 | 2.6 | Summary 20 |

INTRODUCTION

Part 3 of this book addresses the exam objectives that apply to JavaServer Pages (JSP). For those of you who are just learning about JSP technology, this chapter will give you all the information you need to get started.

2.1 **WHAT IS A JSP PAGE?**

A JSP page is a web page that contains Java code along with the HTML tags. Like any other web page, a JSP page has a unique URL, which is used by the clients to access the page. When accessed by a client, the Java code within the page is executed on the server side, producing textual data. This data, which is surrounded by HTML tags, is sent as a normal HTML page to the client. Since the Java code embedded in a JSP page is processed on the server side, the client has no knowledge of the code. The code is replaced by the HTML generated by the Java code before the page is sent to the client. Before we discuss how to create JSP pages, let's discuss the need for such a technology.

2.1.1 **Server-side includes**

HTML is a markup language that specifies how to label different parts of data for visual presentation. The hyperlinks provide a way to jump from one piece of information to another. However, the content is already inside the HTML tags. The tags do not create it; they merely decorate it for presentation. HTML by itself produces static web pages, but today, it is necessary for most web sites to have dynamic content. To generate the content dynamically, we need something that can allow us to specify business logic and that can generate data in response to a request. The data can then be formatted using HTML.

A dynamic web page consists of markup language code as well as programming language code. Instead of serving the page as is to the clients, a server processes the programming language code, replaces the code with the data generated by the code, and then sends the page to the client. This methodology of embedding programming languages within HTML is called the *server-side include* and the programming language that is embedded within the HTML is called the *scripting language*. For example, Netscape's Server-Side JavaScript (SSJS) and Microsoft's Active Server Pages (ASP) are examples of server-side includes. They use JavaScript and VBScript, respectively, as the scripting languages. JavaServer Pages is the name of the technology that provides a standard specification for combining Java as the scripting language with HTML. It forms the presentation layer of Sun's Java 2 Enterprise Edition (J2EE) architecture.

The JSP specification lists the syntax and describes the semantics of the various elements that make up a JSP page. These elements are called *JSP tags*. Thus, a JSP page is an HTML template made up of intermixed active JSP tags and passive HTML tags. At runtime, the template is used to generate a purely HTML page, which is sent to the client.

2.2 **HELLO USER**

To see the benefits of JSP, let's look at the following example. We have written it three times: first as an HTML page, then as a servlet, and finally as a JSP page. The purpose of the example is to greet the visitors to a web page with the word `Hello`.

2.2.1 The HTML code

Let's start with some simple HTML code, shown in listing 2.1.

Listing 2.1 Hello.html

```
<html>
<body>
<h3>Hello User</h3>
</body>
</html>
```

When accessed with the URL `http://localhost/chapter02/Hello.html`, the code in listing 2.1 prints `Hello User`. However, since HTML is static, it cannot print the user's name. For example, printing either `Hello John` or `Hello Mary` (depending on the user's input) is not possible when using a pure HTML page. It will print the same two words—`Hello User`—regardless of the user.

2.2.2 The servlet code

The `HelloServlet.java` servlet implements this example by modifying the `service()` method. This is shown in listing 2.2.

Listing 2.2 HelloServlet.java

```
import java.io.*;
import javax.servlet.*;
import javax.servlet.http.*;
public class HelloServlet extends HttpServlet
{
    public void service(HttpServletRequest request,
                        HttpServletResponse response)
                        throws ServletException,
                        IOException
    {
        String userName = request.getParameter("userName");

        PrintWriter pw = response.getWriter();
        pw.println("<html>");
        pw.println("<head>");
        pw.println("</head>");
        pw.println("<body>");
        pw.println("<h3>Hello " + userName + "</h3>");
        pw.println("</body>");
        pw.println("</html>");
    }
}
```

When accessed with the URL `http://localhost/chapter02/servlet/HelloServlet?userName=John`, the code in listing 2.2 prints `Hello John`. The user's name is passed to the servlet as part of the URL. The `service()` method sends it back to the browser as part of the generated HTML.

2.2.3 The JSP code

Listing 2.3 contains the JSP code that is equivalent to the previous servlet code.

Listing 2.3 Hello.jsp

```
<html>
<body>
<h3>Hello ${param.userName} </h3>
</body>
</html>
```

When accessed with the URL `http://localhost/chapter02/Hello.jsp?userName=John`, the code in listing 2.3 prints `Hello John`. Again, the user's name is passed to the JSP page as part of the URL.

As you can see from this example, a JSP page contains standard HTML tags. Unlike servlets, it does not involve the explicit writing and compilation of a Java class by the page author. What gives it the power of dynamically generating the greeting is the small amount of JSP code enclosed within the characters `${` and `}`.

2.3 *SERVLET OR JSP?*

Well, if servlets can do whatever JSP pages can, and vice versa, what is the difference between them? And if JSP pages are that easy to write, why bother learning about servlets?

You will recall from the first chapter that servlets are server extensions and provide extra functionality to the main server. This could include implementation of specialized services, such as authentication, authorization, database validation, and transaction management. Servlets act as controller components that control the business logic. They are developed by Java programmers with strong object-oriented programming skills.

On the other hand, JavaServer Pages are web pages. They are similar in structure to HTML pages at design time. Any web page designer who has some knowledge of JSP tags and the basics of Java can write JSP pages.

Web applications typically consist of a combination of servlets and JSP pages. A user-authentication process that accepts login and password information is a good example. The code that generates the HTML FORM, success and error messages, and so forth should be in a JSP page, while the code that accesses the database, validates the password, and authenticates the user should be in a servlet.

Keep these conventions in mind:

- JSP pages are meant for visual presentation.
- Business logic is deferred to servlets.

2.4 JSP ARCHITECTURE MODELS

The JSP tutorials from Sun describe two architectural approaches for building applications using the JSP and servlet technology. These approaches are called JSP Model 1 and JSP Model 2 architectures. The difference between the two lies in the way they handle the requests.

2.4.1 The Model 1 architecture

In Model 1 architecture, the target of every request is a JSP page. This page is completely responsible for doing all the tasks required for fulfilling the request. This includes authenticating the client, using JavaBeans to access the data, managing the state of the user, and so forth. This architecture is illustrated in figure 2.1.

As you can see in figure 2.1, there is no central component that controls the workflow of the application. This architecture is suitable for simple applications. However, it has some serious drawbacks that limit its usage for complex applications. First, it requires embedding business logic using big chunks of Java code into the JSP page. This creates a problem for the web page designers who are usually not comfortable with server-side programming. Second, this approach does not promote reusability of application components. For example, the code written in a JSP page for authenticating a user cannot be reused in other JSP pages.

2.4.2 The Model 2 architecture

This architecture follows the Model-View-Controller (MVC) design pattern (which we will discuss in chapter 18, “Design patterns.”). In this architecture, the targets of all the requests are servlets that act as the controller for the application. They analyze the request and collect the data required to generate a response into JavaBeans objects, which act as the model for the application. Finally, the controller servlets dispatch the request to JSP pages. These pages use the data stored in the JavaBeans to generate a response. Thus, the JSP pages form the view of the application. Figure 2.2 illustrates this architecture.

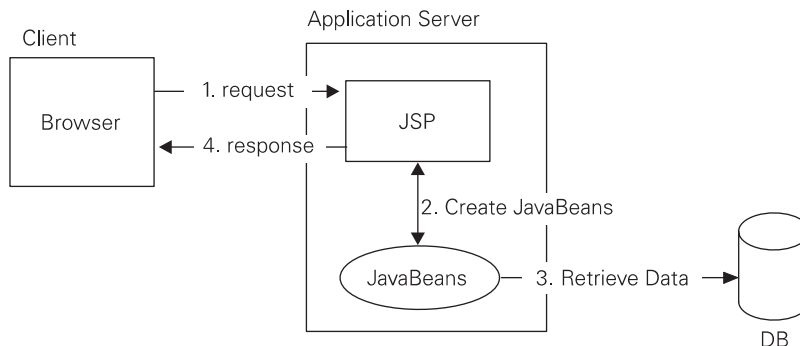


Figure 2.1 The JSP Model 1 architecture

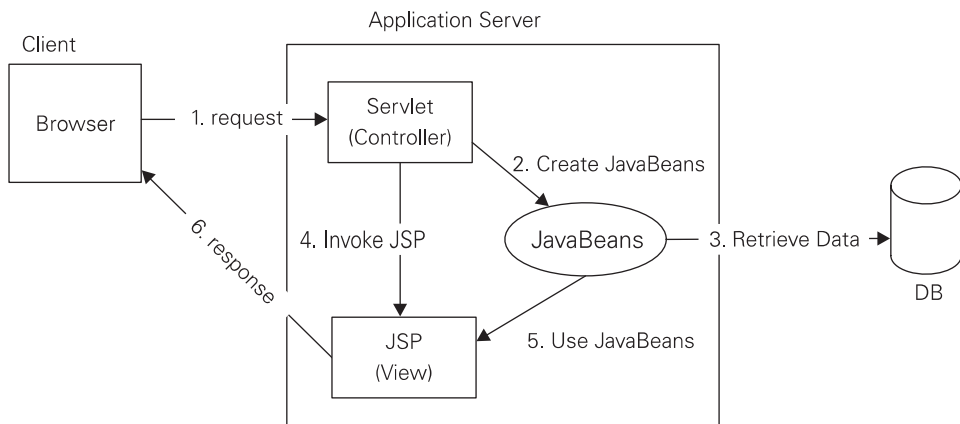


Figure 2.2 The JSP Model 2 architecture

The biggest advantage of this model is the ease of maintenance that results from the separation of responsibilities. The Controller presents a single point of entry into the application, providing a cleaner means of implementing security and state management; these components can be reused as needed. Then, depending on the client's request, the Controller forwards the request to the appropriate presentation component, which in turn replies to the client. This helps the web page designers by letting them work only with the presentation of the data, since the JSP pages do not require any complex business logic. In this way, it satisfactorily solves the problems associated with the Model 1 architecture.

2.5 A NOTE ABOUT JSP SYNTAX

Since this book is specifically meant for the SCWCD exam, its chapters are designed according to the exam objectives specified by Sun. The JSP syntax elements are spread over multiple sections in the exam specification, and therefore, we have spread out the explanations of the elements over several chapters in the book. Table 2.1 contains all of the JSP elements and points out which of them are covered in the exam and which are not. It also documents in which exam objective sections these elements are addressed and where you can find explanations in this book.

Table 2.1 JSP syntax elements

Elements	Exam objective section/subsection	Book section
Directives	6.1	10.1.1
page	6.2	10.3
include	6.2	12.1

continued on next page

Table 2.1 JSP syntax elements (continued)

Elements		Exam objective section/subsection	Book section
	taglib	6.2	15, 16, and 17
Declarations		6.1	10.1.2 and 12.1.1
Scriptlets		6.1	10.1.3 and 12.1.1
	Conditional	6.1	11.1.2
	Iteration	6.1	11.1.2
Expressions		8.1	10.1.4 and 11.1.3
Actions			10.1.5
	jsp:include	8.2	12.2.1
	jsp:forward	8.2	12.2.2
	jsp:useBean	8.1	14.2.1
	jsp:setProperty	8.1	14.2.2
	jsp:getProperty	8.1	14.2.3
	jsp:plugin	NC	10.1.5
Expression Language	implicit variables	7.1	13.1
	operators	7.2, 7.3	13.2
	functions	7.4	13.3
Comments		NC	10.1.6
XML-based syntax		6.3	11.4

NC = Not covered on the exam

2.6 SUMMARY

In this chapter, we learned about the basics of JavaServer Pages technology and server-side includes. We briefly compared JSP pages to servlets and discussed when it is appropriate to use one or the other. We also discussed the two JSP architectural models and how they differ in their request-handling process.



CHAPTER 3

Web application and HTTP basics

- 3.1 What is a web application? 22
- 3.2 Understanding the HTTP protocol 23
- 3.3 Summary 27

INTRODUCTION

In the early years of the Internet, most web sites were constructed entirely of HTML pages. HTML pages are called *static web pages*, since they have all of their content embedded within them and they cannot be modified at execution time. As web technology became more sophisticated, web sites started to incorporate various techniques to create or modify the pages at the time of the user's visit to the site, often in response to the user's input. These are called *dynamic pages*. Today, web sites come in all kinds of styles, and most of them offer at least some type of dynamic features on their pages. The web technologies used to create these dynamic pages include plug-in web components, such as Java Applets or Microsoft ActiveX Controls; programs to build dynamic web pages, such as CGI programs or ASP pages; and n-tier web/distributed systems based on Java Servlets and JavaServer Pages.

3.1 WHAT IS A WEB APPLICATION?

An obvious but still accurate definition of a web application is that it is an application that is accessible from the Web! A common example of a web application is a web site that provides free e-mail service. It offers all the features of an e-mail client such as Outlook Express, but is completely web based. A key benefit of web applications is the ease with which the users can access the applications. All a user needs is a web browser; there is nothing else to be installed on the user's machine. This increases the reach of the applications tremendously while alleviating versioning and upgrading issues.

A *web application* is built of *web components* that perform specific tasks and are able to expose their services over the Web. For example, the `HelloWorldServlet` that we developed in chapter 1 is a *web component*. Since it is complete in itself, it is also a *web application*. In real life, however, a web application consists of multiple servlets, JSP pages, HTML files, image files, and so forth. All of these components coordinate with one another and provide a complete set of services to users.

3.1.1 Active and passive resources

One way of categorizing web resources is that they are either *passive* or *active*. A resource is passive when it does not have any processing of its own; active objects have their own processing capabilities.

For example, when a browser sends a request for `www.myserver.com/myfile.html`, the web server at `myserver.com` looks for the `myfile.html` file, a passive resource, and returns it to the browser. Similarly, when a browser sends a request for `www.myserver.com/reportServlet`, the web server at `myserver.com` forwards the request to `reportServlet`, an active resource. The servlet generates the HTML text on the fly and gives it to the web server. The web server, in turn, forwards it to the browser. A passive resource is also called a static resource, since its contents do not change with requests.

A web application is usually a mixture of active and passive resources, but it is the presence of the active resources that make a web application nearly as interactive as normal applications. Active resources in a web application typically provide dynamic content to users and enable them to execute business logic via their browsers.

3.1.2 Web applications and the web application server

A web application resides in a web application server (or application server). The application server provides the web application with easy and managed access to the resources of the system. It also provides low-level services, such as the HTTP protocol implementation and database connection management. A servlet container is just a part of an application server. In addition to the servlet container, an application server may provide other J2EE components, such as an EJB container, a JNDI server, and a JMS server. You can find detailed information about J2EE and application servers at <http://java.sun.com/j2ee>. Examples of J2EE application servers include BEA Systems' WebLogic, IBM's WebSphere, and Sun's Java System Application Server.

A web application is described using a *deployment descriptor*. A deployment descriptor is an XML document named `web.xml`, and it contains the description of all the dynamic components of the web application. For example, this file has an entry for every servlet used in the web application. It also declares the security aspects of the application. An application server uses the deployment descriptor to initialize the components of the web application and to make them available to the clients.

3.2 UNDERSTANDING THE HTTP PROTOCOL

Simply put, the Hypertext Transfer Protocol is a request-response-based stateless protocol. A client sends an HTTP request for a resource and the server returns an HTTP response with the desired resource, as shown in figure 3.1.

A client opens a connection to the server and sends an HTTP request message. The client receives an HTTP response message sent by the server and closes the connection. It is stateless because once the server sends the response it forgets about the client. In other words, the response to a request does not depend on any previous requests that the client might have made. From the server's point of view, any request is the first request from the client.

In the case of the Internet, the web browser is an HTTP client, the web server is an HTTP server, and the resources are HTML files, image files, servlets, and so forth. Each resource is identified by a unique *Uniform Resource Identifier* (URI). You will frequently hear three terms used interchangeably: URI, URL, and URN. Although they are similar, they have subtle differences:

- *Uniform Resource Identifier*—A URI is a string that identifies any resource. Identifying the resource may not necessarily mean that we can retrieve it. URI is a superset of URL and URN.
- *Uniform Resource Locator*—URIs that specify common Internet protocols such as HTTP, FTP, and mailto are also called URLs. URL is an informal term and is not used in technical specifications.
- *Uniform Resource Name*—A URN is an identifier that uniquely identifies a resource but does not specify how to access the resource. URNs are standardized by official institutions to maintain the uniqueness of a resource.

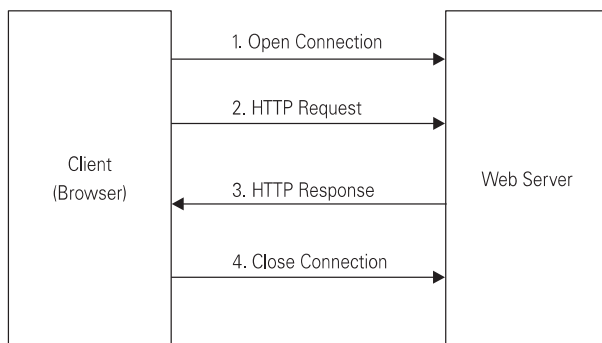


Figure 3.1
HTTP is a request-response-based stateless protocol.

Here are some examples:

- `files/sales/report.html` is a URI, because it identifies some resource. However, it is not a URL because it does not specify how to retrieve the resource. It is not a URN either, because it does not identify the resource uniquely.
- `http://www.manning.com/files/sales/report.html` is a URL because it also specifies how to retrieve the resource.
- `ISBN:1-930110-59-6` is a URN because it uniquely identifies this book, but it is not a URL because it does not indicate how to retrieve the book.

For more details on these terms, visit www.w3c.org.

3.2.1 HTTP basics

An HTTP message is any request from a client to a server, or any response from a server to a client.

The formats of the request and response messages are similar and are in plain English. Table 3.1 lists the parts of an HTTP message.

Table 3.1 The parts of an HTTP message

Message part	Description
The initial line	Specifies the purpose of the request or response message
The header section	Specifies the meta-information, such as size, type, and encoding, about the content of the message
A blank line	
An optional message body	The main content of the request or response message

All the lines end with CRLF—that is, ASCII values 13 (Carriage Return) and 10 (Line Feed).

Let's now look at the individual structures of the request and response messages.

3.2.2 The structure of an HTTP request

An HTTP message sent by a client to a server is called an *HTTP request*. The initial line for an HTTP request has three parts, separated by spaces:

- A method name
- The local path of the requested resource (URI)
- The version of HTTP being used

A typical request line is

```
GET /reports/sales/index.html HTTP/1.1
```

Here, `GET` is the method name, `/report/sales/index.html` is the resource URI, and `HTTP/1.1` is the HTTP version of the request.

The method name specifies the action that the client is requesting the server to perform. HTTP 1.1 requests can have only one of the following three methods: GET, HEAD, or POST. HTTP 1.1 adds five more: PUT, OPTIONS, DELETE, TRACE, and CONNECT.

GET

The HTTP GET method is used to retrieve a resource. It means “*get* the resource identified by this URI.” The resource is usually a passive resource. A GET request may be used for an active resource if there are few or no parameters to be passed. If parameters are required, they are passed by appending a query string to the URI. For example, figure 3.2 illustrates the initial request line for passing `john` as a `userid`.




Figure 3.2 An initial request line using GET and a query string

The part after the question mark is called a *query string*. It consists of parameter name-value pairs separated by an ampersand (&), as in

```
name1=value1&name2=value2&...&nameM=valueM
```

In the example in figure 3.2, `userid` is the parameter name and `john` is the value.

HEAD

An HTTP HEAD request is used to retrieve the meta-information about a resource. Therefore, the response for a HEAD request contains only the header. The structure of a HEAD request is exactly the same as that of a GET request.

HEAD is commonly used to check the time when the resource was last modified on the server before sending it to the client. A HEAD request can save a lot of bandwidth, especially if the resource is very big, since the actual resource would not have to be sent if the client already had the latest version.

POST

A POST request is used to send data to the server in order to be processed. It means “post the data to the active resource identified by this URI.” The block of data is sent in the message body. Usually, to describe this message body, extra lines are present in the header, such as `Content-Type` and `Content-Length`.

HTML pages use POST to submit HTML FORM data. Figure 3.3 shows an example of an HTTP POST request generated by a typical form submission. The value of `Content-Type` is `application/x-www-form-urlencoded`, and the value of `Content-Length` is the length of the URL-encoded form data.



Figure 3.3 A POST request as generated by a form submission

Observe the data line of the request in figure 3.3. In POST, the parameters are sent in the message body, unlike in GET, in which they are a part of the request URI.

PUT

A PUT request is used to add a resource to the server. It means, “*put* the data sent in the message body and associate it with the given Request-URI.” For example, when we PUT a local file named `sample.html` to the server `myhome.com` using the URI `http://www.myhome.com/files/example.html`, the file becomes a resource on that server and is associated with the URI `http://www.myhome.com/files/example.html`. The name of the file (`sample.html`) on the client machine is irrelevant on the server. This request is mainly used to publish files on the server.

NOTE There is a subtle difference between a POST and a PUT request. POST means we are sending some data to a resource for processing. On the other hand, a PUT request means we are sending some data that we want to be associated with a URI.

If you want to learn more about HTTP, read the specification at www.w3.org/Protocols/rfc2616/rfc2616.

3.2.3 The structure of an HTTP response

An HTTP message sent by a server to a client is called an HTTP response. The initial line of an HTTP response is called the status line. It has three parts, separated by spaces: the HTTP version, a response status code that tells the result of the request, and an English phrase describing the status code. HTTP defines many status codes; common ones that you may have noticed are 404 and 500. Here are two examples of a status line that could be sent in the response:

```

HTTP/1.1 404 Not Found
HTTP/1.1 500 Internal Error
  
```

When the browser receives a status code that implies a problem, it displays an appropriate message to the user. If some data is associated with the response, headers like `Content-Type` and `Content-Length` that describe the data may also be present.

A typical HTTP response looks like this:

```
HTTP/1.1 200 OK
Date: Tue, 01 Sep 2004 23:59:59 GMT
Content-Type: text/html
Content-Length: 52
<html>
<body>
  <h1>Hello, John!</h1>
</body>
</html>
```

3.3 SUMMARY

A web application is a collection of web components that perform specific tasks and allow the users to access business logic via their browsers.

In this chapter, we introduced the basic concepts of HTTP, the Hypertext Transfer Protocol. We examined the structure of the HTTP request, including GET, HEAD, POST, and PUT, as well as the structure of the HTTP response.

P A R T

2

Servlets

In the Java world, servlets are the cornerstone of web component technology. In this part of the book, we discuss aspects of the Servlet technology that you need to know, as specified by the exam objectives.



CHAPTER 4

The servlet model

- | | | | | | |
|-----|---|----|------|-------------------------------|----|
| 4.1 | Sending requests: Web browsers and HTTP methods | 32 | 4.5 | Servlet life cycle | 45 |
| 4.2 | Handling HTTP requests in an HttpServlet | 35 | 4.6 | ServletConfig: a closer look | 50 |
| 4.3 | Analyzing the request | 36 | 4.7 | ServletContext: a closer look | 53 |
| 4.4 | Sending the response | 40 | 4.8 | Beyond servlet basics | 54 |
| | | | 4.9 | Summary | 63 |
| | | | 4.10 | Review questions | 63 |

EXAM OBJECTIVES

- 1.1** For each of the HTTP Methods (such as GET, POST, HEAD, and so on):
- Describe the purpose of the method and the technical characteristics of the HTTP Method protocol,
 - List triggers that might cause a Client (usually a Web browser) to use the method; and
 - Identify the HttpServlet method that corresponds to the HTTP Method.
- (Sections 4.1 and 4.2)
- 1.2** Using the HttpServletRequest interface, write code to
- Retrieve HTML form parameters from the request,
 - Retrieve HTML request header information, or
 - Retrieve cookies from the request
- (Section 4.3)
- 1.3** Using the HttpServletResponse interface, write code to
- Set up an HTTP response header,
 - Set the content type of the response,
 - Acquire a text stream for the response,

- Acquire a binary stream for the response,
- Redirect an HTTP request to another URL, or
- Add cookies to the response

(Section 4.4)

1.4 Describe the purpose and event sequence of the servlet life cycle:

- Servlet class loading,
- Servlet instantiation,
- Call the `init()` method,
- Call the service method, and
- Call the `destroy()` method

(Section 4.5)

3.5 Describe the `RequestDispatcher` mechanism;

- Write servlet code to create a request dispatcher,
- Write servlet code to forward or include the target resource, and
- Identify and describe the additional request-scoped attributes provided by the container to the target resource.

(Section 4.8)

INTRODUCTION

Java servlet technology is commonly used to handle the business logic of a web application, although servlets may also contain presentation logic. We discussed the basics of Java servlets in chapter 1. In this chapter, we will take a closer look at the servlet model.

The Servlet specification applies to any protocol, but in practice, most servlets are written for the HTTP protocol, which is why the SCWCD exam focuses on HTTP servlets. In this context, whenever we talk about servlets, we mean `HttpServlet`s. Similarly, by client and server, we mean HTTP client and HTTP server, respectively.

This chapter is lengthy, and while it introduces many concepts about servlets, it will not provide in-depth discussions. Don't worry; at this point, we want you to get familiar with the servlet model without getting lost in the details. We will cover all of these concepts in detail in later chapters.

4.1 SENDING REQUESTS: WEB BROWSERS AND HTTP METHODS

As we discussed in chapter 3, the HTTP protocol consists of requests from the client to the server, and the responses from the server back to the client. Let's look at the request first. A web browser sends an HTTP request to a web server when any of the following events happen:

- A user clicks on a hyperlink displayed in an HTML page.
- A user fills out a form in an HTML page and submits it.
- A user enters a URL in the browser's address field and presses Enter.

Other events trigger a browser to send a request to a web server; for instance, a JavaScript function may call the `reload()` method on the current document. Ultimately, however, all such triggers boil down to one of the three events listed above, because such method calls are nothing but programmatic simulations of the user's actions.

By default, the browser uses the HTTP GET method in all of the above events. However, we can customize the browser's behavior to use different HTTP methods. For example, the following HTML FORM forces the browser to use the HTTP POST method via the method attribute:

```
<FORM name='loginForm' method='POST' action='/loginServlet'>
<input type='text' name='userid'>
<input type='password' name='passwd'>
<input type='submit' name='loginButton' value='Login'>
</FORM>
```

NOTE If you do not specify the method attribute in a `<FORM>` tag, the browser uses GET by default. If you require a POST, you must explicitly specify `METHOD='POST'` in the `<FORM>` tag.

In the following section, we will look at situations in which we might need to force the browser to use POST instead of GET.

4.1.1 Comparing HTTP methods

For the exam, you will be required to demonstrate that you understand both the benefits and functionality of these HTTP methods:

- GET
- POST
- HEAD

We have already seen their basic structure and meaning in chapter 3. Now we will look at the difference between their uses, and identify the situations in which one method is preferred over the other.

Table 4.1 compares the features of GET and POST.

Table 4.1 Comparison of GET and POST methods

Feature	GET method	POST method
Target resource type	Active or passive.	Active.
Type of data	Text.	Text as well as Binary.
Amount of data	Although the HTTP protocol does not limit the length of the query string, some older browsers and web servers may not be able to handle more than 255 characters.	Unlimited.

continued on next page

Table 4.1 Comparison of GET and POST methods (continued)

Feature	GET method	POST method
Visibility	Data is part of the URL and is visible to the user in the URL field of the browser.	Data is not a part of the URL and is sent as the request message body. It is not visible to the user in the URL field of the browser.
Caching	Data can be cached in the browser's URL history.	Data is not cached in the browser's URL history.

Based on table 4.1, we can make some generalizations about when to use each method.

Use GET

- To retrieve an HTML file or an image file, because only the filename needs to be sent.

Use POST

- To send a lot of data; for example, POST is well suited for an online survey, since the length of the query string may exceed 255 characters.
- To upload a file.
- To capture the username and password, because we want to prevent users from seeing the password as a part of the URL.

Recall from chapter 3 that HEAD is the same as GET except that for a HEAD request, the server returns only the response header and not the message body. This makes HEAD more efficient than GET in cases where we need only the response header. For example, a response header contains the modification timestamp, which can be used to determine the staleness of a resource.

In general, clicking on a hyperlink or using the browser's address field causes the browser to send a GET request. We can, of course, attach a JavaScript function, `onClick()`, to programmatically submit a form, thereby causing a POST request to be sent. However, that is not what we are concerned about for the purpose of the exam.

Quizlet

Q: A developer wants to upload a file from the browser to the server. The following is the HTML snippet from the HTML page that she wrote:

```
<FORM name='uploader' action='/saveServlet'
      enctype='multipart/form-data' >
  <input type='file' name='file'>
  <input type='submit' name='uploadButton' value='Upload'>
</FORM>
```

What is wrong with this code snippet?

- A:** The contents of the file must be sent to the server using a POST request. However, the HTML FORM used in this code does not have any method attribute; therefore, a GET request will be sent. The developer must specify the <FORM> tag like this:

```
<FORM name='uploader' action='/saveServlet'  
      enctype='multipart/form-data' method='POST'>
```

4.2 **HANDLING HTTP REQUESTS IN AN HTTPSERVLET**

In the previous section, we discussed three commonly used HTTP methods, their features and limitations, and the situations in which these methods are used. In this section, we will explore what happens when an HTTP request reaches an HTTP servlet.

For every HTTP method, there is a corresponding method in the `HttpServlet` class having the general signature

```
protected void doXXX(HttpServletRequest, HttpServletResponse)  
                throws ServletException, IOException;
```

where `doXXX()` depends on the HTTP method, as shown in table 4.2.

Table 4.2 HTTP methods and the corresponding servlet methods

HTTP method	HttpServlet method
GET	doGet()
HEAD	doHead()
POST	doPost()
PUT	doPut()
DELETE	doDelete()
OPTIONS	doOptions()
TRACE	doTrace()

The `HttpServlet` class provides empty implementations for each of the `doXXX()` methods. We should override the `doXXX()` methods to implement our business logic.

Understanding the sequence of events in HttpServlet

You may now wonder who calls the `doXXX()` methods. Here is the flow of control from the servlet container to the `doXXX()` methods of a servlet:

- 1 The servlet container calls the `service(ServletRequest, ServletResponse)` method of `HttpServlet`.
- 2 The `service(ServletRequest, ServletResponse)` method of `HttpServlet` calls the `service(HttpServletRequest, HttpServletResponse)` method of the same class. Observe that the `service` method is overloaded in the `HttpServlet` class.

- 3 The `service(HttpServletRequest, HttpServletResponse)` method of `HttpServlet` analyzes the request and finds out which HTTP method is being used. Depending on the HTTP method, it calls the corresponding `doXXX()` method of the servlet. For example, if the request uses the POST method, it calls the `doPost()` method of the servlet.

NOTE If you override the service methods in your servlet class, you will lose the functionality provided by the `HttpServlet` class, and the `doXXX()` methods will not be called automatically. In your implementation, you will have to determine the HTTP method used in the request, and then you will have to call the appropriate `doXXX()` method yourself. For this reason, it's recommended to only override the `doPost()` or `doGet()` methods.

All of the `doXXX()` methods take two parameters: an `HttpServletRequest` object and an `HttpServletResponse` object. We will learn about these objects in the following sections.

But first, here's a note about the Servlet API: Most of the important components of the Servlet API, including `HttpServletRequest` and `HttpServletResponse`, are interfaces. The servlet container provides the classes that implement these interfaces. So, whenever we say something like "an `HttpServletRequest` object," we mean "an object of a class that implements the `HttpServletRequest` interface." The name of the actual class is not significant and is, in fact, unknown to the developer.

Quizlet

Q: Which method of `TestServlet` will be called when a user clicks on the following URL?

```
<a href="/servlet/TestServlet" method="POST">Test URL</a>
```

A: The `method="POST"` attribute-value pair does not make sense in the `<a href>` tag. Clicking on a hyperlink always sends a GET request and thus, the `doGet()` method of the servlet will be called.

4.3 ANALYZING THE REQUEST

Both `ServletRequest` and its subclass, `HttpServletRequest`, allow us to analyze a request. They provide us with a view of the data sent by the browser. The data includes parameters, meta-information, and a text or binary data stream.

`ServletRequest` provides methods that are relevant to any protocol, while `HttpServletRequest` extends `ServletRequest` and adds methods specific to HTTP. It is for this reason that the `ServletRequest` interface belongs to the `javax.servlet` package and the `HttpServletRequest` interface belongs to the `javax.servlet.http` package.

We always use the `HttpServletRequest` class, but it is important to know which methods are implemented in the `HttpServletRequest` class and which methods are inherited from the `ServletRequest` class.

4.3.1 Understanding ServletRequest

The primary use of `ServletRequest` is to retrieve the parameters sent by a client. Table 4.3 describes the methods provided to retrieve the parameters.

Table 4.3 Methods provided by `ServletRequest` for retrieving client-sent parameters

Method	Description
<code>String getParameter (String paramName)</code>	This method returns just one of the values associated with the given parameter.
<code>String[] getParameterValues (String paramName)</code>	This method returns all the values associated with the parameter. For example, while doing a job search, you might have seen a “location” list box that allows you to select multiple states. In this case, the parameter “location” may have multiple values.
<code>Enumeration getParameterNames ()</code>	This method is useful when you don’t know the names of the parameters. You can iterate through the Enumeration of Strings returned by this method and for each element you can call <code>getParameter ()</code> or <code>getParameterValues ()</code> .

4.3.2 Understanding HttpServletRequest

The class that implements the `HttpServletRequest` interface implements all of the methods of `ServletRequest` in an HTTP-specific manner. It parses and interprets HTTP messages and provides the relevant information to the servlet.

Let’s look at an example of how we can use these methods. Figure 4.1 shows an HTML page that allows a user to send two parameters to the server.



Figure 4.1
An HTML page
containing a FORM

Listing 4.1 is the HTML code for this page.

Listing 4.1 HTML page snippet

```
<form action="/servlet/TestServlet" method="POST"> ← Uses HTTP POST
Technology : <input type="text" name="searchstring" value="java">
<br><br>
State : <select name="state" size="5" multiple> ← Allows selection of
      <option value="NJ">New Jersey</option>           multiple values
      <option value="NY">New York</option>
      <option value="KS">Kansas</option>
      <option value="CA">California</option>
      <option value="TX">Texas</option>
</select>
<br><br>
<input type="submit" value="Search Job">
</form>
```

This FORM displays a text field, a list box, and a submit button. The action attribute of the FORM specifies that `TestServlet` should handle the request. Observe that the method attribute of the FORM is set to `POST`, and so the parameters will be sent to the server using an HTTP `POST` request.

Once the request is sent to the server, `TestServlet` is invoked. Listing 4.2 shows how `TestServlet`'s `doPost()` method retrieves the parameters that were sent by submitting the form (listing 4.1).

Listing 4.2 Code for `doPost` to retrieve parameter values

```
public void doPost(HttpServletRequest req,
                  HttpServletResponse res)
{
    String searchString = req.getParameter("searchstring"); ← Retrieves
                                                            searchstring
                                                            parameter
                                                            value
    String[] stateList = req.getParameterValues("state"); ← Retrieves all
                                                            the values
                                                            selected in the
                                                            state list
    //use the values and generate appropriate response
}
```

In the above code, we know the names of the parameters (`searchstring` and `state`) sent with the request, so we can use the `getParameter()` and `getParameterValues()` methods to retrieve the parameter values. When the parameter values are not known, we can use `getParameterNames()` to retrieve the parameter names.

Retrieving request headers

Just as there are methods to retrieve request parameters, there are methods to retrieve names and values from request headers. We'd like to point out one

difference, though; unlike parameters, headers are specific to the HTTP protocol and so the methods that deal with the headers belong to `HttpServletRequest` and not to `ServletRequest`.

`HttpServletRequest` provides the methods shown in table 4.4 to help us retrieve the header information.

Table 4.4 `HttpServletRequest` methods for managing request headers

Method	Description
<code>String getHeader (String headerName)</code>	This method returns just one of the values associated with the given header.
<code>Enumeration getHeaders (String headerName)</code>	This method returns all the values associated with the header as an Enumeration of String objects.
<code>Enumeration getHeaderNames ()</code>	This method is useful when you don't know the names of the headers. You can iterate through the enumeration returned by this method, and for each element you can call <code>getHeader ()</code> or <code>getHeaders ()</code> .

Let's see how we can use the methods described in table 4.4. The `service ()` method code shown in listing 4.3 prints out all the headers present in a request.

Listing 4.3 Printing out all the headers on the console

```
public void service(HttpServletRequest req,
                    HttpServletResponse res)
{
    Enumeration headers = req.getHeaderNames();    ← Retrieves header names

    while (headers.hasMoreElements())
    {
        String header = (String) headers.nextElement();
        String value = req.getHeader(header);    ← Retrieves header values
        System.out.println(header+" = "+value);
    }
}
```

Retrieving cookies from the request

The final `HttpServletRequest` method that you'll need to know for the exam is `getCookies ()`. This method returns an array of `Cookie` objects that provide information about the current client/server interaction called a *session*. We'll discuss these objects and their uses in chapter 8.

There are other convenience methods in `ServletRequest` and in `HttpServletRequest` that we will not discuss here, since they are not required for the exam. To learn more about them, refer to the Servlet API documentation.

Quizlet

- Q:** Which method would you use to retrieve the number of parameters present in a request?
- A:** Neither `ServletRequest` nor `HttpServletRequest` provides any method to retrieve the number of parameters directly. You'll have to use `ServletRequest.getParameterNames()`, which returns an `Enumeration`, and count the number of parameters yourself.

4.4 SENDING THE RESPONSE

The `HttpServletResponse` object is a servlet's gateway to send information back to the browser. It accepts the data that the servlet wants to send to the client and formats it into an HTTP message as per the HTTP specification.

`ServletResponse` provides methods that are relevant to any protocol, while `HttpServletResponse` extends `ServletResponse` and adds HTTP-specific methods. Not surprisingly, the `ServletResponse` interface belongs to the `javax.servlet` package and the `HttpServletResponse` interface belongs to the `javax.servlet.http` package.

4.4.1 Understanding ServletResponse

`ServletResponse` declares several generic methods, including `getWriter()`, `getOutputStream()`, `setContentType()`, and so forth. We will now discuss two important methods.

Using PrintWriter

Let's first look at the `getWriter()` method of `ServletResponse`. This method returns an object of class `java.io.PrintWriter` that can be used to send character data to the client. `PrintWriter` is extensively used by servlets to generate HTML pages dynamically. Listing 4.4 demonstrates its use by sending the header information of a request to the browser.

Listing 4.4 Writing HTML code dynamically

```
import java.io.*;
import java.util.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class ShowHeadersServlet extends HttpServlet
{
    public void doGet(HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException
    {
        PrintWriter pw = res.getWriter(); ← Gets the PrintWriter object
```

```

pw.println("<html>");
pw.println("<head>");
pw.println("</head>");
pw.println("<body>");

pw.println("<h3>Following are the headers that the
           server received.</h3><p>");

Enumeration headers = req.getHeaderNames();
while(headers.hasMoreElements())
{
    String header = (String) headers.nextElement();
    String value = req.getHeader(header);

    pw.println(header+" = "+value+"<br>");
}

pw.println("</body>");
pw.println("</html>");
}
}

```

Uses PrintWriter to write the HTML page

In listing 4.4, we use the `getWriter()` method to retrieve the `PrintWriter` object. We use the `getHeaderNames()` and `getHeader()` methods to retrieve header information, and then we write the values using the `PrintWriter` object. We then finish our dynamic HTML page with a closing `</body>` and `</html>` tag.

Using ServletOutputStream

If we want to send a binary file, for example a JAR file, to the client, we will need an `OutputStream` instead of a `PrintWriter`. `ServletResponse` provides the `getOutputStream()` method that returns an object of class `javax.servlet.ServletOutputStream`. In listing 4.5, we have changed the `doGet()` method of the previous example to send a JAR file to the browser.

Listing 4.5 Sending a JAR file to the browser

```

public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
    throws ServletException, IOException
{
    res.setContentType("application/jar");    ← Sets the content type

    File f = new File("test.jar");
    byte[] bytearray = new byte[(int) f.length()];
    FileInputStream is = new FileInputStream(f);
    is.read(bytearray);

    OutputStream os = res.getOutputStream();    ← Gets the OutputStream

```

Reads the file into a byte array

```

os.write(bytearray);    ← Sends the bytes of the byte array to the browser

os.flush();            ← Flushes the data
}

```

In listing 4.5, we retrieve the `OutputStream` using the `getOutputStream()` method. We simply read the contents of a JAR file into a byte array and write the byte array to the `OutputStream`.

Observe that we are calling the `setContentType()` method before calling the `getOutputStream()` method. The `setContentType()` method allows us to specify the MIME type of the data we are sending in the response. The content type can also include the character encoding used in the response. If we need to obtain a `PrintWriter` with a nondefault content type or character encoding, we must call `setContentType()` before calling `getWriter()`. The `setContentType()` method belongs to the `ServletResponse` interface and is declared as shown in table 4.5.

Table 4.5 The method provided by `ServletResponse` for setting the content type of the response

Method	Description
<code>public void setContentType (String type)</code>	This method is used to set the content type of the response. The content type may include the type of character encoding used, for example, <code>text/html; charset=ISO-8859-4</code> . If obtaining a <code>PrintWriter</code> , this method should be called first for the charset to take effect. If this method is not called, the content type is assumed to be <code>text/html</code> . The following are some commonly used values for the content type: <code>text/html</code> , <code>image/jpeg</code> , <code>video/quicktime</code> , <code>application/java</code> , <code>text/css</code> , and <code>text/javascript</code> .

You might have also noticed that in the line `File f = new File ("test.jar")`; we are hard-coding the filename to `test.jar`. This requires the file `test.jar` to be in the `bin` directory of Tomcat. We will come back to this later in section 4.7 to learn a better way of specifying the file.

NOTE An important point to note about the `getWriter()` and `getOutputStream()` methods is that you can call only one of them on an instance of `ServletResponse`. For example, if you have already called the `getWriter()` method on a `ServletResponse` object, you cannot call the `getOutputStream()` method on the same `ServletResponse` object. If you do, the `getOutputStream()` method will throw an `IllegalStateException`. You can call the same method multiple times, though.

4.4.2 Understanding HttpServletResponse

In addition to setting the content type of the response, there are three other important capabilities of `HttpServletResponse`s that you need to know for the exam: setting response header information, redirecting HTTP requests to another URL, and adding cookies to the response.

Setting the response headers

We use headers to convey additional information about the response by setting name-value pairs. For example, we can use a header to tell the browser to reload the page it is displaying every 5 minutes, or to specify how long the browser can cache the page. As shown in table 4.6, the `HttpServletResponse` interface provides seven methods for header management.

Table 4.6 `HttpServletResponse` methods for managing response headers

Method	Description
<code>void setHeader (String name, String value)</code>	Used to set the name-value pair for a header in the <code>ServletRequest</code> .
<code>void setIntHeader (String name, int value)</code>	Saves you from converting the int value to string.
<code>void setDateHeader (String name, long millisecs)</code>	Pretty much the same as above.
<code>void addHeader/addIntHeader/addDateHeader</code>	These methods can be used to associate multiple values with the same header.
<code>boolean containsHeader (String name)</code>	Returns a Boolean that tells you whether or not a header with this name is already set.

Table 4.7 shows four important header names. Although on the exam you will not be asked questions based on header names and values, it is good to know some commonly used headers. For a complete list of header names-values, refer to the HTTP specification.

Table 4.7 Typical response header names and their uses

Header name	Description
Date	Specifies the current time at the server.
Expires	Specifies the time when the content can be considered stale.
Last-Modified	Specifies the time when the document was last modified.
Refresh	Tells the browser to reload the page.

Another useful method related to headers is `addCookie (Cookie c)`. This method lets us create `Cookie` objects and set them in the response. We will learn about cookies in chapter 8, “Session management.”

Redirecting the request

After analyzing a request, a servlet may decide that it needs to redirect the browser to another resource. For example, a company web site may be able to provide only company news. For all other kind of news, it may redirect the browser to another web site. The `HttpServletResponse` class provides the `sendRedirect()` method exactly for this purpose, as shown here:

```
if ("companynews".equals(request.getParameter("news_category")))
{
    //retrieve internal company news and generate
    //the page dynamically
}
else
{
    response.sendRedirect("http://www.cnn.com");
}
```

The above code checks the `news_category` parameter to decide whether it should generate a reply on its own or redirect the browser to `cnn.com`. When the browser receives the redirect message, it automatically goes to the given URL.

We should keep in mind a couple of important points about the `sendRedirect()` method. We cannot call this method if the response is committed—that is, if the response header has already been sent to the browser. If we do, the method will throw a `java.lang.IllegalStateException`. For example, the following code will generate an `IllegalStateException`:

```
public void doGet(HttpServletRequest req, HttpServletResponse res)
{
    PrintWriter pw = res.getWriter();
    pw.println("<html><body>Hello World!</body></html>");
    pw.flush();    ← Sends the response
    res.sendRedirect("http://www.cnn.com");    ← Tries to redirect
}
```

In this code, we are forcing the servlet container to send the header and the generated text to the browser immediately by calling `pw.flush()`. The response is said to be *committed* at this point. Calling `sendRedirect()` after committing the response causes the servlet container to throw an `IllegalStateException`.

NOTE Another important point to understand about `sendRedirect()` is that the browser goes to the second resource only after it receives the redirect message from the first resource. In that sense, `sendRedirect()` is not transparent to the browser. In other words, the servlet sends a message telling the browser to get the resource from elsewhere.

Sending status codes for error conditions

HTTP defines status codes for common error conditions such as *Resource not found*, *Resource moved permanently*, and *Unauthorized access*. All such codes are defined in the `HttpServletResponse` interface as constants. `HttpServletResponse` also provides `sendError(int status_code)` and `sendError(int status_code, String message)` methods that send a status code to the client. For example, if a servlet finds out that the client should not have access to its output, it may call

```
response.sendError(HttpServletResponse.SC_UNAUTHORIZED);
```

When the browser receives this status code, it displays an appropriate message to the user.

For a complete list of status codes, refer to the API documentation on `HttpServletResponse`.

Quizlet

Q: Which methods should you use to do the following?

- 1 Write HTML tags to the output.
- 2 Specify that the content of the response is a binary file.
- 3 Send a binary file to the browser.
- 4 Add a header to a response.
- 5 Redirect a browser to another resource.

A:

- 1 First, get the `PrintWriter` using `ServletResponse.getWriter()` and then call `PrintWriter.print("<html tags>");`
- 2 Use `ServletResponse.setContentType(String content-type);`
- 3 Use `ServletResponse.getOutputStream();` and then `OutputStream.write(bytes);`
- 4 Use `HttpServletResponse.setHeader("name", "value");`
- 5 Use `HttpServletResponse.sendRedirect(String url-string);`

4.5 SERVLET LIFE CYCLE

By now, it should be very clear that a servlet receives a request, processes it, and sends a response back using the `doXXX()` methods. There is, however, a little bit more to understand than just the `doXXX()` methods. Before a servlet can service the client requests, a servlet container must take certain steps in order to bring the servlet to a state in which it is ready to service the requests. The first step is loading and instantiating the servlet class; the servlet is now considered to be in the loaded state. The second step is initializing the servlet instance. Once the servlet is in the initialized state,

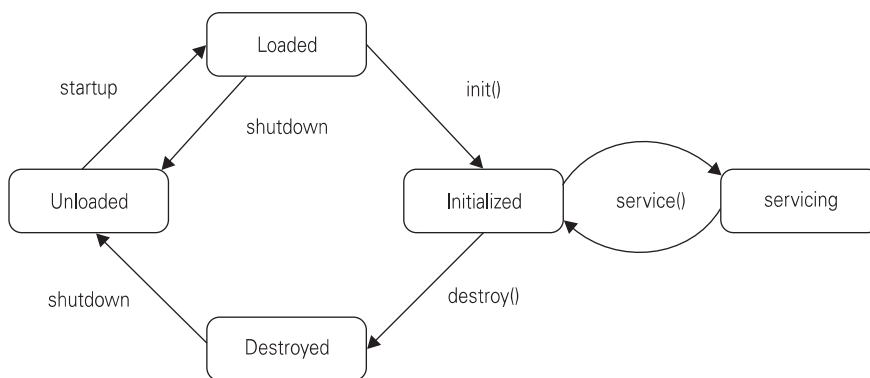


Figure 4.2 Servlet state transition diagram

the container can invoke its `service()` method whenever it receives a request from the client. There may be times when the container will call the `destroy()` method on the servlet instance to put it in the destroyed state. Finally, when the servlet container shuts down, it must unload the servlet instance. Figure 4.2 shows these servlet states and their transitions.

These states constitute the life cycle of a servlet. Let's take a closer look at them.

4.5.1 Loading and instantiating a servlet

When we start up a servlet container, it looks for a set of configuration files, also called the deployment descriptors, that describe all the web applications. Each web application has its own deployment descriptor file, `web.xml`, which includes an entry for each of the servlets it uses. An entry specifies the name of the servlet and a Java class name for the servlet. The servlet container creates an instance of the given servlet class using the method `Class.forName(className).newInstance()`. However, to do this the servlet class must have a public constructor with no arguments. Typically, we do not define any constructor in the servlet class. We let the Java compiler add the default constructor. At this time, the servlet is *loaded*.

4.5.2 Initializing a servlet

It is entirely possible that we will want to initialize the servlet with some data when it is instantiated. How can we do that if we do not define a constructor? Good question. It is exactly for this reason that once the container creates the servlet instance, it calls the `init(ServletConfig)` method on this newly created instance. The `ServletConfig` object contains all the initialization parameters that we specify in the deployment descriptor of the web application. We will see how these parameters can be specified in the `web.xml` file in section 4.6.2. The servlet is *initialized* after the `init()` method returns.

This process of initializing a servlet using the initialization parameters from the `ServletConfig` object is quite important in order to ensure the reusability of a

servlet. For example, if we wanted to create a database connection in the servlet we would not want to hard-code the username/password and the database URL in the servlet. The `init()` method allows us to specify them in the deployment descriptor. The values in the deployment descriptor can be changed as needed without affecting the servlet code. When the servlet initializes, it can read the values in its `init()` method and make the connection.

It does not make sense to initialize an object repeatedly; therefore, the framework guarantees that the servlet container will call the `init()` method once and only once on a servlet instance.

NOTE If you look up the API for the `GenericServlet` class, you will see that it has two `init()` methods: one with a parameter of type `ServletConfig` as required by the `Servlet` interface and one with no parameters. The no parameter `init()` method is a convenience method that you can override in your servlet class. If you override the `init(ServletConfig config)` method, you will have to include a call to `super.init(config)` in the method so that the `GenericServlet` can store a reference to the config object for future use. To save you from doing that, the `GenericServlet`'s `init(ServletConfig)` method makes a call to the `GenericServlet`'s no parameter `init()` method, which you can implement freely. To get the `ServletConfig` object in the no parameter `init()` method, you can call the `getServletConfig()` method implemented by the `GenericServlet` class.

Preinitializing a servlet

Usually, a servlet container does not initialize the servlets as soon as it starts up. It initializes a servlet when it receives a request for that servlet for the first time. This is called *lazy loading*. Although this process greatly improves the startup time of the servlet container, it has a drawback. If the servlet performs many tasks at the time of initialization, such as caching static data from a database on initialization, the client that sends the first request will have a poor response time. In many cases, this is unacceptable. The servlet specification defines the `<load-on-startup>` element, which can be specified in the deployment descriptor to make the servlet container load and initialize the servlet as soon as it starts up. This process of loading a servlet before any request comes in is called *preloading*, or *preinitializing*, a servlet.

4.5.3 Servicing client requests

After the servlet instance is properly initialized, it is ready to service client requests. When the servlet container receives requests for this servlet, it will dispatch them to the servlet instance by calling the `Servlet.service(ServletRequest, ServletResponse)` method.

4.5.4 Destroying a servlet

If the servlet container decides that it no longer needs a servlet instance, it calls the `destroy()` method on the servlet instance. In this method, the servlet should clean up the resources, such as database connections that it acquired in the `init()` method. Once this method is called, the servlet instance will be out of service and the container will never call the `service()` method on this instance. The servlet container cannot reuse this instance in any way. From this state, a servlet instance may only go to the unloaded state. Before calling the `destroy()` method, the servlet container waits for the remaining threads that are executing the servlet's `service()` method to finish.

A servlet container may destroy a servlet if it is running low on resources and no request has arrived for a servlet in a long time. Similarly, if the servlet container maintains a pool of servlet instances, it may create and destroy the instances from time to time as required. A servlet container may also destroy a servlet if it is shutting down.

4.5.5 Unloading a servlet

Once destroyed, the servlet instance may be garbage collected, in which case the servlet instance is said to be *unloaded*. If the servlet has been destroyed because the servlet container is shutting down, the servlet class will also be unloaded.

4.5.6 Servlet state transition from the servlet container's perspective

Figure 4.3 illustrates the relationship between the servlet container and the life-cycle phases of a servlet.

A servlet goes from the unloaded to the loaded state when a servlet container loads the servlet class and instantiates an object of the class. The servlet container initializes the servlet object by calling its `init()` method, thereby putting it in the initialized state.

The servlet stays in the initialized state until the servlet container decides to destroy the servlet. From the initialized state, the servlet enters the servicing state whenever the servlet container calls its `service()` method in order to process client requests.

The servlet enters the destroyed state when the servlet container calls its `destroy()` method. Finally, the servlet goes to the unloaded state when the servlet instance is garbage collected.

Table 4.8 summarizes all the servlet life-cycle methods.

Table 4.8 Servlet life-cycle methods as defined in servlet interface

Method	Description
<code>void init(ServletConfig)</code>	The servlet container calls this method to initialize the servlet.
<code>void service(ServletRequest, ServletResponse)</code>	The servlet container calls this method to service client requests.
<code>void destroy()</code>	The servlet container calls this method when it decides to unload the servlet.

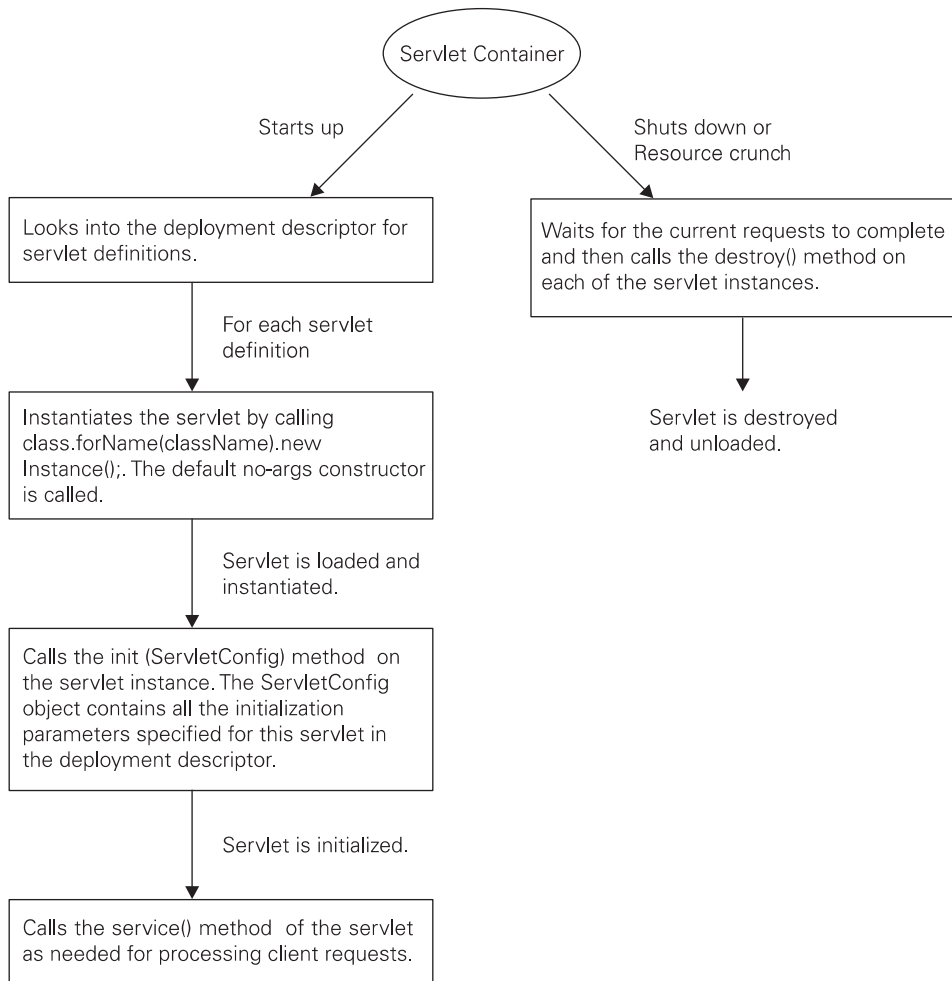


Figure 4.3 The servlet life cycle from the servlet container's perspective

NOTE A servlet container calls the `init (ServletConfig)` method on a servlet object only once. However, it is possible that it will create multiple servlet objects of the same servlet class if more than one `<servlet>` element is defined in the `web.xml` file having the same servlet class names. You can do this if you want to have multiple sets of initialization parameters. For example, you may want one instance to connect to one database and a second instance to connect to another database.

Quizlet

Q: Which method does the servlet container call on a servlet to initialize the servlet?

- A:** The `init (javax.servlet.ServletConfig)` method of the `javax.servlet.Servlet` interface. The `javax.servlet.GenericServlet` class implements this method.

4.6 **SERVLETCONFIG: A CLOSER LOOK**

We learned in the previous section that the servlet container passes a `ServletConfig` object in the `init (ServletConfig)` method of the servlet. In this section, we will look at the details of `ServletConfig` that you need to understand for the exam.

4.6.1 **ServletConfig methods**

The `ServletConfig` interface is defined in the `javax.servlet` package and is rather simple to use. It provides four methods, as shown in table 4.9.

Table 4.9 `ServletConfig` methods for retrieving initialization parameters

Method	Description
<code>String getInitParameter (String name)</code>	Returns the value of the parameter or null if no such parameter is available.
<code>Enumeration getInitParameterNames ()</code>	Returns an Enumeration of Strings for all the parameter names.
<code>ServletContext getServletContext ()</code>	Returns the <code>ServletContext</code> for this servlet.
<code>String getServletName ()</code>	Returns the servlet name as specified in the configuration file.

Notice that `ServletConfig` provides methods only to retrieve parameters. You cannot add or set parameters to the `ServletConfig` object.

A servlet container takes the information specified about a servlet in the deployment descriptor and wraps it into a `ServletConfig` object. This information can then be retrieved by the servlet at the time of its initialization.

4.6.2 **Example: a servlet and its deployment descriptor**

To really understand the `ServletConfig` methods, you first need to understand how to specify the initialization parameters in the deployment descriptor. The `web.xml` file shown in listing 4.6 declares a servlet and specifies four initialization parameters for the servlet. Later, we will use these parameters in our servlet to make a connection to the database.

Listing 4.6 The `web.xml` file specifying init parameters

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE web-app
    PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
    "http://java.sun.com/j2ee/dtds/web-app_2_3.dtd">
```

```

<web-app>
  <servlet>  ← Defines a servlet                Defines a parameter
  <servlet-name>TestServlet</servlet-name>          and specifies its
  <servlet-class>TestServlet</servlet-class>        name and value
  <init-param>
    <param-name>driverclassname</param-name>
    <param-value>sun.jdbc.odbc.JdbcOdbcDriver</param-value>
  </init-param>
  <init-param>
    <param-name>dburl</param-name>
    <param-value>jdbc:odbc:MySQLODBC</param-value>
  </init-param>
  <init-param>
    <param-name>username</param-name>
    <param-value>testuser</param-value>
  </init-param>
  <init-param>
    <param-name>password</param-name>
    <param-value>test</param-value>
  </init-param>
  <load-on-startup>1</load-on-startup>
</servlet>
</web-app>

```

We will discuss the complete structure of a deployment descriptor in chapter 5, but for now, just note that the above listing has one `<servlet>` element, which defines a servlet named `TestServlet`. The `<servlet>` element has four `<init-param>` elements, which define four parameters: `driverclassname`, `dburl`, `username`, and `password`. Notice the `<load-on-startup>` element, which ensures that this servlet will be loaded as soon as the container starts up.

Now we are ready to examine the `ServletConfig` methods in action. Listing 4.7 shows the complete code for the `TestServlet` servlet, which uses the initialization parameters we defined in the deployment descriptor (listing 4.6) to connect to a database.

Listing 4.7 TestServlet.java; making use of init parameters

```

import java.io.*;
import java.util.*;
import java.sql.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class TestServlet extends HttpServlet
{
    Connection dbConnection;

```

```

public void init()    ← Creates connection in init()
{
    System.out.println(getServletName()+" : Initializing...");

    ServletConfig config = getServletConfig();

    String driverClassName =
        config.getInitParameter("driverclassname");

    String dbURL = config.getInitParameter("dburl");
    String username = config.getInitParameter("username");
    String password = config.getInitParameter("password");

    //Load the driver class
    Class.forName(driverClassName);

    //get a database connection
    dbConnection =
        DriverManager.getConnection(dbURL,username,password);

    System.out.println("Initialized.");
}

public void service(HttpServletRequest req,
                    HttpServletResponse res)
                    throws ServletException, java.io.IOException
{
    //get the requested data from the database and
    //generate an HTML page.
}

public void destroy() ← Cleans up the resources
{
    try
    {
        dbConnection.close();
    }
    catch(Exception e)
    {
        e.printStackTrace();
    }
}
}

```

**Retrieves
parameters**

The above servlet reads the initialization parameters specified in the deployment descriptor and makes a connection to the database in the `init()` method. It also uses the `getServletName()` method to print a debug statement on the console. The `getServletName()` method returns the name of the servlet as defined in the `<servlet-name>` element of the deployment descriptor. Observe the use of the `destroy()` method to close the connection.

4.7 **SERVLETCONTEXT: A CLOSER LOOK**

We can think of the `ServletContext` interface as a window for a servlet to view its environment. A servlet can use this interface to get information, such as initialization parameters for the web application or the servlet container's version. This interface also provides utility methods for retrieving the Multipurpose Internet Mail Extensions (MIME) type for a file, for retrieving shared resources (such as property files), for logging, and so forth. Every web application has one and only one `ServletContext`, and it is accessible to all the active resources of that application. It is also used by the servlets to share data with one another.

It is important to have a thorough understanding of `ServletContext`, since the exam contains many questions about this interface. This book is organized according to the exam objectives, and we will discuss the various `ServletContext` methods in different chapters as they apply to those objectives.

In this section, we will work with the `getResource()` and `getResourceAsStream()` methods. These methods are used by a servlet to access any resource without worrying about where the resource actually resides.

For a detailed description of these methods, you should refer to the API documentation. Table 4.10 provides a brief description of each.

Table 4.10 `ServletContext` methods for retrieving a resource

Method	Description
<code>java.net.URL getResource (String path)</code>	This method returns a <code>java.net.URL</code> object for the resource that is mapped to the given path. Although the path should start with <code>/</code> it is not an absolute path. It is relative to the document root of this web application. For instance, if you pass a path to a JSP file, it will give you the unprocessed data, i.e., in this case, the JSP source code, when you read the contents.
<code>java.io.InputStream getResourceAsStream (String path)</code>	This is a shortcut method if you just want to get an <code>InputStream</code> out of the resource. It is equivalent to <code>getResource(path).openStream()</code> .

If you recall from listing 4.5, in which we sent a JAR file to the browser, we had hard-coded the JAR filename in the line:

```
File f = new File("test.jar");
```

The code in listing 4.8 uses the `getResource()` method to specify the filename independently of the file system.

Listing 4.8 Making use of `ServletContext.getResource()`

```
public void service(HttpServletRequest req,
                    HttpServletResponse res)
    throws ServletException,
           java.io.IOException
{
```



```

res.setContentType("application/jar");

OutputStream os = res.getOutputStream();

//1K buffer
byte[] bytearray = new byte[1024];

ServletContext context = getServletContext();
URL url = context.getResource("/files/test.jar");
InputStream is = url.openStream();

int bytesread = 0;
while( (bytesread = is.read(bytearray) ) != -1 )
{
    os.write(bytearray, 0, bytesread);
}
os.flush();
is.close();
}

```

← Returns a URL object to the file

In listing 4.8, we have specified a relative path to the file `test.jar`. This allows us to deploy this servlet anywhere without worrying about the absolute location of the file. As long as `test.jar` is available under the `<webappdirectory>\files` directory, it can be found.

Here are the limitations of the `getResource()` and `getResourceAsStream()` methods:

- You cannot pass a URL of any active resource—for example, a JSP page or servlet—to this method.
- If used improperly, this method can become a security hole; it can read all of the files that belong to this web application, including the files under the `WEB-INF` directory of this web application.

A servlet can, of course, access a resource directly by converting a relative path to an absolute path using the `getRealPath(String relativePath)` method of `ServletContext`. However, the problem with this approach is that it is not helpful when the resource is inside a JAR file. It is also useless when the servlet is running in a distributed environment where the resource may reside on a different machine. In such situations, the `getResource()` method comes handy.

4.8 BEYOND SERVLET BASICS

Until now, we have been discussing servlets from the point of view of just one servlet. But in the real world, having just one servlet to do all the tasks is not practical. Typically, we divide the business process into multiple tasks. For example, consider a grossly simplified business process of a bank. A user should be able to

- Open an account
- View the account balance
- Make deposits
- Make withdrawals
- Close the account

Besides these activities, many other business rules need to be addressed as well; for example, a user should not be able to view or withdraw from anyone else's account.

We usually break up the whole business process into different tasks and have one servlet focus on one task. In the example described above, we can have a `LoginServlet` that allows a user to sign up and log in/out, and an `AccountServlet` that allows users to view their account balance and deposit or withdraw money.

To implement the required functionality, the servlets will have to coordinate their processes and share the data. For example, if a user directly tries to access an account, `AccountServlet` should be able to determine the user's login status, and it should redirect the user to the login page if he is not logged in. On the other hand, once a user logs in, `LoginServlet` should be able to share the `userid` with `AccountServlet` so that `AccountServlet` can display the status of the account without asking for the `userid` again.

The Servlet API provides elegant ways to share data and to coordinate the servlet processes. We will discuss these ways in the following sections.

4.8.1 Sharing the data (attribute scopes)

Data is shared between the servlets using the rendezvous concept. One servlet puts the data in a well-known place, which acts as a container, and other servlets access the data from that place. These well-known containers are the `ServletRequest` object, the `HttpSession` object, and the `ServletContext` object. All three objects provide a `setAttribute(String name, Object value)` method (to put the data in the container) and an `Object getAttribute(String name)` method (to access the data).

Although data can be shared using any of these containers, there is a difference in the visibility of the data present in these containers. Simply put, objects shared using `ServletRequest` are accessible only for the life of a request, objects shared using `HttpSession` are accessible only for the life of the session, and objects shared using `ServletContext` are accessible for the life of the web application. To understand this difference clearly, consider the following situations where we have different requirements for sharing the data:

- The banking application that we described earlier needs to provide credit reports for the users. So, we add a `ReporterServlet` which, given a social security number (SSN), can generate a credit report for any user. When a user asks for his credit report, `AccountServlet` should be able to retrieve the SSN

and pass it on to `ReporterServlet`. In this case, `AccountServlet` should be able to share the SSN with `ReporterServlet` only for that request. Once the request is serviced, `ReporterServlet` should not be able to access the SSN anymore.

- As described before, `LoginServlet` should be able to share the `userid` with `AccountServlet`, but `AccountServlet` should be able to access only the `userid` for the user whose request it is servicing. Further, it should be able to access it for as long as the user is logged in.
- All three servlets—`LoginServlet`, `AccountServlet`, and `ReporterServlet`—need to access the same database, and so `driverclassname`, `dburl`, `dbusername`, and `dbpassword` should be shared with all the servlets all of the time.

The three containers that we mentioned earlier help us in these situations:

- If we put an object in a `javax.servlet.ServletRequest` object, it can be shared with any servlet that processes that request. We will see how to pass this request object around in section 4.8.2.
- If a servlet puts an object in a `javax.servlet.http.HttpSession` object, it can be accessed by any servlet anytime but only while the servlet is servicing a request for the same client who put the object into the session, and only while that session is valid. We will learn about this in detail in chapter 8, “Session management.”
- If a servlet puts an object in the `java.servlet.ServletContext` object, it can be accessed anytime by any servlet of the same web application. We will work with the `ServletContext` object in chapter 6, “The servlet container model.”

All three interfaces provide the same set of three methods for setting and getting attributes, as shown in table 4.11.

Table 4.11 The methods available in `ServletRequest`, `HttpSession`, and `ServletContext` for getting and setting attributes

Method	Description
Object <code>getAttribute</code> (String name)	This method returns the value mapped to this name or null if no such name exists.
Enumeration <code>getAttributeNames</code> ()	This method returns an Enumeration of Strings for all the names that are available in this container.
void <code>setAttribute</code> (String name, Object value)	The method adds the given name-value pair to this container. If the name is already present, then the old value is removed.

We will build a simple web application that uses these concepts in section 4.8.3.

4.8.2 Coordinating servlets using RequestDispatcher

Again, with respect to the banking application, if a user is not logged in, `AccountServlet` should forward the request to `LoginServlet`. Similarly, once a user enters her user ID/password, `LoginServlet` should forward the request to `AccountServlet`.

The Servlet API includes the `javax.servlet.RequestDispatcher` interface, which allows us to do this. It has the two methods shown in table 4.12.

Table 4.12 Methods provided by `RequestDispatcher` for forwarding/including a request to/from another resource

Method	Description
<code>void forward</code> (<code>ServletRequest</code> <code>request</code> , <code>ServletResponse</code> <code>response</code>)	This method allows a servlet to process a request partially and then forward the request to another servlet for generating the final response. It can also be used to forward a request from one active resource (a servlet or a JSP page) to another resource (servlet, JSP file, or HTML file) on the server. This method can be called only if the response is not committed; otherwise, it will throw an <code>IllegalStateException</code> .
<code>void include</code> (<code>ServletRequest</code> <code>request</code> , <code>ServletResponse</code> <code>response</code>)	This method allows the contents of another resource to be included with the response being generated by the calling resource. Unlike forwarding, control is not permanently passed to another resource. Instead, it is passed temporarily so that the other resource can partially process the request, and then the including servlet/JSP page can take over the request again and service it to completion. The included resource cannot set the headers or status code of the response; attempts to do so are ignored.

NOTE An important difference between `RequestDispatcher.forward()` and `HttpServletResponse.sendRedirect()` (which we discussed in section 4.4) is that `RequestDispatcher.forward()` is completely handled on the server side while `HttpServletResponse.sendRedirect()` sends a redirect message to the browser. In that sense, `RequestDispatcher.forward()` is transparent to the browser while `HttpServletResponse.sendRedirect()` is not.

This sounds good, but how do we obtain a `RequestDispatcher` in the first place? Simple: both `javax.servlet.ServletContext` and `javax.servlet.ServletRequest` have the method shown in table 4.13.

Table 4.13 The method in `ServletContext` and `ServletRequest` for getting a `RequestDispatcher`

Method	Description
<code>public RequestDispatcher</code> <code>getRequestDispatcher</code> (<code>String path</code>)	The path parameter is the path to the resource; for example, <code>request.getRequestDispatcher("/servlet/AccountServlet")</code> . It will not accept a path outside the current web application.

Besides the `getRequestDispatcher()` method, the `ServletContext` interface also provides a `getNamedDispatcher()` method, which allows us to dispatch requests to a component by specifying its name (as given in the deployment descriptor) instead of a full URI path.

There is an important difference between the `getRequestDispatcher()` method of `ServletContext` and that of `ServletRequest`: you can pass a relative path to the `getRequestDispatcher()` method of `ServletRequest` but not to the `getRequestDispatcher()` method of `ServletContext`. For example, `request.getRequestDispatcher("../html/copyright.html")` is valid, and the `getRequestDispatcher()` method of `ServletRequest` will evaluate the path relative to the path of the request. For the `getRequestDispatcher()` method of `ServletContext`, the path parameter cannot be relative and must start with `/`. This makes sense because `ServletRequest` has a current request path to evaluate the relative path while `ServletContext` does not.

NOTE You cannot directly forward or include a request to a resource in another web application. To do this, you need to get a reference to the `ServletContext` of the other web application using `this.getServletContext().getContext(uripath)`. Using this servlet context reference, you can retrieve an appropriate `RequestDispatcher` object as usual.

4.8.3 Accessing request-scoped attributes with `RequestDispatcher`

One important aspect of the new servlet specification involves the ability of an included or forwarded servlet to access request information through attributes. The name of the attribute depends on whether `RequestDispatcher.include()` or `RequestDispatcher.forward()` was invoked. The attributes are listed in table 4.14.

Table 4.14 Attributes available to included/forwarded servlets for obtaining information about the request

Attribute for included servlet	Attribute for forwarded servlet
<code>javax.servlet.include.request_uri</code>	<code>javax.servlet.forward.request_uri</code>
<code>javax.servlet.include.context_path</code>	<code>javax.servlet.forward.context_path</code>
<code>javax.servlet.include.servlet_path</code>	<code>javax.servlet.forward.servlet_path</code>
<code>javax.servlet.include.path_info</code>	<code>javax.servlet.forward.path_info</code>
<code>javax.servlet.include.query_string</code>	<code>javax.servlet.forward.query_string</code>

The values of these attributes provide the same information as the `HttpServletRequest` methods `getRequestURI`, `getContextPath`, `getServletPath`, `getPathInfo`, and `getQueryString`. They can be accessed with `getAttribute()` just like regular request attributes. These values correspond to the parameters of the original request, without regard to further forwarding.

For example, if a forwarded servlet invokes `req.getServletPath()`, the result will be the same as if the `java.servlet.include.servlet_path` attribute had been acquired with `req.getAttribute(?javax.servlet.include.servlet_path?)`.

NOTE These attributes can't be set if the `RequestDispatcher` was obtained with the `getNamedDispatcher()` method.

4.8.4 Putting it all together: A simple banking application

Let's build the banking application that we have been discussing, using two servlets:

- `LoginServlet`
- `AccountServlet`

LoginServlet

Listing 4.9 contains the code for `LoginServlet` that verifies the user ID and password and forwards the request to `AccountServlet`.

Listing 4.9 `LoginServlet.java` for a simple banking web application

```
package chapter4;

import java.io.*;
import java.util.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class LoginServlet extends HttpServlet
{
    Hashtable users = new Hashtable();

    //This method will be called if somebody types the URL
    //for this servlet in the address field of the browser.
    public void doGet(HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException
    {
        doPost(req, res);
    }

    //This method retrieves the userid and password, verifies them,
    //and if valid, it forwards the request to AccountServlet.
    //Otherwise, it forwards the request to the login page.
    public void doPost(HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException
    {
        String userid = req.getParameter("userid");
        String password = req.getParameter("password");

        if( userid != null && password != null &&
            password.equals(users.get(userid)) )
        {
            req.setAttribute("userid", userid);
        }
    }
}
```

Retrieves the
user ID/password

Sets the userid
in the request

```

ServletContext ct = getServletContext();
RequestDispatcher rd =
    ct.getRequestDispatcher("/servlet/AccountServlet");
rd.forward(req, res);
return;
}
else
{
    RequestDispatcher rd =
        req.getRequestDispatcher("../login.html");
    rd.forward(req, res);
    return;
}
}

//initialize some userids and passwords
public void init()
{
    users.put("ann", "aaa");
    users.put("john", "jjj");
    users.put("mark", "mmm");
}
}

```

Gets the Request-Dispatcher for AccountServlet

Forwards the request to AccountServlet

Gets the Request-Dispatcher for login.html

The logic of authenticating a user in the servlet in listing 4.9 is simple. We initialize a `Hashtable` to store some user IDs and passwords. In the `doPost()` method, the servlet validates the credentials given by the user and forwards the request to either `AccountServlet` or to the login page.

Observe the use of an absolute path used in the creation of `RequestDispatcher` for `AccountServlet` and the use of a relative path in the creation of `RequestDispatcher` for the login page.

Login.html

To access our application, a user will go to the login page. This page will allow the user to enter her user ID and password. Listing 4.10 contains the code for the login page.

Listing 4.10 login.html for capturing the userid and password

```

<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN">
<html>
<head>
    <title>SCWCD_Example_1_3</title>
</head>
<body>
<h3>Please enter your userid and password to see your account statement:</h3><p>
<form action="/servlet/LoginServlet" method="POST">

```

Sends the data to LoginServlet using POST

```

Userid : <input type="text" name="userid"><br><br>
Password : <input type="password" name="password"><br><br>
<input type="submit" value="Show Statement">
</form>

</body>
</html>

```

AccountServlet

Listing 4.11 contains the code for AccountServlet. Its job is to generate an HTML page that displays the user's account information.

Listing 4.11 AccountServlet.java

```

package chapter4;

import java.io.*;
import java.util.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class AccountServlet extends HttpServlet
{
    Hashtable data = new Hashtable();

    //This method will be called if somebody types the URL
    //for this servlet in the address field of the browser.
    public void doGet(HttpServletRequest req, HttpServletResponse res)
        throws javax.servlet.ServletException, java.io.IOException
    {
        doPost(req, res);
    }

    public void doPost(HttpServletRequest req, HttpServletResponse res)
        throws javax.servlet.ServletException, java.io.IOException
    {
        String userid = (String) req.getAttribute("userid");
        if(userid != null )
        {
            // Retrieve the data and generate the page dynamically.
            String[] records = (String[]) data.get(userid);

            PrintWriter pw = res.getWriter();
            pw.println("<html>");
            pw.println("<head>");
            pw.println("</head>");
            pw.println("<body>");
            pw.println("<h3>Account Status for "+userid+"
                at the start of previous three months...</h3><p>");
            for(int i=0; i<records.length; i++)

```

← Gets the
userid set by
LoginServlet


```

        {
            pw.println(records[i]+"<br>");
        }

        pw.println("</body>");
        pw.println("</html>");
    }
    else
    {

        //No user ID. Send login.html to the user.
        //observe the use of relative path.
        RequestDispatcher rd =
            req.getRequestDispatcher("../login.html");
        rd.forward(req, res);
    }
}

//initialize some data.
public void init()
{
    data.put("ann", new String[]{ "01/01/2002 : 1000.00",
        "01/02/2002 : 1300.00", "01/03/2002 : 900.00" } );
    data.put("john", new String[]{ "01/01/2002 : 4500.00",
        "01/02/2002 : 2100.00", "01/03/2002 : 2600.00" } );
    data.put("mark", new String[]{ "01/01/2002 : 7800.00",
        "01/02/2002 : 5200.00", "01/03/2002 : 1900.00" } );
}
}
}

```

← Creates a request dispatcher using the relative path

Running the application

We have provided the above examples on the Manning web site. Just copy the `chapter04` directory from the web site to the `webapps` directory of your Tomcat installation. For example, if you have installed Tomcat to `c:\jakarta-tomcat-5.0.25`, copy the `chapter04` directory to `c:\jakarta-tomcat-5.0.25\webapps`. Once you restart Tomcat, you can go to `http://localhost:8080/chapter04/login.html`.

You should observe the following:

- If you enter an invalid user ID/password, you get the login page.
- If you enter a valid user ID/password, you get the statement page.

You may also notice that the user needs to enter the user ID and password every time he tries to access `AccountServlet`. This is indeed annoying. There should be some way for `AccountServlet` to remember the user ID during the time in which it is interacting with the user. There is: that's where the concept of *sessions* comes into the picture. We will discuss sessions in detail in chapter 8.

4.9 SUMMARY

In this chapter, we discussed the basics of the servlet model. An `HttpServlet` has methods, which can be overridden, that correspond to the HTTP methods of the request. The `service(HttpServletRequest, HttpServletResponse)` method of `HttpServlet` is responsible for calling the appropriate method on the servlet depending on the request.

Using `HttpServletRequest` and `HttpServletResponse`, we learned how to create a dynamic response by analyzing the request. We also discussed the life-cycle phases of a servlet, including loading, initializing, destroying, and unloading the servlet. The `ServletRequest`, `HttpSession`, and `ServletContext` objects are the containers used to share data within the three scopes of a servlet: request, session, and application.

Finally, we developed a small web application, in which we shared data between two servlets and coordinated their execution.

You should now be ready to answer exam questions based on the HTTP methods GET, POST, and HEAD and the methods of a servlet that correspond to these HTTP methods. You should be able to answer the questions based on the servlet life cycle and the usage of servlets and related classes to retrieve request header information, form parameters, and text and binary streams. Finally, you should be able to answer questions based on attribute sharing using the request, session, and application scopes.

In the next chapter, we will take a closer look at the structure of a web application, the deployment descriptor, and the way in which a request is mapped to a servlet.

4.10 REVIEW QUESTIONS

1. Which method in the `HttpServlet` class services the HTTP POST request? (Select one)

- a `doPost(ServletRequest, ServletResponse)`
- b `doPOST(ServletRequest, ServletResponse)`
- c `servicePost(HttpServletRequest, HttpServletResponse)`
- d `doPost(HttpServletRequest, HttpServletResponse)`

2. Consider the following HTML page code:

```
<html><body>  
<a href="/servlet/HelloServlet">POST</a>  
</body></html>
```

Which method of `HelloServlet` will be invoked when the hyperlink displayed by the above page is clicked? (Select one)

- a `doGet`
- b `doPost`

- c doForm
- d doHref
- e serviceGet

3. Consider the following code for the doGet () method:

```
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
{
    PrintWriter out = res.getWriter();
    out.println("<html><body>Hello</body></html>");

    //1

    if(req.getParameter("name") == null)
    {
        res.sendError(HttpServletResponse.SC_UNAUTHORIZED);
    }
}
```

Which of the following lines can be inserted at //1 so that the above code does not throw any exception? (Select one)

- a if (! res.isSent())
 - b if (! res.isCommitted())
 - c if (! res.isDone())
 - d if (! res.isFlushed())
 - e if (! res.flush())
4. Which of the following lines would initialize the out variable for sending a Microsoft Word file to the browser? (Select one)

- a PrintWriter out = response.getServletOutput();
- b PrintWriter out = response.getPrintWriter();
- c OutputStream out = response.getWriter();
- d PrintWriter out = response.getOutputStream();
- e OutputStream out = response.getOutputStream();
- f ServletOutputStream out = response.getServletOutputStream();

5. You need to send a GIF file to the browser. Which of the following lines should be called after (or before) a call to response.getOutputStream()? (Select one)

- a response.setContentType("image/gif"); Before
- b response.setContentType("image/gif"); After
- c response.setContentType("image/gif"); Before
- d response.setContentType("image/gif"); After
- e response.setStreamType("image/gif"); Before
- f response.setStreamType("image/gif"); After

6. Consider the following HTML page code:

```
<html><body>
<form name="data" action="/servlet/DataServlet" method="POST">
<input type="text" name="name">
<input type="submit" name="submit">
</form>
</body></html>
```

Identify the two methods that can be used to retrieve the value of the name parameter when the form is submitted.

- a `getParameter("name");`
- b `getParameterValue("name");`
- c `getParameterValues("name");`
- d `getParameters("name");`
- e `getValue("name");`
- f `getName();`

7. Which of the following methods would you use to retrieve header values from a request? (Select two)

- a `getHeader()` Of `ServletRequest`
- b `getHeaderValue()` Of `ServletRequest`
- c `getHeader()` Of `HttpServletRequest`
- d `getHeaders()` Of `ServletRequest`
- e `getHeaders()` Of `HttpServletRequest`

8. Consider the following code:

```
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
    throws IOException
{
    if(req.getParameter("switch") == null)
    {
        //1
    }
    else
    {
        //other code
    }
}
```

Which of the following lines can be inserted at `//1` so that the request is redirected to `collectinfo.html` page? (Select one)

- a `req.sendRedirect("collectinfo.html");`
- b `req.redirect("collectinfo.html");`
- c `res.direct("collectinfo.html");`

```
d res.sendRedirect("collectinfo.html");  
e this.sendRedirect("collectinfo.html");  
f this.send("collectinfo.html");
```

9. Consider the following code:

```
public void doGet(HttpServletRequest req,  
                  HttpServletResponse res)  
{  
    HttpSession session = req.getSession();  
    ServletContext ctx = this.getServletContext();  
  
    if(req.getParameter("userid") != null)  
    {  
        String userid = req.getParameter("userid");  
        //1  
    }  
}
```

You want the `userid` parameter to be available only to the requests that come from the same user. Which of the following lines would you insert at `//1`? (Select one)

```
a session.setAttribute("userid", userid);  
b req.setAttribute("userid", userid);  
c ctx.addAttribute("userid", userid);  
d session.addAttribute("userid", userid);  
e this.addParameter("userid", userid);  
f this.setAttribute("userid", userid);
```

10. Which of the following lines would you use to include the output of `DataServlet` into any other servlet? (Select one)

```
a RequestDispatcher rd =  
    request.getRequestDispatcher("/servlet/DataServlet");  
    rd.include(request, response);  
b RequestDispatcher rd =  
    request.getRequestDispatcher("/servlet/DataServlet");  
    rd.include(response);  
c RequestDispatcher rd = request.getRequestDispatcher();  
    rd.include("/servlet/DataServlet", request, response);  
d RequestDispatcher rd = request.getRequestDispatcher();  
    rd.include("/servlet/DataServlet", response);  
e RequestDispatcher rd = request.getRequestDispatcher();  
    rd.include("/servlet/DataServlet");
```



CHAPTER 5

Structure and deployment

- 5.1 Directory structure of a web application 68
- 5.2 The deployment descriptor: an overview 71
- 5.3 Summary 80
- 5.4 Review questions 80

EXAM OBJECTIVES

- 2.1** Construct the file and directory structure of a Web application that may contain

- Static content,
- JSP pages,
- Servlet classes,
- The deployment descriptor (web.xml),
- Tag libraries,
- JAR files, and
- Java class files; and

Describe how to protect resource files from HTTP access
(Section 5.1)

- 2.2** Describe the purpose and semantics of the deployment descriptor
(Section 5.2)

- 2.3** Construct the correct structure of the deployment descriptor
(Section 5.2)

- 2.4** Explain the purpose of a WAR file, describe the contents of a WAR file, and describe how one may be constructed (Section 5.1.3)

INTRODUCTION

A web application consists of many resources, including servlets, JSP pages, utility classes, third-party JAR files, HTML files, and so forth. Managing so many resources can be a difficult task in itself; to complicate matters, the resources have dependencies. For example, a servlet may depend on third-party JAR files containing ready-made components, or a servlet may redirect a request to a JSP page. This requires the resources to learn the location of the other resources. Furthermore, a web application must be portable across different servlet containers.

Fortunately, to satisfy the above requirements, the Java Servlet Specification mandates that web applications be packaged in a standard way. In this chapter, we will discuss the way we package and deploy web applications.

5.1 DIRECTORY STRUCTURE OF A WEB APPLICATION

The resources of a web application are kept in a structured hierarchy of directories. The directory structure is well defined in terms of the placement of the resources and files. Figure 5.1 shows a hypothetical web application named `helloapp`. To put the structure in perspective, we've shown it in relationship to its location within the Tomcat directory structure.

The `webapps` directory under the Tomcat installation is the home directory of all the web applications that it hosts. In the directory structure shown in figure 5.1, the `webapps` directory contains the `app1`, `sampleapp`, and `helloapp` web applications. In the following sections, we will take a closer look at the directory structure of the `helloapp` web application.

5.1.1 Understanding the document root directory

In figure 5.1, the `helloapp` directory is the document root for the `helloapp` web application. A request for `http://www.myserver.com/helloapp/index.html` will refer to the `index.html` file in the `helloapp` directory. All publicly accessible files should go in this directory. It is very common to organize these files into multiple sub-directories. A typical root directory looks like this:

```
| -  
|- helloapp  
    |- html (contains all the HTML files)  
    |- jsp (contains all the JSP files)  
    |- images (contains all the GIFs, JPEGs, BMPs)  
    |- javascripts (contains all *.js files)  
    |- index.html (default HTML file)  
    |- WEB-INF
```

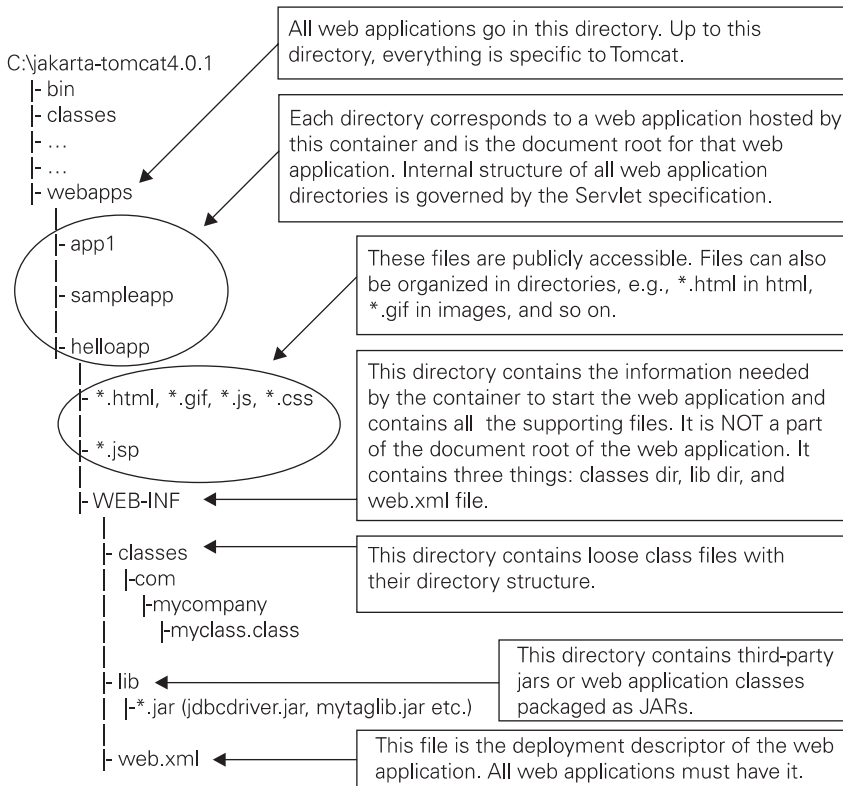


Figure 5.1 The directory structure of a web application

In the above structure, an HTML file named `hello.html` can be accessed through the URL `http://www.myserver.com/helloapp/html/hello.html`.

5.1.2 Understanding the WEB-INF directory

Every web application must have a `WEB-INF` directory directly under its root directory. Although it is physically located inside the document root directory, it is not considered a part of the document root; i.e., files in the `WEB-INF` directory are not served to the clients. This directory contains three things:

- *classes directory*—The servlet class files and the class files needed to support the servlets or JSP pages of this web application go in this directory if they have not been included in a JAR file. The class files should be organized according to their packages. At runtime, the servlet container adds this directory to the class-path for this web application.
- *lib directory*—All the JAR files used by the web application, including the third-party JAR files, go in this directory. For example, if a servlet uses JDBC to connect to a database, the JDBC driver JAR file should go here. We can also package

the servlet classes in a JAR file and keep that file in this directory. At runtime, the servlet container adds all the JAR files from this directory to the classpath for this web application.

- *web.xml file (also known as the deployment descriptor)*—This file is the heart of a web application, and every web application must have it. It contains the information needed by the servlet container in order to run the web application, such as servlet declarations and mappings, properties, authorization and security constraints, and so forth. We will learn more about this file in section 5.2.

Quizlet

- Q:** Your web application includes an applet packaged as a JAR file. Which directory would you keep the JAR file in?
- A:** Because an applet is only run on the client side, the applet JAR file should be accessible to the clients. This means that it may be kept anywhere in the document root of the application except in the `WEB-INF` directory and its subdirectories.

5.1.3 The web archive (WAR) file

Since a web application contains many files, it can be cumbersome to migrate the application from one environment to another—for instance, from development to production. To simplify the process, these files can be bundled into a single JAR file but with the extension `.war` instead of `.jar`. The extension `.war` stands for *web archive* and signifies that the file should be treated differently than a JAR file. For example, if we place a WAR file in Tomcat's `webapps` directory, Tomcat can be configured to automatically extract its contents to a directory under `webapps`. The name of the new directory is the same as the name of the WAR file without the extension.

In essence, a servlet container can install a WAR file as a web application without manual intervention.

Creating a WAR file is simple. For example, to create a WAR file for the `helloapp` web application, follow these steps:

- 1 From the DOS prompt (or `$` prompt), go to the `webapps\helloapp` directory (`c:\jakarta-tomcat-5.0.25\webapps\helloapp`).
- 2 Jar the `helloapp` directory by using the `jar` utility:

```
c:\jakarta-tomcat-5.0.25\webapps\helloapp>jar -cvf helloapp.war *
```

This will create a `helloapp.war` file in the `webapps\helloapp` directory.

5.1.4 Resource files and HTML access

When you create web applications, you may need to keep clients from accessing particular resources while allowing the web container to find them. To protect these files from HTTP access, you should store them in the `WEB-INF` directory of a web application or the `META-INF` directory of a WAR. Files in these directories remain visible to the web container but can't be served to the client.

5.1.5 The default web application

Besides the web applications created by the users, a servlet container maintains a default web application. This application handles all requests that do not match any of the user-created web applications. It is similar to any other web application except that we can access its resources without specifying its name or context path. In Tomcat, the `webapps\ROOT` directory is set as the document root for the default web application.

A default web application allows you to deploy individual JSPs, servlets, and static content without prepackaging them into a separate application. For example, if you want to test an individual JSP file named `test.jsp`, you can place it in the `ROOT` directory instead of creating a separate application. You can access it through the URL `http://localhost:8080/test.jsp`, and you can modify the deployment descriptor of this application to add your own components, such as servlets, as needed.

5.2 THE DEPLOYMENT DESCRIPTOR: AN OVERVIEW

The deployment descriptor (`web.xml`) of a web application describes the web application to the servlet container. As is evident from the extension of `web.xml`, it is an XML file. To ensure portability across the servlet containers, the document type definition (DTD) for this XML file is standardized by Sun. If you are new to XML technology, you should read the brief tutorial that we have provided in appendix B. It will help you to understand this section.

Table 5.1 shows the properties that can be defined in a deployment descriptor.

Table 5.1 Properties defined in a deployment descriptor

Web Application Properties	Short Description	Discussed in:
Servlet Declarations	Used to specify servlet properties.	Chapter 5
Servlet Mappings	Used to specify URL to servlet mapping.	Chapter 5
Application Lifecycle Listener classes	Used to specify listener classes for <code>HttpSession</code> Events and <code>ServletContextAttributeEvent</code> .	Chapter 6
ServletContext Init Parameters	Used to specify initialization parameters for the web application.	Chapter 6
Filter Definitions and Filter Mappings	Used to specify the filter.	Chapter 7
Session Configuration	Used to specify session timeout.	Chapter 8
Security Constraints	Used to specify security requirements of the web application.	Chapter 9
Tag libraries	Used to specify the tag libraries required by JSP pages.	Chapter 15

continued on next page

Table 5.1 Properties defined in a deployment descriptor (continued)

Web Application Properties	Short Description	Discussed in:
Welcome File list	Used to specify the welcome files for the web application.	Not needed for the exam
MIME Type Mappings	Used to specify MIME types for common file extensions.	Not needed for the exam
JNDI names	Used to specify JNDI names of the EJBs.	Not needed for the exam

We'll discuss many of these properties throughout the book as they apply to the exam objectives. In this section, we will look at the general structure of the deployment descriptor, and we will learn how to define servlets and servlet mappings in a deployment descriptor.

5.2.1 Example: A simple deployment descriptor

Listing 5.1 shows the general structure of a simple deployment descriptor.

Listing 5.1 Simple deployment descriptor

```

<?xml version="1.0" encoding="ISO-8859-1" ?>
<web-app version="2.4"
  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_4.xsd" >
  <display-name>Test Webapp</display-name>
    <context-param>
      <param-name>author</param-name>
      <param-value>john@abc.com</param-value>
    </context-param>
    <servlet>
      <servlet-name>test</servlet-name>
      <servlet-class>com.abc.TestServlet</servlet-class>
      <init-param>
        <param-name>greeting</param-name>
        <param-value>Good Morning</param-value>
      </init-param>
    </servlet>
    <servlet-mapping>
      <servlet-name>test</servlet-name>
      <url-pattern>/test/*</url-pattern>
    </servlet-mapping>
  </web-app>

```

Declares the XML version and character set used in this file

Declares the schema definition for this file

Specifies a parameter for this web application

Specifies a servlet

Specifies a parameter for this servlet

Maps /test/* to test servlet

```

        <extension>zip</extension>
        <mime-type>application/zip</mime-type>
    </mime-mapping>
</web-app>

```

A `web.xml` file, like all XML files, starts with the line `<?xml version="1.0" encoding="ISO-8859-1">`, which specifies the version of XML and the character set it is using. What comes next depends on the version of the Servlet or JSP specification you are targeting. If you do not use any specific features of JSP2.0 (for instance, EL), and stick only to methods found in the 2.3 version of the Servlet spec, then a DOCTYPE declaration specifying the location of the DTD would follow the first line:

```

<!DOCTYPE web-app PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application
2.3//EN" "http://java.sun.com/j2ee/dtds/web-app_2_3.dtd">

```

If you want to use the latest features, you need to use the XML schema notation shown here: `version="2.4"`

```

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_4.xsd" >

```

The rest of the content must go under the `<web-app>` element, which is the root of this XML file.

Now let's look at the servlet-specific elements of the deployment descriptor that are required by the exam.

5.2.2 Using the `<servlet>` element

Each `<servlet>` element under `<web-app>` defines a servlet for that web application. The following is the definition of the `<servlet>` element as given by the DTD for `web.xml`:

```

<!ELEMENT servlet (icon?, servlet-name, display-name?,
description?, (servlet-class|jsp-file), init-param*,
load-on-startup?, security-role-ref*)>

```

The code that follows demonstrates a typical use of the `<servlet>` element within a deployment descriptor:

```

<servlet>
    <servlet-name>us-sales</servlet-name>    ← The servlet name
    <servlet-class>com.xyz.SalesServlet</servlet-class>    ← The
    servlet
    class
    <init-param>
        <param-name>region</param-name>
        <param-value>USA</param-value>
    </init-param>
    <init-param>

```

The servlet parameters

```
<param-name>limit</param-name>
<param-value>200</param-value>
<init-param>
</servlet>
```

The servlet container instantiates the class of the servlet and associates it with the given servlet name. Every parameter should be specified using the `<init-param>` element. The above servlet definition tells the servlet container to create a servlet named `us-sales` using the class `com.xyz.SalesServlet`. The container passes `region` and `limit` as the initialization parameters through the `ServletConfig` object.

servlet-name

This element defines the name for the servlet. Clients can use this name to access the servlet if Tomcat's invoker servlet is turned on. For example, the servlet defined above can be accessed through the URL `http://www.myserver.com/servlet/us-sales`. This name is also used to define the URL to the servlet mapping for the servlet. This element is mandatory, and the name should be unique across the deployment descriptor. We can retrieve the name of a servlet by using the `ServletConfig.getServletName()` method.

servlet-class

This element specifies the Java class name that should be used by the servlet container to instantiate this servlet. In the previous example, the servlet container will use the `com.xyz.SalesServlet` class. This element is mandatory. This class, as well as all the classes that it depends on, should be available in the classpath for this web application. Remember that the classes directory and JAR files in the `lib` directory inside `WEB-INF` are automatically added to the classpath by the servlet container, so there is no need to set your classpath if you put the classes in either of these two places.

init-param

This element is used to pass initialization parameters to the servlet. We can have any number of `<init-param>` elements in the `<servlet>` element. Each `<init-param>` element must have one and only one set of `<param-name>` and `<param-value>` subelements. `<param-name>` defines the name of the parameter and must be unique across the servlet element. `<param-value>` defines the value for that parameter. A servlet can retrieve the initialization parameters using the method `ServletConfig.getInitParameter("paramname")`.

Notice the name `us-sales` given to the above servlet. You can define another servlet named `euro-sales` with the same servlet class and set the value of the `region` parameter to `europa`. In such cases, multiple instances of the servlet class will be created, one for each name.

Quizlet

- Q:** How can you associate an array of values for an initialization parameter of a servlet?
- A:** You can't—at least not directly! The deployment descriptor does not allow you to specify multiple parameters with the same name. So you have to do something like this:

```
<init-param>
  <param-name>countries</param-name>
  <param-value>Australia, Brazil, India, UK, US</param-value>
</init-param>
```

You would then have to parse the `param-value` string in the servlet and interpret the multiple values listed in the string.

5.2.3 Using the `<servlet-mapping>` element

Simply put, servlet mappings specify which URL patterns should be handled by which servlet. The servlet container uses these mappings to invoke the appropriate servlets depending on the actual URL. Here is the definition of the `<servlet-mapping>` element:

```
<!ELEMENT servlet-mapping (servlet-name, url-pattern)>
```

In an instance of `servlet-mapping`, `servlet-name` should contain the name of one of the servlets defined using the `<servlet>` element, and `url-pattern` can contain any string that we want to associate with this servlet.

The following are examples of using the `<servlet-mapping>` element in a deployment descriptor:

```
<servlet-mapping>
  <servlet-name>accountServlet</servlet-name>
  <url-pattern>/account/*</url-pattern>
</servlet-mapping>

<servlet-mapping>
  <servlet-name>accountServlet</servlet-name>
  <url-pattern>/myaccount/*</url-pattern>
</servlet-mapping>
```

In these mappings, we are associating `/account` and `/myaccount` URL patterns to `accountServlet`. Whenever the container receives a request URL that starts with `<webapp name>/account` or `<webapp name>/myaccount`, it will send that request to `accountServlet`.

A servlet container interprets the `url-pattern` according to the following rules:

- A string beginning with a `/` and ending with the `/*` characters is used for determining a *servlet path* mapping. We will discuss servlet paths in section 5.2.4.
- A string beginning with a `*` prefix is used to map the request to a servlet that handles the extension specified in the string. For example, the following mapping will direct all the requests ending with `.pdf` to `pdfGeneratorServlet`:

```

<servlet-mapping>
  <servlet-name>pdfGeneratorServlet</servlet-name>
  <url-pattern>*.pdf</url-pattern>
</servlet-mapping>

```

- A string containing only the / character indicates that servlet specified by the mapping becomes the default servlet of the application. In this case, the servlet path is the request URI minus the context path and the path info is null. We will discuss context path and path info in the next section.
- All other strings are used as exact matches only. For example, the following mapping will direct `http://www.mycompany.com/report` to `reportServlet`. However, it will not direct `http://www.mycompany.com/report/sales` to `reportServlet`.

```

<servlet-mapping>
  <servlet-name>reportServlet</servlet-name>
  <url-pattern>/report</url-pattern>
</servlet-mapping>

```

5.2.4 Mapping a URL to a servlet

In the previous section, we learned how to specify the servlet mappings in the deployment descriptor of a web application. Now let's look at how the container uses these mappings to route a request to the appropriate servlet. Routing a request to a servlet is a two-step process. First, the servlet container identifies the web application that the request belongs to, and then it finds an appropriate servlet of that web application to handle the request.

Both steps require the servlet container to break up the request URI into three parts: the context path, the servlet path, and the path info. Figure 5.2 shows these three components of a URL.




Figure 5.2 The context path, servlet path, and path info

Let's take a look at each component:

- *Context path*—The servlet container tries to match the longest possible part of the request URI, starting from the beginning, with the available web application names. This part is called the *context path*. For example, if the request URI is `/autobank/accountServlet/personal`, then `/autobank` is the context path (assuming that a web application named `autobank` exists within the servlet container). If there is no match, the context path is empty; in this case, it associates the request with the default web application.

- *Servlet path*—After taking out the context path, the servlet container tries to match the longest possible part of the remaining URI with the servlet mappings defined for the web application that was specified as the context path. This part is called the *servlet path* . For example, if the request URI is `/autobank/accountServlet/personal`, then `/accountServlet` is the servlet path (assuming that a servlet named `accountServlet` is defined in the `autobank` web application). If it is unable to find any match, it returns an error page. We will see how the servlet container determines this path shortly.
- *Path info*—Anything that remains after determining the servlet path is called *path info* . For example, if the request URI is `/autobank/accountServlet/personal`, `/personal` is the path info.

NOTE Remember the following three points:

- Request URI = context path + servlet path + path info.
- Context paths and servlet paths start with a `/` but do not end with it.
- `HttpServletRequest` provides three methods—`getContextPath()`, `getServletPath()` and `getPathInfo()`—to retrieve the context path, the servlet path, and the path info, respectively, associated with a request.

Identifying the servlet path

To match a request URI with a servlet, the servlet container follows a simple algorithm. Once it identifies the context path, if any, it evaluates the remaining part of the request URI with the servlet mappings specified in the deployment descriptor, in the following order. If it finds a match at any step, it does not take the next step.

- 1 The container tries to match the request URI to a servlet mapping. If it finds a match, the complete request URI (except the context path) is the servlet path. In this case, the path info is null.
- 2 It tries to recursively match the longest path by stepping down the request URI path tree a directory at a time, using the `/` character as a path separator, and determining if there is a match with a servlet. If there is a match, the matching part of the request URI is the servlet path and the remaining part is the path info.
- 3 If the last node of the request URI contains an extension (`.jsp`, for example), the servlet container tries to match it to a servlet that handles requests for the specified extension. In this case, the complete request URI, minus the context path, is the servlet path, and the path info is null.
- 4 If the container is still unable to find a match, it will forward the request to the default servlet. If there is no default servlet, it will send an error message indicating the servlet was not found.

Understanding this process is very important for the exam, because you will be asked to map a given request URI to a servlet. Although the process looks complicated, it is actually not. The following detailed example will help you to understand this process. We will assume that the following servlet mappings are defined for the `colorapp` web application in `web.xml`:

```

<servlet-mapping>
  <servlet-name>RedServlet</servlet-name>
  <url-pattern>/red/*</url-pattern>
</servlet-mapping>

<servlet-mapping>
  <servlet-name>RedServlet</servlet-name>
  <url-pattern>/red/red/*</url-pattern>
</servlet-mapping>

<servlet-mapping>
  <servlet-name>RedBlueServlet</servlet-name>
  <url-pattern>/red/blue/*</url-pattern>
</servlet-mapping>

<servlet-mapping>
  <servlet-name>BlueServlet</servlet-name>
  <url-pattern>/blue/</url-pattern>
</servlet-mapping>

<servlet-mapping>
  <servlet-name>GreenServlet</servlet-name>
  <url-pattern>/green</url-pattern>
</servlet-mapping>

<servlet-mapping>
  <servlet-name>ColorServlet</servlet-name>
  <url-pattern>*.col</url-pattern>
</servlet-mapping>

```

Table 5.2 shows the separation of request URIs into servlet path and path info. For simplicity, we have kept the context path as `colorapp` for all the examples. The table also shows the servlets used for handling the requests.

Table 5.2 Mapping a request URI to a servlet

Request URI	Servlet used	Servlet path	Path info	Comments
<code>/colorapp/red</code>	RedServlet	<code>/red</code>	null	See Step 1.
<code>/colorapp/red/</code>	RedServlet	<code>/red</code>	<code>/</code>	See Step 2.
<code>/colorapp/red/aaa</code>	RedServlet	<code>/red</code>	<code>/aaa</code>	See Step 2.
<code>/colorapp/red/blue/aa</code>	RedBlueServlet	<code>/red/blue</code>	<code>/aa</code>	See Step 2.

continued on next page

Table 5.2 Mapping a request URI to a servlet (continued)

Request URI	Servlet used	Servlet path	Path info	Comments
/colorapp/red/red/aaa	RedServlet	/red/red	/aaa	Longest matching URL mapping is chosen. So, servletpath is /red/red instead of /red. See Step 2.
/colorapp/aa.col	ColorServlet	/aa.col	null	*.col is mapped to ColorServlet. See Step 3.
/colorapp/hello/aa.col	ColorServlet	/hello/aa.col	null	/hello/aa.col matches with *.col, so the servlet path is /hello/aa.col and the path info is null. See Step 3.
/colorapp/red/aa.col	RedServlet	/red	/aa.col	RedServlet is chosen because there is a path (/red) matching with a url-mapping (/red/*). Extension mapping (*.col) is considered only if there is no match for path. See Step 2.
/colorapp/blue	NONE (Error message)			The url-pattern for BlueServlet is /blue/. Note the trailing /.
/colorapp/hello/blue/	NONE (Error message)			/hello/blue does not start with /blue.
/colorapp/blue/mydir	NONE (Error message)			There is no * in the mapping for BlueServlet.
/colorapp/blue/dir/aa.col	ColorServlet	/blue/dir/aa.col	null	There is no mapping for blue/*, so extension mapping *.col is considered. See Step 3.
/colorapp/green	GreenServlet	/green	null	See Step 1.

Quizlet

- Q:** In the above example, which servlet will handle the request with a request URI of /colorapp/blue/cool.col?
- A:** The only URI pattern that applies to BlueServlet is /blue/. In this case, our URI is /blue/cool.col, which does not match /blue/. However, it matches *.col, which maps to the ColorServlet. Therefore, ColorServlet will handle this request.

5.3 SUMMARY

We began the chapter with a look at the directory structure of a web application. The Java Servlet Specification mandates the way we package all the components, files, and other resources to ensure portability and ease of deployment. Every web application must have a deployment descriptor, named `web.xml`, which contains the information about the web application that the servlet container needs, such as servlet declarations and mappings, properties, authorization and security constraints, and so forth. We discussed the contents of the deployment descriptor and the manner in which it defines servlets and their initialization parameters. Finally, we looked at the way the servlet container uses the servlet mappings of the deployment descriptor to map a request URI to a servlet.

At this point, you should be able to answer the questions about the structure of a web application and the elements of the deployment descriptor and their uses. You should also be able to determine the servlet used for processing a request by looking at the servlet-mapping elements and the request URI.

In the next chapter, we will discuss how the components of a web application interact by describing the web container model.

5.4 REVIEW QUESTIONS

1. Which element is used to specify useful information about an initialization parameter of a servlet in the deployment descriptor? (Select one)

- a param-description
- b description
- c info
- d param-info
- e init-param-info

2. Which of the following deployment descriptor snippets correctly associates a servlet implemented by a class named `com.abc.SalesServlet` with the name `SalesServlet`? (Select one)

- a

```
<servlet>
  <servlet-name>com.abc.SalesServlet</servlet-name>
  <servlet-class>SalesServlet</servlet-class>
</servlet>
```
- b

```
<servlet>
  <servlet-name>SalesServlet</servlet-name>
  <servlet-package>com.abc.SalesServlet</servlet-package>
</servlet>
```
- c

```
<servlet>
  <servlet-name>SalesServlet</servlet-name>
  <servlet-class>com.abc.SalesServlet</servlet-class>
</servlet>
```

```
d <servlet name="SalesServlet" class="com.abc.SalesServlet">
    <servlet>
        <servlet-class name="SalesServlet">
            com.abc.SalesServlet
        </servlet-class>
    </servlet>
    <servlet>
        <servlet-name class="com.abc.SalesServlet">
            SalesServlet
        </servlet-name>
    </servlet>
```

3. A web application is located in a directory named `sales`. Where should its deployment descriptor be located? (Select one)

- a `sales`
- a `sales/deployment`
- a `sales/WEB`
- a `sales/WEB-INF`
- a `WEB-INF/sales`
- a `WEB-INF`
- b `WEB/sales`

4. What file is the deployment descriptor of a web application named `BankApp` stored in? (Select one)

- a `BankApp.xml`
- b `bankapp.xml`
- c `server.xml`
- d `deployment.xml`
- e `WebApp.xml`
- f `web.xml`

5. Your servlet class depends on a utility class named `com.abc.TaxUtil`. Where would you keep the `TaxUtil.class` file? (Select one)

- a `WEB-INF`
- b `WEB-INF/classes`
- c `WEB-INF/lib`
- d `WEB-INF/jars`
- e `WEB-INF/classes/com/abc`

6. Your web application, named `simpletax`, depends on a third-party JAR file named `taxpackage.jar`. Where would you keep this file? (Select one)

- a `simpletax`
- b `simpletax/WEB-INF`
- c `simpletax/WEB-INF/classes`

- d** simpletax/WEB-INF/lib
- e** simpletax/WEB-INF/jars
- f** simpletax/WEB-INF/thirdparty

7. Which of the following deployment descriptor elements is used to specify the initialization parameters for a servlet named `TestServlet`? (Select one)

- a** No element is needed because initialization parameters are specified as attributes of the `<servlet>` element.
- b** `<servlet-param>`
- c** `<param>`
- d** `<initialization-param>`
- e** `<init-parameter>`
- f** `<init-param>`

8. Assume that the following servlet mapping is defined in the deployment descriptor of a web application:

```
<servlet-mapping>
  <servlet-name>TestServlet</servlet-name>
  <url-pattern>*.asp</url-pattern>
</servlet-mapping>
```

Which of the following requests will not be serviced by `TestServlet`? (Select one)

- a** `/hello.asp`
- b** `/gui/hello.asp`
- c** `/gui/hello.asp/bye.asp`
- d** `/gui/*.asp`
- e** `/gui/sales/hello.asp`
- f** `/gui/asp`



CHAPTER 6

The servlet container model

- | | | | | | |
|-----|--|----|-----|---|----|
| 6.1 | Initializing ServletContext | 84 | 6.4 | Adding listeners in the deployment descriptor | 90 |
| 6.2 | Adding and listening to scope attributes | 85 | 6.5 | Web applications in a distributed environment | 92 |
| 6.3 | Servlet life-cycle events and listeners | 88 | 6.6 | Summary | 94 |
| | | | 6.7 | Review questions | 94 |

EXAM OBJECTIVES

- 3.1** For the ServletContext initialization parameters:
- Write servlet code to access initialization parameters; and
 - Create the deployment descriptor elements for initialization parameters (Sections 6.1 and 6.4)
- 3.2** For the fundamental servlet attribute scopes (request, session, and context):
- Write servlet code to add, retrieve, and remove attributes;
 - Given a usage scenario, identify the proper scope for an attribute; and
 - Identify multi-threading issues associated with each scope. (Sections 6.2 and 6.5)
- 3.4** Describe the Web container life cycle event model for requests, sessions, and web applications;
- Create and configure listener classes for each scope life cycle;
 - Create and configure scope attribute listener classes; and

- Given a scenario, identify the proper attribute listener to use.
(Section 6.3)

INTRODUCTION

Within a web application, all of the servlets share the same environment. The servlet container exposes the environment to the servlets through the `javax.servlet.ServletContext` interface. The Servlet API also defines interfaces that allow the servlets and the servlet container to interact with each other. In this chapter, we will learn about these interfaces, and we will see how to configure the environment using the deployment descriptor. We will also examine the behavior of the servlets and the servlet container in a distributed environment.

6.1 INITIALIZING SERVLETCONTEXT

Every web application has exactly one instance of `javax.servlet.ServletContext` (assuming that the servlet container is not distributed across multiple JVMs). The context is initialized at the time that the web application is loaded. Just as we have initialization parameters for a servlet, we have initialization parameters for a servlet context. These parameters are defined in the deployment descriptor of the web application, contained within a `<context-param>` element. Here is an example:

```
<web-app>
    ...
    <context-param>
        <param-name>dburl</param-name>
        <param-value>jdbc:databaseurl</param-value>
    </context-param>
    ...
</web-app>
```

The servlets of a web application can retrieve initialization parameters like these using the methods of the `ServletContext` interface, shown in table 6.1.

Table 6.1 `ServletContext` methods for retrieving the initialization parameters

Method	Description
<code>String getInitParameter(String name)</code>	Returns a <code>String</code> containing the value of the parameter, or null if the parameter does not exist
<code>java.util.Enumeration getInitParameterNames()</code>	Returns an <code>Enumeration</code> of the names of the context's initialization parameters

The servlet context initialization parameters are used to specify application-wide information, such as the developer's contact information and the database connection information. Of course, before we can use these methods we must get a reference to `ServletContext`. The following code snippet from the `init()` method of a servlet demonstrates this:

```

public void init()
{
    ServletContext context =
        getServletConfig().getServletContext();
    //ServletContext context =
        getServletContext();
    String dburl = context.getInitParameter("dburl");
    //use the dburl to create database connections
}

```

**Uses ServletConfig to
get ServletContext**

**Uses GenericServlet.get-
ServletContext**

The `ServletContext` object is contained in the `ServletConfig` object. The previous code uses the `getServletContext()` method of `ServletConfig` to get the `ServletContext`. You can also use the `getServletContext()` method of the `GenericServlet` class. `GenericServlet` provides this method since it implements the `ServletConfig` interface.

NOTE There is a difference between servlet context initialization parameters and servlet initialization parameters. Servlet context parameters belong to the web application and are accessible to all servlets and JSP pages of that web application. On the other hand, servlet initialization parameters belong to the servlet for which they are defined, and cannot be accessed by any other component of the web application.

6.2 **ADDING AND LISTENING TO SCOPE ATTRIBUTES**

You can provide a servlet context with initial parameters inside a deployment descriptor, but you can't add them programmatically. If you want to add objects to a scope (request, session, or application) inside regular code, you need *attributes*. These objects provide information during the servlet's execution and can be used to communicate between scopes and servlets.

You may need to pay attention to which attributes have been added or removed to a given scope. For this purpose, the servlet specification provides listener interfaces. These receive notification when events related to a web application occur. To respond to this notification, you must create a class that implements the corresponding listener interface and specify the class name in the deployment descriptor. Then, the servlet container will call the appropriate methods on objects of this class when the events occur.

In this section, we'll cover the process of adding and removing attributes and show how to listen to these operations from inside your code.

6.2.1 **Adding and removing scope attributes**

You can add attribute information to context objects, session objects, and request objects. Fortunately, each of these classes contains the same four methods for performing attribute operations. These are listed in table 6.2.

Table 6.2 Methods for getting and setting scope attributes

Method	Description
Object <code>getAttribute(String name)</code>	Returns the attribute object associated with the given name
Enumeration <code>getAttributeNames()</code>	Returns an Enumeration of the names of the scope's attributes
void <code>setAttribute(String name, Object object)</code>	Associates the object with the scope and identifies it with a String name
void <code>removeAttribute(String name)</code>	Removes the attribute from the scope

The first two methods return objects added to the specified scope. You need to keep three important points in mind:

- `HttpServletRequest` attributes reset after each request, but session attributes are available to any servlet in the application that receives a request as part of the session.
- Any attribute associated with the context is available to all servlets in the application.
- You can only associate one attribute with a given name.

Before setting an attribute for a scope, you need to acquire its corresponding scope object. For example, if you want to add a username to a session, you could use

```
HttpSession session = req.getSession(true);
session.setAttribute("username", "Joe Programmer");
```

This process is straightforward and works similarly for requests and servlet contexts.

6.2.2 Listening to attribute events

Now that you understand how to add and remove attributes from scope objects, we need to show you how to receive these attribute events. Each scope (context, session, and request) provides different listeners and events for keeping track of attributes. In this section, we'll look at each of them.

Listening to request attribute events

If you need to keep track of events involving the addition or removal of request attributes, you should create a class that implements the `ServletRequestAttributeListener` interface. This class will respond to attribute events as long as the request hasn't left the application's scope. Table 6.3 lists the methods available for handling these events.

These methods are simple to understand. When you add an attribute to a request, the `attributeAdded()` method will execute with information taken from the `ServletRequestAttributeEvent`. The `attributeRemoved()` method will activate when an attribute is removed from the request, and `attributeReplaced()` will function when an attribute's value is changed.

Table 6.3 ServletRequestAttributeListener methods

Method	Description
<code>void attributeAdded(ServletRequestAttributeEvent sre)</code>	Called when an attribute is added to a request
<code>void attributeRemoved(ServletRequestAttributeEvent sre)</code>	Called when an attribute is removed from a request
<code>void attributeReplaced(ServletRequestAttributeEvent sre)</code>	Called when an attribute is replaced in a request

Listening to session attribute events

Three listener interfaces are available for keeping track of sessions and their attributes: `HttpSessionAttributeListener`, `HttpSessionBindingListener`, and `HttpSessionActivationListener`. We'll cover each of them in chapter 8, when we discuss the overall framework of session operation.

Listening to context attribute events

The `ServletContextAttributeListener` interface is used to receive notifications about the changes to the attribute list of a servlet context. It has three methods, as shown in table 6.4.

Table 6.4 ServletContextAttributeListener methods

Method	Description
<code>void attributeAdded(ServletContextAttributeEvent scae)</code>	Called when a new attribute is added to the servlet context
<code>void attributeRemoved(ServletContextAttributeEvent scae)</code>	Called when an existing attribute is removed from the servlet context
<code>void attributeReplaced(ServletContextAttributeEvent scae)</code>	Called when an attribute of the servlet context is replaced

To use this capability, write a class that implements the interface and specify the name of the class in the deployment descriptor. The servlet container will call its methods automatically when relevant events occur.

In addition to these attribute listeners, the servlet specification provides interfaces for keeping track of the servlet's life cycle. These are important to understand, and we'll cover them in the next section.

6.3 **SERVLET LIFE-CYCLE EVENTS AND LISTENERS**

We saw in the previous section that the servlet container creates events whenever an attribute is added to or removed from a scope object. Similarly, it creates events whenever these objects are created and destroyed.

Many times, it is useful to pay attention to these events. For example, we can log an entry in the log file when the context is created, or we can page the support people when the context is destroyed. The Servlet Specification 2.4 defines the three listener interfaces that make this possible.

6.3.1 **javax.servlet.ServletContextListener**

This interface allows a developer to know when a servlet context is initialized or destroyed. For example, we might want to create a database connection as soon as the context is initialized and close it when the context is destroyed.

Table 6.5 shows the two methods of the `ServletContextListener` interface.

Table 6.5 `ServletContextListener` methods

Method	Description
<code>void contextDestroyed(ServletContextEvent sce)</code>	Called when the context is destroyed
<code>void contextInitialized(ServletContextEvent sce)</code>	Called when the context is initialized

Implementing the interface is rather trivial. Listing 6.1 shows how we can write a class that implements the interface in order to use these notifications to open and close a database connection.

Listing 6.1 Implementing `ServletContextListener` to create a database connection

```
package com.abcinc;

import javax.servlet.*;
import java.sql.*;

public class MyServletContextListener implements
    ServletContextListener
{
    public void contextInitialized(ServletContextEvent sce)
    {
        try
        {
            Connection c = //create connection to database;
                sce.getServletContext().setAttribute("connection", c);
        } catch (Exception e) { }
    }

    public void contextDestroyed(ServletContextEvent sce)
    {
```

```

        try
        {
            Connection c = (Connection)
                sce.getServletContext().getAttribute("connection");
            c.close();
        } catch (Exception e) { }
    }
}

```

In listing 6.1, we create a database connection in the `contextInitialized()` method and store it in `ServletContext`. Since `ServletContext` is accessible to all of the servlets of the web application, the database connection is also available to them. The `contextDestroyed()` method will be called when the servlet container takes the web application out of service and is thus an ideal place to close the database connection.

Notice the use of the `ServletContextEvent` object that is passed in the `contextInitialized()` and `contextDestroyed()` methods. We use this object to retrieve a reference to the `ServletContext` object of the web application. `ServletContextEvent` extends `java.util.EventObject`.

6.3.2 **javax.servlet.Http.HttpSessionListener**

`HttpSessionListener` contains the same methods and capabilities as `ServletContextListener`, but provides notification whenever a session is created or destroyed. Like the listeners used for session attributes, we'll cover this important interface in chapter 8, when we go over the process of session operation.

6.3.3 **javax.servlet.Http.HttpServletRequestListener**

The last listener interface we'll present keeps track of the life cycle of an `HttpServletRequest`. It receives notification whenever a request comes into scope. Table 6.6 lists its methods.

Table 6.6 `ServletRequestListener` methods

Method	Description
<code>void requestDestroyed (ServletRequestEvent sce)</code>	Called when the request is destroyed
<code>void requestInitialized (ServletRequestEvent sce)</code>	Called when the request is initialized

It's important to inform the web container about your listener classes. This involves setting parameters within the deployment descriptor. We'll investigate this next.

Quizlet

- Q:** Which application event listeners are notified when a web application starts up?
- A:** When the application starts up, the servlet context of the application is created. Therefore, only `ServletContextListeners` are notified.

6.4 ADDING LISTENERS IN THE DEPLOYMENT DESCRIPTOR

We can configure the properties of a web application context by using the deployment descriptor. The following is the definition of the `<web-app>` element. You don't have to memorize this definition, but it's helpful to see all the elements in one place.

```
<!ELEMENT web-app (icon?, display-name?, description?, distributable?,
context-param*, filter*, filter-mapping*, listener*, servlet*,
servlet-mapping*, session-config?, mime-mapping*, welcome-file-list?,
error-page*, taglib*, resource-env-ref*, resource-ref*,
security-constraint*, login-config?, security-role*, env-entry*, ejb-ref*,
ejb-local-ref*)>
```

The properties of the web application are accessible through `ServletContext`. Since the properties apply to all of the components of a web application, it is logical that the elements used to configure the properties come directly under the `<web-app>` element. Let's look at these elements briefly:

- *display-name*—Defines a short name that can be used by development tools, such as IDEs.
- *description*—Defines the usage and any important information that the developer might want to convey to the deployer.
- *distributable*—Indicates that the application can be distributed across multiple JVMs.
- *context-param*—Specifies initialization parameters for the web application. It contains a `<param-name>`, a `<param-value>`, and an optional `<description>` element. In section 6.1, we saw how the `ServletContext` initialization parameters can be used to create a database connection. The following lines specify a `dburl` parameter used by a servlet:

```
<context-param>
  <param-name>dburl</param-name>
  <param-value>jdbc:odbc:MySQLODBC</param-value>
</context-param>
```

We can have as many `<context-param>` elements as we need.

- *listener*—Specifies the classes that listen for the application events that we discussed in section 6.2. It contains one and only one `listener-class` element that specifies the fully qualified class name of the class that implements

the listener interface. The following lines show how we can configure two classes that implement the `ServletContextListener` and `ServletContextAttributeListener` interfaces:

```
<listener>
  <listener-class>
    com.abcinc.MyServletContextListener
  </listener-class>
</listener>

<listener>
  <listener-class>
    com.abcinc.MyServletContextAttributeListener
  </listener-class>
</listener>
```

Observe that we did not specify which class should be used for which event; that is because the servlet container will figure that out on its own. It instantiates the specified class and checks all the interfaces that the class implements. For each relevant interface, it adds the instance to its list of respective listeners. The container delivers the events to the listeners in the order the classes are specified in the deployment descriptor. These classes must be present in the `WEB-INF\classes` directory or packaged in a JAR file with other servlet classes.

NOTE You can also implement multiple listener interfaces in the same class and configure just this class to receive the various notifications through the methods of the respective interfaces. In this case, you will need only one `<listener>` element in the deployment descriptor. The servlet container will create only one instance of this class and will send all the notifications to this instance.

Quizlet

Q: You have written a class named `MyServletRequestListener` to listen for `ServletRequestEvents`. How will you configure this class in the deployment descriptor?

A: By adding a `<listener>` element in the deployment descriptor as shown here:

```
<web-app>
  ...
  <listener>
    <listener-class>MySessionAttributeListener</listener-class>
  </listener>
  ...
</web-app>
```

6.5 **WEB APPLICATIONS IN A DISTRIBUTED ENVIRONMENT**

An industrial-strength web application is expected to service thousands of simultaneous users with high reliability. It is common to distribute the applications across multiple server machines that are configured to work as a cluster. Server applications, such as the web server and the servlet container, are spread over these machines and thus work in a distributed mode. For example, one logical servlet container may actually run on multiple JVMs on multiple machines. Distributing an application has the following advantages:

- *Fail-over support*—If a server machine breaks down, another server machine can take over transparently to the users.
- *Load balancing*—Requests are assigned to the least busy server of the cluster to be serviced.

Distributing an application is not an easy task, though. Configuring the machines and the servers to work in a cluster is quite complicated. Moreover, the servlets need to be designed to run within the constraints imposed by a distributed environment. More often than not, it is easier to upgrade the machine than to distribute the application across multiple machines. However, certain requirements, like fail-over support, can only be met by clustering.

Many assumptions that we make while developing web applications for a single JVM no longer hold in a distributed environment. For example, we cannot assume that there is only one instance of a servlet; there may be multiple instances of a servlet running under different JVMs, and so we cannot use static or instance members to share data. We cannot directly use the local file system—the absolute path of the files may be different on different machines. We also have to keep the application state in a database instead of the `ServletContext`, because there will be different `ServletContext`s on different machines.

The Java Servlet Specification helps us by guaranteeing the behavior of some of the important aspects and features of a servlet container in a distributed environment.

6.5.1 **Behavior of a ServletContext**

Each web application has one and only one `ServletContext` instance on each JVM. The `ServletContext` for the default web application, however, exists on only one JVM—that is, it is not distributed.

On the exam, you will find questions based on the following points regarding the behavior of a `ServletContext` in a distributed environment:

- `ServletContext` attributes set on one JVM are not visible on another JVM. We must use a database or the session to share the information.
- A servlet container is not required to propagate `ServletContextEvents` and `ServletContextAttributeEvents` to different JVMs. This means that

changes to the `ServletContext` in one JVM may not trigger a method call on a `ServletContextListener` or a `ServletContextAttributeListener` in another JVM.

- `ServletContext` initialization parameters are available in all of the JVMs. Recall that `ServletContext` initialization parameters are specified in the deployment descriptor.

Quizlet

- Q:** Your web application uses a `ServletContextListener` to page support personnel whenever it goes down. What would be the impact on this functionality if the web application were deployed in a distributed environment?
- A:** There will be no impact on this functionality. Because an instance of `ServletContext` will be created on all the servers, the support personnel will be paged whenever any instance is destroyed.
- Q:** You maintain a list of users who are logged into the system in `ServletContext`. You print the list of these users upon request. How would this functionality be affected if your web application were deployed in a distributed environment?
- A:** This functionality will not work properly in a distributed environment. Remember that each server machine will have a separate instance of `ServletContext`. Therefore, a `ServletContext` will only know of the users who logged in through the server machine on which it resides. Obviously, a request to print the list of users will show only a partial list of the users.

6.5.2 Behavior of an `HttpSession`

In a distributed environment, the semantics of an `HttpSession` are a little different than those of a `ServletContext`. The specification mandates that requests belonging to a session must be serviced by only one JVM at a time. However, the container may migrate the session to another JVM for load balancing or fail-over.

You should remember the following points regarding the behavior of an `HttpSession` in a distributed environment:

- An `HttpSession` can reside on only one JVM at a time.
- A servlet container is not required to propagate `HttpSessionEvents` to different JVMs.
- Attributes of a session that implement the `java.io.Serializable` interface are migrated appropriately when the session migrates. This does not mean that if the attributes implement the `readObject()` and `writeObject()` methods, they will definitely be called.

- A container notifies all of the session attributes that implement the `HttpSessionActivationListener` interface when it migrates the session.
- A container may throw an `IllegalArgumentException` in the `setAttribute()` method of `HttpSession` if the attribute is not `Serializable`.

6.6 SUMMARY

The servlets of a web application share the application's environment through the methods of the `ServletContext` object. When a web application is loaded, `ServletContext` is initialized using the parameters that have been defined in the deployment descriptor. In this chapter, we learned how to use the initialization parameters of `ServletContext`.

Listener interfaces are implemented in order to receive notifications of certain events in a web application. We discussed the uses of the `ServletContextListener`, the `ServletContextAttributeListener`, the `ServletRequestAttributeListener`, and their configuration in the deployment descriptor.

A web application can be distributed across multiple servers to improve performance and reliability. We discussed the ways that `ServletContext` and `HttpSession` function in a distributed environment.

You should now be able to answer the questions about servlet context initialization parameters, the application event listener classes, and the behavior of `ServletContext` and `HttpSession` in a distributed environment.

In the next chapter, we'll explore the topic of filters.

6.7 REVIEW QUESTIONS

1. Which of the following methods will be invoked when a `ServletContext` is destroyed? (Select one)
 - a `contextDestroyed()` Of `javax.servlet.ServletContextListener`
 - b `contextDestroyed()` Of `javax.servlet.HttpServletContextListener`
 - c `contextDestroyed()` Of `javax.servlet.http.ServletContextListener`
 - d `contextDestroyed()` Of `javax.servlet.http.HttpServletContextListener`
2. Which of the following methods will be invoked when a `ServletContext` is created? (Select one)
 - a `contextInstantiated()` Of `javax.servlet.ServletContextListener`
 - b `contextInitialized()` Of `javax.servlet.ServletContextListener`
 - c `contextInited()` Of `javax.servlet.ServletContextListener`
 - d `contextCreated()` Of `javax.servlet.ServletContextListener`
3. Consider the following class:


```
import javax.servlet.*;
public class MyListener implements ServletContextAttributeListener
{
```

```

    public void attributeAdded(ServletContextAttributeEvent scab)
    {
        System.out.println("attribute added");
    }
    public void attributeRemoved(ServletContextAttributeEvent scab)
    {
        System.out.println("attribute removed");
    }
}

```

Which of the following statements about the above class is correct? (Select one)

- a This class will compile as is.
 - b This class will compile only if the `attributeReplaced()` method is added to it.
 - c This class will compile only if the `attributeUpdated()` method is added to it.
 - d This class will compile only if the `attributeChanged()` method is added to it.
4. Which method is used to retrieve an attribute from a `ServletContext`? (Select one)
- a `String getAttribute(int index)`
 - b `String getObject(int index)`
 - c `Object getAttribute(int index)`
 - d `Object getObject(int index)`
 - e `Object getAttribute(String name)`
 - f `String getAttribute(String name)`
 - g `String getObject(String name)`
5. Which method is used to retrieve an initialization parameter from a `ServletContext`? (Select one)
- a `Object getInitParameter(int index)`
 - b `Object getParameter(int index)`
 - c `Object getInitParameter(String name)`
 - d `String getInitParameter(String name)`
 - e `String getParameter(String name)`
6. Which deployment descriptor element is used to specify a `ServletContextListener`? (Select one)
- a `<context-listener>`
 - b `<listener>`
 - c `<servlet-context-listener>`
 - d `<servletcontextlistener>`
 - e `<servletcontext-listener>`

7. Which of the following `web.xml` snippets correctly specify an initialization parameter for a servlet context? (Select one)

- a `<context-param>`
 `<name>country</name>`
 `<value>USA</value>`
 `<context-param>`
- b `<context-param>`
 `<param name="country" value="USA" />`
 `<context-param>`
- c `<context>`
 `<param name="country" value="USA" />`
 `<context>`
- d `<context-param>`
 `<param-name>country</param-name>`
 `<param-value>USA</param-value>`
 `<context-param>`

8. Which of the following is not a requirement of a distributable web application? (Select one)

- a It cannot depend on the notification events generated due to changes in the `ServletContext` attribute list.
- b It cannot depend on the notification events generated due to changes in the session attribute list.
- c It cannot depend on the notification events generated when a session is activated or passivated.
- d It cannot depend on the notification events generated when `ServletContext` is created or destroyed.
- e It cannot depend on the notification events generated when a session is created or destroyed.

9. Which of the following is a requirement of a distributable web application? (Select one)

- a It cannot depend on `ServletContext` for sharing information.
- b It cannot depend on the `sendRedirect()` method.
- c It cannot depend on the `include()` and `forward()` methods of the `RequestDispatcher` class.
- d It cannot depend on cookies for session management.



CHAPTER 7

Using filters

- | | |
|------------------------------|---------------------------|
| 7.1 What is a filter? 98 | 7.4 Advanced features 110 |
| 7.2 The Filter API 102 | 7.5 Summary 117 |
| 7.3 Configuring a filter 106 | 7.6 Review questions 117 |

EXAM OBJECTIVES

- 3.3** Describe the web container request processing model;
- Write and configure a filter;
 - Create a request or response wrapper; and
 - Given a design problem, describe how to apply a filter or a wrapper
- (Sections 7.1 through 7.3)

INTRODUCTION

Filters are a recent addition to servlet development. This subject has gained importance, and is one of the new aspects of the SCWCD exam.

7.1 WHAT IS A FILTER?

In technical terms, a *filter* is an object that intercepts a message between a data source and a data destination, and then filters the data being passed between them. It acts as a guard, preventing undesired information from being transmitted from one point to another. For example, a Digital Subscriber Line (DSL) filter sits between the DSL line and the telephone equipment, and allows normal telephone frequencies to pass through the phone line to the telephone but blocks the frequencies meant for DSL modems. A filter in your e-mail system allows genuine e-mail messages to reach your inbox while blocking spam. These filters screen out undesired parts of the original messages. Another example of a filter is in data transmissions over TCP/IP; as it receives the data packets, the lower layer (IP) removes the information that was intended just for that layer from the data packets before sending the packets to the upper layer (TCP).

For a web application, a *filter* is a web component that resides on the web server and filters the requests and responses that are passed between a client and a resource.

Figure 7.1 illustrates the general idea of a filter in a web application. It shows the request passing through a filter on its way to a servlet. The servlet generates the response as usual; the response also passes through the filter on its way to the client. The filter can thus monitor the request and the response before they reach their destination. As shown in figure 7.1, the existence of a filter is transparent to the client as well as to the servlet.

We can also employ a chain of filters, if necessary, in which each filter processes the request and passes it on to the next filter in the chain (or to the actual resource if it is the last filter in the chain). Similarly, each filter processes the response in the reverse order before the response reaches the client. This process is illustrated in figure 7.2. Observe that a request will be processed by the filters in this order: Filter1, Filter2, and Filter3. However, the response will be processed by the filters in this order: Filter3, Filter2, and Filter1.

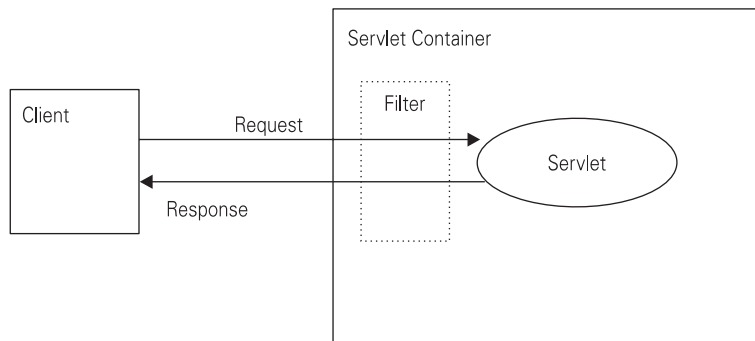


Figure 7.1 A single filter

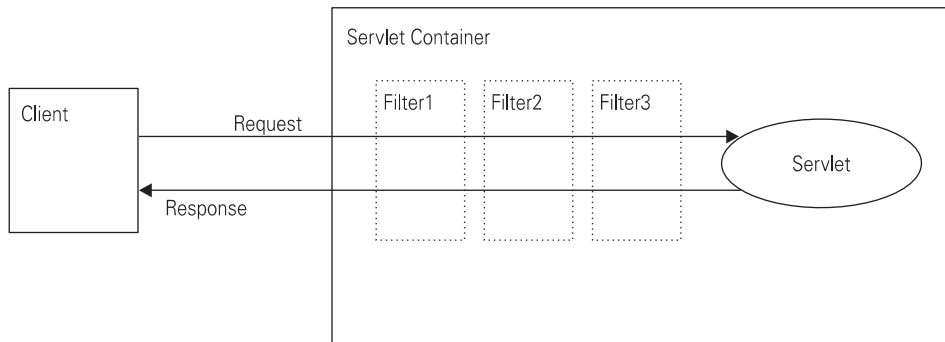


Figure 7.2 Using multiple filters

This is a very simple explanation of filters. As we will learn in the following sections, filters can do much more than just monitor the communication between the client and the server. In general, filters allow us to

- Analyze a request and decide whether to pass on the request to the resource or create a response on its own.
- Manipulate a request, including a request header, by wrapping it into a customized request object before it is delivered to a resource.
- Manipulate a response by wrapping it into a customized response object before it is delivered to the client.

7.1.1 How filtering works

When a servlet container receives a request for a resource, it checks whether a filter is associated with this resource. If a filter is associated with the resource, the servlet container routes the request to the filter instead of routing it to the resource. The filter, after processing the request, does one of three things:

- It generates the response itself and returns it to the client.
- It passes on the request (modified or unmodified) to the next filter in the chain (if any) or to the designated resource if this is the last filter.
- It routes the request to a different resource.

As it returns to the client, the response passes back through the same set of filters in the reverse order. Each filter in the chain may modify the response.

7.1.2 Uses of filters

Some of the common applications of filters identified by the Servlet specification are

- Authentication filters
- Logging and auditing filters
- Image conversion filters

- Data compression filters
- Encryption filters
- Tokenizing filters
- Filters that trigger resource access events
- Extensible Stylesheet Language Transformation (XSLT) filters
- MIME-type chain filters

7.1.3 The Hello World filter

To get a feel for filters, let's write a simple Hello World filter. In this section, we will look at the four steps—coding, compiling, deploying, and running—involved in developing and using a filter. This filter will intercept all of the requests matching the URI pattern `/filter/*` and will respond with the `Hello Filter World` message.

Code

All filters implement the `javax.servlet.Filter` interface. Listing 7.1 shows the code for `HelloWorldFilter.java`. It declares one class, `HelloWorldFilter`, that implements the `Filter` interface and defines three methods—`init()`, `doFilter()`, and `destroy()`—that are declared in the `Filter` interface.

Listing 7.1 `HelloWorldFilter.java`

```
import java.io.*;
import javax.servlet.*;

public class HelloWorldFilter implements Filter
{
    private FilterConfig filterConfig;
    public void init(FilterConfig filterConfig)
    {
        this.filterConfig = filterConfig;
    }

    public void doFilter(
        ServletRequest request,
        ServletResponse response,
        FilterChain filterChain
        ) throws ServletException, IOException
    {
        PrintWriter pw = response.getWriter();
        pw.println("<html>");
        pw.println("<head>");
        pw.println("</head>");
        pw.println("<body>");
        pw.println("<h3>Hello Filter World!</h3>");
        pw.println("</body>");
    }
}
```

```

        pw.println("</html>");
    }

    public void destroy()
    {
    }
}

```

The code that implements the filter in listing 7.1 is similar to the code that implements a servlet. First, we import the required packages, `javax.servlet` and `java.io`. The `ServletRequest`, `ServletResponse`, `ServletException`, `FilterConfig`, `Filter`, and `FilterChain` classes and interfaces belong to the `javax.servlet` package, while the `PrintWriter` and `IOException` classes belong to the `java.io` package. Since we are not using HTTP-specific features in this code, we don't need to import the `javax.servlet.http` package.

Next, we declare the `HelloWorldFilter` class. It implements all of the methods declared in the `Filter` interface. We will learn more about these methods in section 7.2.

Compilation

As usual, we include the `servlet.jar` (located under the directory `c:\jakarta-tomcat-5.0.25\common\lib\`) in the classpath and compile the `HelloWorldFilter.java` file.

Deployment

Just like with a servlet, the deployment of a filter is a two-step process:

- 1 Copy the file `HelloWorldFilter.class` to the `WEB-INF\classes` directory of the web application corresponding to this chapter:

```
c:\jakarta-tomcat-5.0.25\webapps\chapter07\WEB-INF\classes
```

- 2 Specify the filter class and map the required request URLs to this filter in the deployment descriptor:

```

<web-app>
    <!-- specify the Filter name and the Filter class -->
    <filter>
        <filter-name>HelloWorldFilter</filter-name>
        <filter-class>HelloWorldFilter</filter-class>
    </filter>

    <!-- associate the Filter with a URL pattern -->
    <filter-mapping>
        <filter-name>HelloWorldFilter</filter-name>
        <url-pattern>/filter/*</url-pattern>
    </filter-mapping>
</web-app>

```


We will see the details of these elements in section 7.4.

You could also copy the `chapter07` directory directly from the Manning web site to your `c:\jakarta-tomcat-5.0.25\webapps` directory. This directory contains all the files needed to run the example.

Execution

Start Tomcat and enter this URL in your browser's navigation bar:

```
http://localhost:8080/chapter07/filter
```

The browser should display the message `Hello Filter World!`. Notice that we have not put any resource on the server with the above URL. You can enter any URL matching the pattern `/filter/*` and the filter will still execute without any problem. Thus, the resource to which a filter is mapped does not have to exist.

7.2 THE FILTER API

The Filter API is not in a separate package. The set of classes and interfaces used by filters is part of the `javax.servlet` and `javax.servlet.http` packages. Table 7.1 describes the three interfaces and four classes used by filters.

Table 7.1 Interfaces used by filters

Interface/Class	Description
Interfaces of the package <code>javax.servlet</code>	
<code>javax.servlet.Filter</code>	We implement this interface to write filters.
<code>javax.servlet.FilterChain</code>	The servlet container provides an object of this interface to the filter developer at request time. This object gives the developer a view into the invocation chain of a filtered request for a resource.
<code>javax.servlet.FilterConfig</code>	Similar to <code>ServletConfig</code> . The servlet container provides a <code>FilterConfig</code> object that contains initialization parameters for this filter.
Classes of the package <code>javax.servlet</code>	
<code>javax.servlet. ServletRequestWrapper</code>	Provides a convenient implementation of the <code>ServletRequest</code> interface that can be subclassed by developers wanting to adapt the request to a servlet/JSP.
<code>javax.servlet. ServletResponseWrapper</code>	Provides a convenient implementation of the <code>ServletResponse</code> interface that can be subclassed by developers wanting to adapt the response from a servlet/JSP.

continued on next page

Table 7.1 Interfaces used by filters (continued)

Interface/Class	Description
Classes of the package <code>javax.servlet.http</code>	
<code>javax.servlet.http.HttpServletRequestWrapper</code>	Provides a convenient implementation of the <code>HttpServletRequest</code> interface that can be subclassed by developers wanting to adapt the request to a servlet/JSP.
<code>javax.servlet.http.HttpServletResponseWrapper</code>	Provides a convenient implementation of the <code>HttpServletResponse</code> interface that can be subclassed by developers wanting to adapt the response from a servlet/JSP.

7.2.1 The Filter interface

This is the heart of the Filter API. Just as all servlets must implement the `javax.servlet.Servlet` interface (either directly or indirectly), all filters must implement the `javax.servlet.Filter` interface. It declares three methods, as shown in table 7.2.

Table 7.2 Methods of the `javax.servlet.Filter` interface

Method	Description
<code>void init(FilterConfig)</code>	Called by the container during application startup
<code>void doFilter(ServletRequest, ServletResponse, FilterChain)</code>	Called by the container for each request whose URL is mapped to this filter
<code>void destroy()</code>	Called by the container during application shutdown

The three methods of the `Filter` interface are also the life-cycle methods of a filter. Since, unlike the Servlet API, the Filter API does not provide any implementation for the `Filter` interface, all filters must implement all three methods explicitly.

The `init()` method

The servlet container calls the `init()` method on a filter instance once and only once during the lifetime of the filter. The container does not dispatch any request to a filter before this method finishes. This method gives the filter object a chance to initialize itself if required. Here is the signature of the `init()` method:

```
public void init(FilterConfig filterConfig)
           throws ServletException;
```

This method is analogous to the `init(ServletConfig)` method of the `Servlet` interface. This method is typically implemented to save the `FilterConfig` parameter for later use. We will learn more about `FilterConfig` in section 7.2.2. If the initialization fails, the `init()` method may throw a `ServletException` or a subclass of the `ServletException` to indicate the problem.

The doFilter() method

The `doFilter()` method is analogous to the `service()` method of the `Servlet` interface. The servlet container calls this method for each request with the URL that is mapped to this filter. This is the signature of the `doFilter()` method:

```
public void doFilter(ServletRequest request,
                    ServletResponse response,
                    FilterChain chain)
    throws java.io.IOException, ServletException;
```

This gives the `Filter` object a chance to process the request, forward the request to the next component in the chain, or reply to the client itself.

Note that the `request` and `response` parameters are declared of type `ServletRequest` and `ServletResponse`, respectively. Thus, the `Filter` API is not restricted to only HTTP servlets. However, if the filter is used in a web application, which uses the HTTP protocol, these variables refer to objects of type `HttpServletRequest` and `HttpServletResponse`, respectively. Casting these parameters to their corresponding HTTP types before using them is a typical implementation of this method.

The uses of this method vary from filter to filter. A simple auditing filter may retrieve the request URL, the request parameters, and the request headers, and then log them to a file. A security filter may authenticate the request and decide to either forward the request to the resource or reject access to the designated resource. Yet another type of filter may wrap the `ServletRequest` and `ServletResponse` parameter objects with wrapper classes, and alter the request and response messages partially or completely.

A rather odd implementation is to route the request to another resource using the `include()` and `forward()` methods of `RequestDispatcher`. The `RequestDispatcher` object can be obtained using `request.getRequestDispatcher()`.

In case of an irrecoverable error during the processing, `doFilter()` may decide to throw an `IOException`, a `ServletException`, or a subclass of either of these exceptions.

The destroy() method

The `destroy()` method of the `Filter` interface is analogous to the `destroy()` method of the `Servlet` interface. The servlet container calls this method as the last method on the filter object. This is the signature of the `destroy()` method:

```
public void destroy();
```

This gives the filter object a chance to release the resources acquired during its lifetime and to perform cleanup tasks, if any, before it goes out of service. This method does not declare any exceptions.

7.2.2 The FilterConfig interface

Just as a servlet has a `ServletConfig`, a filter has a `FilterConfig`. This interface provides the initialization parameters to the filter. It declares four methods, as shown in table 7.3.

Table 7.3 Methods of the `javax.servlet.FilterConfig` interface

Method	Description
<code>String getFilterName()</code>	Returns the name of the filter specified in the deployment descriptor.
<code>String getInitParameter(String)</code>	Returns the value of the parameter specified in the deployment descriptor.
<code>Enumeration getInitParameterNames()</code>	Returns the names of all the parameters specified in the deployment descriptor.
<code>ServletContext getServletContext()</code>	Returns the <code>ServletContext</code> object associated with the web application. Filters can use it to get and set application-scoped attributes.

The servlet container provides a concrete implementation of the `FilterConfig` interface. It creates an instance of this implementation class, initializes it with the initialization parameter values, and passes it as a parameter to the `Filter.init()` method. The name and initialization parameters are specified in the deployment descriptor, which we will discuss in section 7.3.

Most important, `FilterConfig` also provides a reference to the `ServletContext` in which the filter is installed. A filter can use the `ServletContext` to share application-scoped attributes with other components of the web application.

7.2.3 The FilterChain interface

The `FilterChain` interface has just one method, described in table 7.4.

Table 7.4 The `javax.servlet.FilterChain` interface method

Method	Description
<code>void doFilter(ServletRequest, ServletResponse);</code>	We call this method from the <code>doFilter()</code> method of a <code>Filter</code> object to continue the process of filter chaining. It passes the control to the next filter in the chain or to the actual resource if this is the last filter in the chain.

The servlet container provides an implementation of this interface and passes an instance of it in the `doFilter()` method of the `Filter` interface. Within the `doFilter()` method, we can use this interface to pass the request to the next component in the chain, which is either another filter or the actual resource if this is the last filter in the chain. The two parameters of type `ServletRequest` and

ServletResponse that we pass in this method are received by the next component in the chain in its `doFilter()` or `service()` method.

7.2.4 The request and response wrapper classes

`ServletRequestWrapper` and `HttpServletRequestWrapper` provide a convenient implementation of the `ServletRequest` and `HttpServletRequest` interfaces, respectively, that we can subclass if we want to alter the request before sending it to the next component of the filter chain. Similarly, `ServletResponseWrapper` and `HttpServletResponseWrapper` are used if we want to alter the response received from the previous component. These objects can be passed as parameters to the `doFilter()` method of the `FilterChain` interface. We will see how to do this in section 7.4.

7.3 CONFIGURING A FILTER

A filter is configured using two deployment descriptor elements: `<filter>` and `<filter-mapping>`. Each `<filter>` element introduces a new filter into the web application, while each `<filter-mapping>` element associates a filter with a set of request URIs. Both elements come directly under `<web-app>` and are optional. These elements are similar to the `<servlet>` and `<servlet-mapping>` elements.

7.3.1 The `<filter>` element

Here is the definition of the `<filter>` element:

```
<!ELEMENT filter (icon?, filter-name, display-name?, description?,
                 filter-class, init-param*)>
```

As you can see from the above definition, each filter requires a `<filter-name>` and a `<filter-class>` that implements the filter. Other elements—`<icon>`, `<display-name>`, `<description>`, and `<init-param>`—serve the usual purposes and are optional.

The following example illustrates the use of the `<filter>` element:

```
<filter>
  <filter-name>ValidatorFilter</filter-name>
  <description>Validates the requests</description>
  <filter-class>com.manning.filters.ValidatorFilter</filter-class>
  <init-param>
    <param-name>locale</param-name>
    <param-value>USA</param-value>
  </init-param>
</filter>
```

This code introduces a filter named `ValidatorFilter`. The servlet container will create an instance of the `com.manning.filters.ValidatorFilter` class and associate it with this name. At the time of initialization, the filter can

retrieve the `locale` parameter by calling `filterConfig.getParameterValue("locale")`.

7.3.2 The `<filter-mapping>` element

This element works exactly like the `<servlet-mapping>` element that we discussed in detail in chapter 5, “Structure and deployment.” The `<filter-mapping>` element is defined as follows:

```
<!ELEMENT filter-mapping (filter-name, (url-pattern | servlet-name))>
```

The `<filter-name>` element is the name of the filter as defined in the `<filter>` element, `<url-pattern>` is used to apply the filter to a set of requests identified by a particular URL pattern, and `<servlet-name>` is used to apply the filter to all the requests that are serviced by the servlet identified by this servlet name. In the case of `<url-pattern>`, the pattern matching follows the same rules for servlet mapping that we described in chapter 5.

The following examples illustrate the use of the `<filter-mapping>` element:

```
<filter-mapping>
  <filter-name>ValidatorFilter</filter-name>
  <url-pattern>*.doc</url-pattern>
</filter-mapping>

<filter-mapping>
  <filter-name>ValidatorFilter</filter-name>
  <servlet-name>reportServlet</servlet-name>
</filter-mapping>
```

The first filter mapping shown above associates `ValidatorFilter` with all the requests that try to access a file with the extension `.doc`, while the second filter mapping associates `ValidatorFilter` with all the requests that are to be serviced by the servlet named `reportServlet`. The servlet name used here must refer to a servlet defined using the `<servlet>` element in the deployment descriptor.

7.3.3 Configuring a filter chain

In some cases, you may need to apply multiple filters to the same request. Such filter chains can be configured using multiple `<filter-mapping>` elements. When the servlet container receives a request, it finds all the filter mappings with a URL pattern that matches the request URI. This becomes the first set of filters in the filter chain. Next, it finds all the filter mappings with a servlet name that matches the request URI. This becomes the second set of filters in the filter chain. In both sets, the order of the filters is the order in which they appear in the deployment descriptor.

To understand this process, consider the filter mappings and servlet mapping in the `web.xml` file shown in listing 7.2.

Listing 7.2 A web.xml file for illustrating filter chaining

```
<web-app>

  <filter>
    <filter-name>FilterA</filter-name>
    <filter-class>TestFilter</filter-class>
  </filter>
  <filter>
    <filter-name>FilterB</filter-name>
    <filter-class>TestFilter</filter-class>
  </filter>
  <filter>
    <filter-name>FilterC</filter-name>
    <filter-class>TestFilter</filter-class>
  </filter>
  <filter>
    <filter-name>FilterD</filter-name>
    <filter-class>TestFilter</filter-class>
  </filter>
  <filter>
    <filter-name>FilterE</filter-name>
    <filter-class>TestFilter</filter-class>
  </filter>

  <!-- associate FilterA and FilterB to RedServlet -->
  <filter-mapping>
    <filter-name>FilterA</filter-name>
    <servlet-name>RedServlet</servlet-name>
  </filter-mapping>
  <filter-mapping>
    <filter-name>FilterB</filter-name>
    <servlet-name>RedServlet</servlet-name>
  </filter-mapping>

  <!-- associate FilterC to a request matching /red/* -->
  <filter-mapping>
    <filter-name>FilterC</filter-name>
    <url-pattern>/red/*</url-pattern>
  </filter-mapping>

  <!-- associate FilterD to a request matching /red/red/* -->
  <filter-mapping>
    <filter-name>FilterD</filter-name>
    <url-pattern>/red/red/*</url-pattern>
  </filter-mapping>

  <!-- associate FilterE to a request matching *.red -->
  <filter-mapping>
    <filter-name>FilterE</filter-name>
    <url-pattern>*.red</url-pattern>
  </filter-mapping>

  <servlet>
```

```

        <servlet-name>RedServlet</servlet-name>
        <servlet-class>RedServlet</servlet-class>
    </servlet>

    <servlet-mapping>
        <servlet-name>RedServlet</servlet-name>
        <url-pattern>/red/red/red/*</url-pattern>
    </servlet-mapping>
    <servlet-mapping>
        <servlet-name>RedServlet</servlet-name>
        <url-pattern>*.red</url-pattern>
    </servlet-mapping>
</web-app>

```

In `web.xml`, shown in listing 7.2, we have associated

- 1 `FilterA` and `FilterB` with `RedServlet` using the `servlet` name in the filter mapping elements
- 2 `FilterC` to a request whose URI matches `/red/*`
- 3 `FilterD` to a request whose URI matches `/red/red/*`
- 4 `FilterE` to a request whose URI matches `*.red`

We have also configured `RedServlet` to service requests having the URI pattern of `/red/red/red/*` and `*.red`.

Table 7.5 shows the order of filter invocations for various request URIs. We have not shown the base URIs for any of the request URIs in the table, since it is the same for all: `http://localhost:8080/chapter07/`.

Table 7.5 Order of filter invocation in filter chaining

Request URI	Filter Invocation Order	Reason		
		Request Serviced by RedServlet Because Of	Matching Filter Mappings with URL Pattern	Matching Filter Mappings with Servlet Name
aaa.red	FilterE, FilterA, FilterB	*.red	FilterE	FilterA, FilterB
red/aaa.red	FilterC, FilterE, FilterA, FilterB	*.red	FilterC, FilterE	FilterA, FilterB
red/red/aaa.red	FilterC, FilterD, FilterE, FilterA, FilterB	*.red	FilterC, FilterD, FilterE	FilterA, FilterB
red/red/red/aaa.red	FilterC, FilterD, FilterE, FilterA, FilterB	*.red and /red/red/red/*	FilterC, FilterD, FilterE,	FilterA, FilterB
red/red/red/aaa	FilterC, FilterD, FilterA, FilterB	/red/red/red/*	FilterC, FilterD	FilterA, FilterB

continued on next page

Table 7.5 Order of filter invocation in filter chaining (continued)

Request URI	Filter Invocation Order	Reason		
		Request Serviced by RedServlet Because Of	Matching Filter Mappings with URL Pattern	Matching Filter Mappings with Servlet Name
red/red/aaa	FilterC, FilterD	NONE (404 Error)	FilterC, Filter D	
red/aaa	FilterC	NONE (404 Error)	FilterC	
red/red/red/aaa.doc	FilterC, FilterD, FilterA, FilterB	/red/red/red/*	FilterC, FilterD	FilterA, FilterB
aaa.doc	None	NONE (404 Error)		

In table 7.5, observe the following points:

- The container will call the filters that match the request URI (`url-pattern`) before it calls the filters that match the servlet name to which the request will be delegated (`servlet-name`). Thus, `FilterC`, `FilterD`, and `FilterE` are always called before `FilterA` and `FilterB`.
- Whenever called, `FilterC`, `FilterD`, and `FilterE` are always called in this order since they are configured in this order in the `web.xml` file.
- Whenever `RedServlet` is invoked, `FilterA` and `FilterB` are called in this order, since they are configured in this order in the `web.xml` file.

We have provided this test application on the Manning web site. You can try it out with different request URIs and see the results.

7.4 ADVANCED FEATURES

In addition to monitoring the communication between the clients and the web application components, filters can manipulate the requests and alter the responses. In this section, we will learn about these features.

7.4.1 Using the request and response wrappers

All four wrapper classes—`ServletRequestWrapper`, `ServletResponseWrapper`, `HttpServletRequestWrapper`, and `HttpServletResponseWrapper`—work in the same way. They take a request or a response object in their constructor and delegate all the method calls to that object. This allows us to extend these classes and override any methods to provide a customized behavior.

In this section, we will use these classes in a filter to solve a simple problem. We have a legacy system that generates reports in a plain ASCII text format and stores them in a text file with an extension of `.txt`. We want these reports to be accessible from

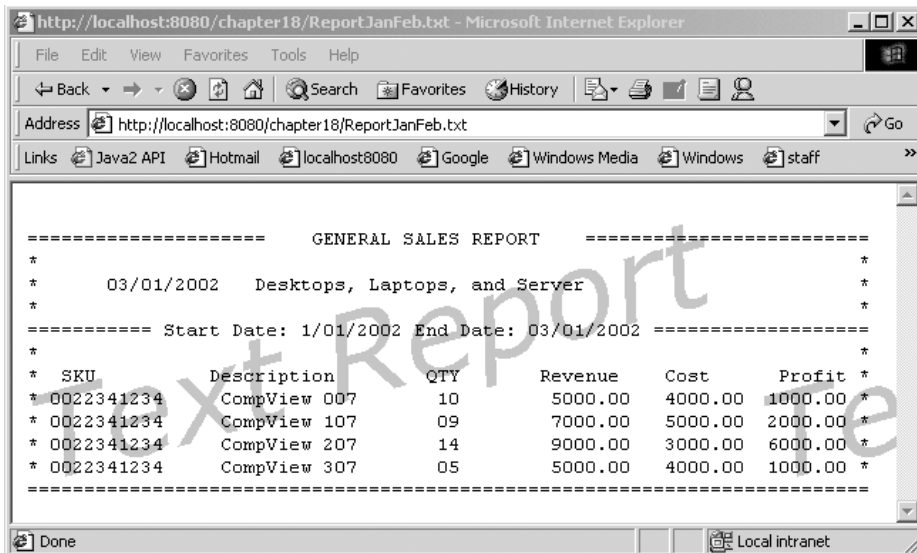


Figure 7.3 A sample report with a background image

browsers with an image displayed as the background of the report. For example, figure 7.3 shows how a sample report should display on the browser.

At the same time, we also do not want the browser to cache the report files. These two problems can be easily solved if we are able to do the following:

- 1 Embed the text of the report into `<html>` and `<body>` tags with an appropriate image as the background:

```

<html>
  <body background="textReport.gif">
    <pre>
      text of the report here.
    </pre>
  </body>
</html>

```

The background attribute of the `<body>` element will display the given image as the background of the report, while the `<pre>` tag will keep the formatting of the textual data intact.

- 2 Override the `If-Modified-Since` header. Browsers send this header so that the server can determine whether the resource needs to be sent. If the resource has not been modified after the period specified by the `If-Modified-Since` value, then the server does not send the resource at all.

For this purpose, we will filter all the requests for files with the extension `.txt`. Our filter will do two things:

- 1 Wrap the request into an `HttpServletRequestWrapper` and override the `getHeader()` method to return null for the `If-Modified-Since` header. A null value for this header ensures that the server does send the file.
- 2 Wrap the response object into an `HttpServletRequestResponseWrapper` so that the filter can modify the response and append the required HTML before sending it to the client.

Let's now look at the code that implements this. Listing 7.3 shows the code for `NonCachingRequestWrapper.java`, which customizes the behavior of `HttpServletRequestWrapper`.

Listing 7.3 Wrapping a request to hide a header value

```
import javax.servlet.*;
import javax.servlet.http.*;
public class NonCachingRequestWrapper extends HttpServletRequestWrapper
{
    public NonCachingRequestWrapper(HttpServletRequest req)
    {
        super(req);
    }

    public String getHeader(String name)
    {
        // hide only the If-Modified-Since header
        // and return the actual value for other headers
        if(name.equals("If-Modified-Since"))
        {
            return null;
        }
        else
        {
            return super.getHeader(name);
        }
    }
}
```

The code for `NonCachingRequestWrapper` is quite simple. It overrides the `getHeader()` method and returns null for the `If-Modified-Since` header. Since this class extends from `HttpServletRequestWrapper`, all other methods are delegated to the underlying request object that is passed in the constructor.

Listing 7.4 shows the code for `TextResponseWrapper`, which customizes the behavior of `HttpServletRequestResponseWrapper`.

Listing 7.4 Wrapping a response to buffer text data

```
import java.io.*;
import javax.servlet.*;
import javax.servlet.http.*;
```

```

public class TextResponseWrapper
    extends HttpServletResponseWrapper
{
    //This inner class creates a ServletOutputStream that
    //dumps everything that is written to it to a byte array
    //instead of sending it to the client.
    private static class ByteArrayServletOutputStream
        extends ServletOutputStream
    {
        ByteArrayOutputStream baos;
        ByteArrayServletOutputStream(ByteArrayOutputStream baos)
        {
            this.baos = baos;
        }

        public void write(int param) throws java.io.IOException
        {
            baos.write(param);
        }
    }

    //the actual ByteArrayOutputStream object that is used by
    //the PrintWriter as well as ServletOutputStream
    private ByteArrayOutputStream baos
        = new ByteArrayOutputStream();

    //This print writer is built over the ByteArrayOutputStream.
    private PrintWriter pw = new PrintWriter(baos);

    //This ServletOutputStream is built over the ByteArrayOutputStream.
    private ByteArrayServletOutputStream basos
        = new ByteArrayServletOutputStream(baos);

    public TextResponseWrapper(HttpServletResponse response)
    {
        super(response);
    }

    public PrintWriter getWriter()
    {
        //Returns our own PrintWriter that writes to a byte array
        //instead of returning the actual PrintWriter associated
        //with the response.
        return pw;
    }

    public ServletOutputStream getOutputStream()
    {
        //Returns our own ServletOutputStream that writes to a
        //byte array instead of returning the actual
        //ServletOutputStream associated with the response.

        return basos;
    }
}

```

```

byte[] toByteArray()
{
    return baos.toByteArray();
}
}

```

The code for `TextResponseWrapper` looks complicated but is actually very straightforward. It creates a `ByteArrayOutputStream` to store all the data that is written by the server. It overrides the `getWriter()` and `getOutputStream()` methods of `HttpServletResponse` to return the customized `PrintWriter` and `ServletOutputStream` that are built over the same `ByteArrayOutputStream`. Thus, no data is sent to the client.

Listing 7.5 shows the code for `TextToHTMLFilter.java`, which converts a textual report into a presentable HTML format as required.

Listing 7.5 Code for `TextToHTMLFilter.java`

```

import java.io.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class TextToHTMLFilter implements Filter
{
    private FilterConfig filterConfig;

    public void init(FilterConfig filterConfig)
    {
        this.filterConfig = filterConfig;
    }

    public void doFilter(
        ServletRequest request,
        ServletResponse response,
        FilterChain filterChain
    ) throws ServletException, IOException
    {
        HttpServletRequest req = (HttpServletRequest) request;
        HttpServletResponse res = (HttpServletResponse) response;

        NonCachingRequestWrapper ncrw
            = new NonCachingRequestWrapper( req );
        TextResponseWrapper trw = new TextResponseWrapper(res);

        //Passes on the wrapped request and response objects
        filterChain.doFilter(ncrw, trw);

        String top = "<html><body background=\"textReport.gif\"><pre>";
        String bottom = "</pre></body></html>";

        //Embeds the textual data into <html>, <body>, and <pre> tags.
        StringBuffer htmlFile = new StringBuffer(top);

```

```

String textFile = new String(trw.toByteArray());
htmlFile.append(textFile);
htmlFile.append("<br>" + bottom);

//Sets the content type to text/html
res.setContentType("text/html");

//Sets the content type to new length
res.setContentLength(htmlFile.length());

//Writes the new data to the actual PrintWriter
PrintWriter pw = res.getWriter();
    pw.println(htmlFile.toString());
}

public void destroy()
{
}
}

```

The code in listing 7.5 wraps the actual request and response objects into the `NonCachingRequestWrapper` and `TextResponseWrapper` objects, respectively, and then passes them on to the next component of the filter chain using the `doFilter()` method.

When the `filterChain.doFilter()` call returns, the text report is already written to the `TextResponseWrapper` object. Our filter retrieves the text data from this object and embeds it into the appropriate HTML tags. Finally, it writes the data to the actual `PrintWriter` object that sends the data to the client.

Deploying the filter

As explained earlier, deploying this filter requires two steps:

- 1 Copy all the class files to the `WEB-INF\classes` directory.
- 2 Set the filter and the filter mapping in the `web.xml` file of the web application as shown here:

```

<filter>
    <filter-name>TextToHTML</filter-name>
    <filter-class>TextToHTMLFilter</filter-class>
</filter>

<filter-mapping>
    <filter-name>TextToHTML</filter-name>
    <url-pattern>*.txt</url-pattern>
</filter-mapping>

```

If you have copied the `chapter07` directory to the `webapps` directory of your Tomcat installation, you can open the given `web.xml` file and see these settings.

Running the application

To run the application, restart Tomcat and request any text file that is available in this web application from the browser. For testing purposes, we have provided a sample text file and a GIF image. You can view them through this URL:

`http://localhost:8080/chapter07/ReportJanFeb.txt`

7.4.2 Important points to remember about filters

You need to understand these points when using filters:

There is one filter per `<filter>` entry in the `web.xml` file, per virtual machine. A servlet container is free to run multiple threads on the same filter object to service multiple requests simultaneously.

As of Servlet spec 2.4, filters can now be configured to be invoked on any combination of the following:

- As a result of calling `RequestDispatcher.forward()`
- As a result of calling `RequestDispatcher.include()`
- On error pages

The execution points used are controlled in the filter mapping in `web.xml`; if not specified, a filter will only be invoked on incoming requests. In the example below, all possible ways to invoke a filter have been enabled:

```
<filter-mapping>
  <filter-name>AccessLog</filter-name>
  <url-pattern>/*</url-pattern>
  <dispatcher>REQUEST</dispatcher>
  <dispatcher>FORWARD</dispatcher>
  <dispatcher>INCLUDE</dispatcher>
  <dispatcher>ERROR</dispatcher>
</filter-mapping>
```

7.4.3 Using filters with MVC

As we discussed in chapter 2, “Understanding JavaServer Pages,” the JSP Model 2 architecture is based on the MVC design pattern. It breaks a web application down into three distinct areas, where the JavaBeans act as the model, the JSP pages act as the view, and the servlets act as the controller. A request, or a group of related requests, is actually handled by a servlet, which retrieves the data and creates JavaBeans to hold the data. After creating the beans, it uses a `RequestDispatcher` object to forward the request to an appropriate JSP page. The JSP page uses the beans and generates the view.

This approach works well when the view to be displayed is determined by the business rules. For example, the page that should be displayed after a user logs in may depend on the access rights she has. Therefore, the final presentation depends on the servlet code that decides which JSP page the request is to be forwarded to.

Now consider another situation. An application is required to display some reports in either XML or HTML format as requested by the client. For this purpose, you develop two JSP pages: `xmlView.jsp` and `htmlView.jsp`. You also develop a servlet that generates the data needed by both the JSP pages. In the Model 2 architecture, the client request would go to this servlet. The servlet would retrieve the data and then dispatch the request to either `xmlView.jsp` or `htmlView.jsp`. The problem with this solution is that the servlet needs an extra parameter in the request that will indicate which view is being requested. Furthermore, the servlet will have to hard-code the names of the JSP pages. This means that adding a new view will require a code change in the servlet.

Filters are very useful in the situation described here. We can code the servlet's logic of retrieving the data and creating JavaBeans in a filter and apply this filter to both views. The user can directly request the `xmlView.jsp` or `htmlView.jsp` page. Since the filter will be executed first, the necessary beans will be available when the request reaches the JSP page. This architecture eliminates the need for an extra parameter to inform the filter about the view; this information is already present in the name of the resource given in the request. It also provides a clean way of adding a new view. For example, if we need to provide a text view, we can develop a `textView.jsp` file and apply the same filter on this JSP page. There is no code change in the filter. Thus, in such situations filters are a better choice for controllers than servlets.

7.5 SUMMARY

Filters add value to a web application by monitoring the requests and responses that are passed between the client and the server. With filters, we can analyze, manipulate, and redirect requests and responses.

In this chapter, we learned about the Filter API, including the three interfaces—`Filter`, `FilterConfig`, and `FilterChain`—and we saw how to configure filters in the deployment descriptor. We then looked at the way we can use a filter to wrap a request or a response in a customized object before delivering it to its destination. Within the MVC model, filters, rather than servlets, can better play the role of a controller in some situations.

7.6 REVIEW QUESTIONS

1. Which elements are allowed in the `<filter-mapping>` element of the deployment descriptor? (Select three)
 - a `<servlet-name>`
 - b `<filter-class>`
 - c `<dispatcher>`
 - d `<url-pattern>`
 - e `<filter-chain>`

2. What is wrong with the following code?

```
public void doFilter(ServletRequest req, ServletResponse, res,
FilterChain chain)
throws ServletException, IOException {

    chain.doFilter(req, res);
    HttpServletRequest request = (HttpServletRequest)req;
    HttpSession session = request.getSession();
    if (session.getAttribute("login") == null) {
        session.setAttribute("login", new Login());
    }
}
```

- a The doFilter() method signature is incorrect; it should take HttpServletRequest and HttpServletResponse.
- b The doFilter() method should also throw ServletException.
- c The call to chain.doFilter(req, res) should be this.doFilter(req, res, chain).
- d Accessing the request after chain.doFilter() results in an IllegalStateException.
- e Nothing is wrong with this filter.

3. Given these filter mapping declarations:

```
<filter-mapping>
  <filter-name>FilterOne</filter-name>
  <url-pattern>/admin/*</url-pattern>
  <dispatcher>FORWARD</dispatcher>
</filter-mapping>
<filter-mapping>
  <filter-name>FilterTwo</filter-name>
  <url-pattern>/users/*</url-pattern>
</filter-mapping>
<filter-mapping>
  <filter-name>FilterThree</filter-name>
  <url-pattern>/admin/*</url-pattern>
</filter-mapping>
<filter-mapping>
  <filter-name>FilterTwo</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>
```

In what order are the filters invoked for the following browser request?

/admin/index.jsp

- a FilterOne, FilterThree
- b FilterOne, FilterTwo, FilterThree
- c FilterThree, FilterTwo
- d FilterThree, FilterTwo
- e FilterThree
- f None of these filters are invoked.



CHAPTER 8

Session management

- | | |
|--|--------------------------------------|
| 8.1 Understanding state and sessions 120 | 8.4 Implementing session support 131 |
| 8.2 Using HttpSession 121 | 8.5 Summary 136 |
| 8.3 Understanding session timeout 130 | 8.6 Review questions 136 |

EXAM OBJECTIVES

- 4.1** Write servlet code to store objects into a session object and retrieve objects from a session object.
(Section 8.2)
- 4.2** Given a scenario,
- Describe the APIs used to access the session object,
 - Explain when the session object was created, and
 - Describe the mechanisms used to destroy the session object and when it was destroyed
- (Sections 8.2 and 8.3)
- 4.3** Using session listeners,
- Write code to respond to an event when an object is added to a session, and
 - Write code to respond to an event when a session migrates from one VM to another
- (Section 8.2.2)

4.4 Given a scenario,

- Describe which session management mechanism the Web container should employ,
- How cookies might be used to manage sessions,
- How URL rewriting might be used to manage sessions, and
- Write servlet code to perform URL rewriting

(Section 8.4)

INTRODUCTION

Since a web application is normally interacting with more than one user at the same time, it needs to remember each user and his history of transactions. The session provides this continuity by tracking the interaction between a user and the web application. To do well on the exam, you must know how to create and manage a session and how to associate it with its specific user.

In this chapter, we will discuss the session object, three of the session-related listener interfaces (we introduced the fourth session-related listener interface in chapter 6, “The servlet container model”), and the session timeout. We will also learn how to track the sessions using cookies and URL rewriting.

8.1 UNDERSTANDING STATE AND SESSIONS

The ability of a protocol to remember the user and her requests is called its *state*. From this perspective, protocols are divided into two types: stateful and stateless. In chapter 3, we observed that HTTP is a stateless protocol; each request to a web server and its corresponding response is handled as one isolated transaction.

Since all of the requests are independent and unrelated, the HTTP server has no way to determine whether a series of requests came from the same client or from different clients. This means that the server cannot maintain the state of the client between multiple requests; in other words, the server cannot remember the client.

In some cases, there may be no need to remember the client. For example, an online library catalog does not need to maintain the state of the client. While stateless HTTP may work well for this type of simple web browsing, the interaction between a client and a server in a web application needs to be stateful. A classic example is a shopping cart application. A user may add items and remove items from his shopping cart many times. At any time during the process, the server should be able to display the list of items in the cart and calculate their total cost. In order to do this, the server must track all of the requests and associate them with the user. We use a session to do this and turn stateless HTTP pages into a stateful Web application.

A session is an uninterrupted series of request-response interactions between a client and a server. For each request that is a part of this session, the server is able to identify the request as coming from the same client. A session starts when an unknown client sends the first request to the web application server. It ends when either the client explicitly ends the session or the server does not receive any requests from the client

within a predefined time limit. When the session ends, the server conveniently forgets the client as well as all the requests that the client may have made.

It should be made clear at this point that the first request from the client to the web application server may not be the very first interaction between that client and the server. By *first request*, we mean the request that requires a session to be created. We call it the first request because this is the request when the numbering of the requests starts (logically) and this is the request from which the server starts remembering the client. For example, a server can allow a user to browse a catalog of items without creating a session. However, as soon as the user logs in or adds an item to the shopping cart, it is clear that a session must be started.

So, how does a server establish and maintain a session with a client if HTTP does not provide any way to remember the client? There is only one way:

- When the server receives the first request from a client, the server initiates a session and assigns the session a unique identifier.
- The client must include this unique identifier with each subsequent request. The server inspects the identifier and associates the request with the corresponding session.

You may wonder why a server can't just look at the IP address of a request to identify a user. Many users access the Internet through a proxy server, in which case the server gets the IP address of the proxy server and not of the actual user, which makes the IP address nonunique for the set of users using that proxy server. For this reason, the server generates a unique identifier instead of relying on the IP address. This identifier is called a *session ID*, and the server uses this ID to associate the client's requests in a session. The exam requires you to understand the two most commonly used approaches for implementing session support: cookies and URL rewriting. We will discuss both of these approaches later in this chapter, in section 8.4.

In the following sections, we will see how the Servlet API helps us in implementing stateful web applications.

8.2 USING HTTPSESSION

The Servlet API abstracts the concept of session through the `javax.servlet.http.HttpSession` interface. This interface is implemented by the servlet container and provides a simple way to track the user's session.

A servlet container creates a new `HttpSession` object when it starts a session for a client. In addition to representing the session, this object acts as the data store for the information related to that session. In short, it provides a way to store data into memory and then retrieve it when the same user comes back later. Servlets can use this object to maintain the state of the session. As you'll recall, we discussed sharing data within the session scope using the `HttpSession` object in chapter 4, "The servlet model."

To put this in perspective, let's go back to the shopping cart example. The servlet container creates an `HttpSession` object for a user when the user logs in. The

servlet implementing the shopping cart application uses this object to maintain the list of items selected by the user. The servlet updates this list as the user adds or removes the items from his cart. Anytime the user wants to check out, the servlet retrieves the list of items from the session and calculates the total cost. Once the payment is made, the servlet closes the session; if the user sends another request, a new session is started.

Obviously, the servlet container creates as many `HttpSession` objects as there are sessions. In other words, there is an `HttpSession` object corresponding to each session (or user). However, we need not worry about associating the `HttpSession` objects with the users. The servlet container does that for us and, upon request, automatically returns the appropriate session object.

8.2.1 Working with an HttpSession

Using `HttpSession` is usually a three-step process:

- 1 Retrieve the session associated with the request.
- 2 Add or remove name-value pairs of attributes from the session.
- 3 Invalidate the session if required.

Often the client offers no indication that it is ending the session. For example, a user may browse to another site and may not return for a long time. In this case, the server will never know whether the user has ended her session. To help us in such situations, the servlet container automatically invalidates the session after a certain period of inactivity. The period is, of course, configurable through the deployment descriptor and is known as the *session timeout* period. We will learn more about it in section 8.3.

Listing 8.1 contains the `doPost()` method of an imaginary `ShoppingCartServlet`. This listing illustrates a common use of `HttpSession`.

Listing 8.1 Using HttpSession methods

```
//code for the doPost() method of ShoppingCartServlet
public void doPost(HttpServletRequest req,
                   HttpServletResponse res)
{
    HttpSession session = req.getSession(true); ← Retrieves the session
    List listofItems =
        (List) session.getAttribute("listofitems"); ← Retrieves an attribute
    if(listofItems == null)                               from the session
    {
        listofItems = new ArrayList();
        session.setAttribute("listofitems", listofItems); ← Sets an attribute
    }                                                       in the session

    String itemcode = req.getParameter("itemcode");
    String command = req.getParameter("command");
        if("additem".equals(command) )
    {
        listofItems.add(itemcode);
    }
}
```

```

    }
    else if ("removeitem".equals(command) )
    {
        listOfItems.remove(itemcode);
    }
}

```

In listing 8.1, we first get a reference to the `HttpSession` object using `req.getSession(true)`. This will create a new session if a session does not already exist for the user. The `HttpServletRequest` interface provides two methods to retrieve the session, as shown in table 8.1.

Table 8.1 `HttpServletRequest` methods for retrieving the session

Method	Description
<code>HttpSession getSession(boolean create)</code>	This method returns the current <code>HttpSession</code> associated with this request, or if there is no current session and the <code>create</code> parameter is true, then it returns a new session.
<code>HttpSession getSession()</code>	This method is equivalent to calling <code>getSession(true)</code> .

Notice that we have not written any code to identify the user. We just call the `getSession()` method and assume that it will return the same `HttpSession` object each time we process a request from a specific user. It is the job of the implementation of the `getSession()` method to analyze the request and find the right `HttpSession` object associated with the request. We have used `getSession(true)` just to emphasize that we want to create a new session if it does not exist, although `getSession()` would have the same effect. A session will not be available for the first request sent by a user. In that case, the user is a new client and a new session will be created for her.

After retrieving the session, we get the list of items from the session. We use the `HttpSession` methods shown in table 8.2 to set and get the `listofitems` attribute that stores the item codes in the session. If this is a new session, or if this is the first time the user is adding an item, the `session.getAttribute()` will return null, in which case we create a new `List` object and add it to the session. Then, based on the `command` and `itemCode` request parameters, we either add or remove the item from the list.

Table 8.2 `HttpSession` methods for setting/getting attributes

Method	Description
<code>void setAttribute(String name, Object value)</code>	This method adds the passed object to the session, using the name specified.
<code>Object getAttribute(String name)</code>	This method returns the object bound with the specified name in this session, or null if no object is bound under the name.

Quizlet

- Q:** Listing 8.1 shows the `doPost()` method of `ShoppingCartServlet`. As you can see, this servlet only maintains the list of item codes. However, once a user has finished selecting the items, he needs to complete the process by “checking out.” Is it possible to implement this functionality using a different servlet instead of adding the functionality to `ShoppingCartServlet`? Can you retrieve the `listOfItems` attribute associated with a user from another servlet?
- A:** Definitely. An `HttpSession` object associated with a user is accessible from all of the components of a web application (servlets and JSP pages) during the time that the components are serving a request from that user. In this case, we can have a hyperlink named *Check Out* on our page that refers to a different servlet named `CheckoutServlet`. This servlet can access the session information and retrieve the `listOfItems` attribute, as shown here:

```
//code for the doGet() method of CheckoutServlet
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
{
    HttpSession session = req.getSession();

    List listOfItems =
        (List) session.getAttribute("listOfItems");

    //process the listOfItems.
}
```

As we mentioned earlier, the `HttpServletRequest.getSession()` method will return the correct session object for a specific user.

8.2.2 Handling session events with listener interfaces

As we saw in chapter 6, listener interfaces are a way to receive notifications when important events occur in a web application. To receive notification of an event, we need to write a class that implements the corresponding listener interface. The servlet container then calls the appropriate methods on the objects of this class when the events occur.

The Servlet API defines four listeners and two events related to the session in the `javax.servlet.http` package:

- `HttpSessionAttributeListener` and `HttpSessionBindingEvent`
- `HttpSessionBindingListener` and `HttpSessionBindingEvent`
- `HttpSessionListener` and `HttpSessionEvent`
- `HttpSessionActivationListener` and `HttpSessionEvent`.

All four listener interfaces extend `java.util.EventListener`. `HttpSessionEvent` extends `java.util.EventObject` and `HttpSessionBindingEvent` extends `HttpSessionEvent`.

HttpSessionAttributeListener

The `HttpSessionAttributeListener` interface allows a developer to receive notifications whenever attributes are added to, removed from, or replaced in the attribute list of any of the `HttpSession` objects of the web application. We specify the class that implements this interface in the deployment descriptor. We discussed this listener (`javax.servlet.http.HttpSessionAttributeListener`) in detail in chapter 6. Now let's look at the three other listeners.

HttpSessionBindingListener

The `HttpSessionBindingListener` interface is implemented by the classes whose objects need to receive notifications whenever they are added to or removed from a session. We do not have to inform the container about such objects explicitly via the deployment descriptor. Whenever an object is added to or removed from any session, the container introspects the interfaces implemented by that object. If the object implements the `HttpSessionBindingListener` interface, the container calls the corresponding notification methods shown in table 8.3.

Table 8.3 `HttpSessionBindingListener` methods for receiving notification of a change in the attribute list of `HttpSession`

Method	Description
<code>void valueBound (HttpSessionBindingEvent event)</code>	Notifies the object that it is being bound to a session
<code>void valueUnbound (HttpSessionBindingEvent event)</code>	Notifies the object that it is being unbound from a session

The servlet container calls the interface methods even if the session is explicitly invalidated or has timed out. Listing 8.2 illustrates the use of this interface to log `HttpSessionBindingEvents`.

Listing 8.2 Implementing `HttpSessionBindingListener`

```
import javax.servlet.*;
import javax.servlet.http.*;

//An entry will be added to the log file whenever objects of
//this class are added to or removed from a session.
public class CustomAttribute implements
    HttpSessionBindingListener
{
    public Object theValue;
```



```

public void valueBound(HttpSessionBindingEvent e)
{
    HttpSession session = e.getSession();
    session.getServletContext().log("CustomAttribute "+
        theValue+"bound to a session");
}

public void valueUnbound(HttpSessionBindingEvent e)
{
    HttpSession session = e.getSession();
    session.getServletContext().log("CustomAttribute "+
        theValue+" unbound from a session.");
}
}

```

Retrieves the session associated with this event

In the example code in listing 8.2, we retrieve the session object from `HttpSessionBindingEvent`. From the session object, we retrieve the `ServletContext` object and log the messages using the `ServletContext.log()` method.

You may wonder what the difference is between `HttpSessionAttributeListener` and `HttpSessionBindingListener`, since both are used for listening to changes in the attribute list of a session. The difference is that `HttpSessionAttributeListener` is configured in the deployment descriptor and the container creates only one instance of the specified class. `HttpSessionBindingEvents` generated from all the sessions are sent to this object. On the other hand, `HttpSessionBindingListener` is not configured in the deployment descriptor. The servlet container calls methods on an object implementing this interface only if that object is added to or removed from a session. While the `HttpSessionAttributeListener` interface is used to track the activity of all the sessions on an application level, the `HttpSessionBindingListener` interface is used to take actions when certain kinds of objects are added to or removed from a session.

HttpSessionListener

The `HttpSessionListener` interface is used to receive notifications when a session is created or destroyed. The listener class implementing this interface must be configured in the deployment descriptor. It has the methods shown in table 8.4.

Table 8.4 `HttpSessionListener` methods for receiving notification when a session is created or destroyed

Method	Description
<code>void sessionCreated(HttpSessionEvent se)</code>	Called when a session is created
<code>void sessionDestroyed(HttpSessionEvent se)</code>	Called when a session is destroyed

This interface can be used to monitor the number of active sessions, as demonstrated in the implementation of the `HttpSessionListener` interface in listing 8.3.

Listing 8.3 Counting the number of sessions

```
import javax.servlet.http.*;

public class SessionCounter implements HttpSessionListener
{
    private static int activeSessions = 0;

    public void sessionCreated(HttpSessionEvent evt)
    {
        activeSessions++;
        System.out.println("No. of active sessions on:"+
            new java.util.Date()+" : "+activeSessions);
    }

    public void sessionDestroyed (HttpSessionEvent evt)
    {
        activeSessions--;
    }
}
```

In listing 8.3, we increment the session count when a new session is created in the `sessionCreated()` method and decrement the session count when any session is invalidated in the `sessionDestroyed()` method.

At first, it appears that this interface also might be very useful to clean up a user's transient data from the database when that user's session ends. For example, as soon as a user logs in, we can set the user ID in the session. When the user logs out or when the session times out, we can use the `sessionDestroyed()` method to clean up the database, as shown in listing 8.4.

Listing 8.4 Incorrect use of HttpSessionListener

```
import javax.servlet.*;
import javax.servlet.http.*;

public class BadSessionListener implements
    HttpSessionListener
{
    public void sessionCreated(HttpSessionEvent e)
    {
        //can't do much here as the session is just created and
        //does not contain anything yet, except the sessionid
        System.out.println("Session created: "+
            e.getSession().getId());
    }

    public void sessionDestroyed(HttpSessionEvent e)
    {
        HttpSession session = e.getSession();
        String userid = (String) session.getAttribute("userid");
```

**Will
not
work!**

```

        //delete user's transient data from the database
        //using the userid.
    }
}

```

In listing 8.4, the line `session.getAttribute("userid");` will not work because the servlet container calls the `sessionDestroyed()` method after the session is invalidated. Therefore, a call to `getAttribute()` will throw an `IllegalStateException`.

So, how do we solve our problem of cleaning up the database when a session is invalidated? The solution is a little cumbersome. We will create a class that wraps the user ID and implements the `HttpSessionBindingListener` interface. When the user logs in, for instance through a `LoginServlet`, instead of setting the user ID directly in the session, we will set this wrapper in the session. The servlet container will call `valueUnbound()` on the wrapper as soon as the session is invalidated. We will use this method to clean up the database. Listing 8.5 illustrates the process.

Listing 8.5 Cleaning up the database using `HttpSessionBindingListener`

```

import javax.servlet.*;
import javax.servlet.http.*;

public class UseridWrapper implements HttpSessionBindingListener
{
    public String userid = "default";
    public UseridWrapper(String id)
    {
        this.userid = id;
    }
    public void valueBound(HttpSessionBindingEvent e)
    {
        //insert transient user data into the database
    }

    public void valueUnbound(HttpSessionBindingEvent e)
    {
        //remove transient user data from the database
    }
}

```

The following code for the `doPost()` method of `LoginServlet` shows the use of the `UseridWrapper` class:

```

//code for doPost() of LoginServlet
public void doPost(HttpServletRequest req, HttpServletResponse res)
{
    String userid = req.getParameter("userid");
    String password = req.getParameter("password");
}

```

```

boolean valid = //validate the userid/password.
if(valid)
{
    UseridWrapper useridwrapper = new UseridWrapper(userid);
    req.getSession().setAttribute("useridwrapper", useridwrapper); ←
}
else
{
    //forward the user to the login page.
}
...
...
}

```

Sets the UseridWrapper object in the session

HttpSessionActivationListener

This interface is used by the session attributes to receive notifications when a session is being migrated across the JVMs in a distributed environment. This interface declares two methods, as shown in table 8.5.

Table 8.5 HttpSessionActivationListener methods for receiving activation/passivation notification in a distributed environment

Method	Description
void sessionDidActivate(HttpSessionEvent se)	Called just after the session is activated
void sessionWillPassivate(HttpSessionEvent se)	Called when the session is about to be passivated

We will not discuss this interface in detail since it is rarely used and is not required for the exam.

Quizlet

Q: Which interface would you use to achieve the following?

- 1 You want to listen to the HttpSessionBindingEvents but none of your session attributes implement the HttpSessionBindingListener interface.
- 2 You want to monitor the average time users are logged into your web application.

A:

- 1 Use HttpSessionAttributeListener. Remember, you will have to configure it in the deployment descriptor.
- 2 Use HttpSessionListener. You can use the sessionCreated() and sessionDestroyed() methods to calculate how long a user has been logged in.

8.2.3 Invalidating a Session

We observed at the beginning of this chapter that a session is automatically terminated when the user remains inactive for a specified period of time. In some cases, we may also want to end the session programmatically. For instance, in our shopping cart example, we would want to end the session after the payment process is complete so that if the user sends another request, a new session is started with no items in the shopping cart. `HttpSession` provides the method shown in table 8.6 for invalidating a session.

Table 8.6 `HttpSession` method for invalidating a session

Method	Description
<code>void invalidate()</code>	This method invalidates this session and then unbinds any objects bound to it. This means that the <code>valueUnbound()</code> method will be called on all of its attributes that implement <code>HttpSessionBindingListener</code> . It throws an <code>IllegalStateException</code> if the session is already invalidated.

Listing 8.6 shows `LogoutServlet`'s `doGet()` method. It uses the `invalidate()` method to expunge a session.

Listing 8.6 Using `HttpSession.invalidate()` to expunge a session

```
//code for doGet() of LogoutServlet
//This method will be invoked if a user clicks on
//a "Logout" button or hyperlink.
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
{
    ...
    req.getSession().invalidate(); ← Expunges the session
    //forward the user to the main page.
    ...
}
```

8.3 UNDERSTANDING SESSION TIMEOUT

Since the HTTP protocol does not provide any signal for the termination of the session to the server, if the user does not click on some kind of a logout button or hyperlink, the only way to determine whether a client is active or not is to observe the inactivity period. If a user does not perform any action for a certain period of time, the server assumes the user to be inactive and invalidates the session. The `web.xml` in listing 8.7 shows the configuration of the timeout period of a session.

Listing 8.7 Configuring session timeout in web.xml

```
<web-app>
...
<session-config>
  <session-timeout>30</session-timeout> ← Sets timeout to 30 minutes
</session-config>
...
</web-app>
```

The `<session-timeout>` element contains the timeout in minutes. A value of 0 or less means that the session will never expire. The `HttpSession` interface provides the two methods shown in table 8.7 for getting and setting the timeout value of a session.

Table 8.7 `HttpSession` methods for getting/setting the timeout value of a session

Method	Description
<code>void setMaxInactiveInterval(int seconds)</code>	This method specifies the number of seconds between client requests before the servlet container will invalidate this session. A negative value means that the session will never expire.
<code>int getMaxInactiveInterval()</code>	This method returns the maximum time interval, in seconds, that the servlet container will keep this session open between client accesses.

It is important to note that `setMaxInactiveInterval()` affects only the session on which it is called. Other sessions will still have the same timeout period as specified in the deployment descriptor.

NOTE There are two inconsistencies in the way the `session-timeout` tag of the deployment descriptor and the `setMaxInactiveInterval()` method of `HttpSession` work:

- 1 The `session-timeout` value is specified in minutes, while the `setMaxInactiveInterval()` method accepts seconds.
- 2 A `session-timeout` value of 0 or less means that the session will never expire, while if we want to specify that a session will never expire using the `setMaxInactiveInterval()` method, a negative value (not 0) is required.

8.4 IMPLEMENTING SESSION SUPPORT

We have seen how storing attributes in the session object enables us to maintain the state of the application. In this section, we will see how a servlet container associates incoming requests with an appropriate `HttpSession` object. As we observed at the

beginning of this chapter, the way containers provide support for HTTP sessions is to identify each client with a unique ID, which is called a session ID, and force the client to send that ID to the server with each request. Let's go over in detail the steps that a client and a server must take to track a session:

- 1 A new client sends a request to a server. Since this is the first request, it does not contain any session ID.
- 2 The server creates a session and assigns it a new session ID. At this time, the session is said to be in the *new* state. We can use `session.isNew()` to determine whether the session is in this state or not. The server then sends the ID back to the client with the response.
- 3 The client gets the session ID and saves it for future requests. This is the first time the client is aware of the existence of a session on the server.
- 4 The client sends another request, and this time, it sends the session ID with the request.
- 5 The server receives the request and observes the session ID. It immediately associates the request with the session that it had created earlier. At this time, the client is said to have *joined* the session, which means the session is no longer in the new state. Therefore, a call to `session.isNew()` will return false.

Steps 3–5 keep repeating for the life of the session. If the client does not send any requests for a length of time that exceeds the session timeout, the server invalidates the session. Once the session is invalidated, either programmatically or because it has timed out, it cannot be resurrected even if the client sends the same session ID again. After that, as far as the server is concerned, the next request from the client is considered to be the first request (as in step 1) that cannot be associated with an existing session. The server will create a new session for the client and will assign it a new ID (as in step 2).

In the following section, we look at two techniques—cookies and URL rewriting—that a servlet container uses to implement the steps described above to provide session support.

8.4.1 Supporting sessions using cookies

In this technique, to manage the sending and receiving of session IDs, the servlet container uses HTTP headers. As we saw in chapter 3, all HTTP messages—requests as well as responses—contain header lines. While sending a response, a servlet container adds a special header line containing the session ID. The container adds this header line transparently to the servlet developer. The client, which is usually a browser, receives the response, extracts the special header line, and stores it on the local machine. The browser does this transparently to the user. While sending another request, the client automatically adds a header line containing the stored session ID.

The header line that is stored by the browser on the user's machine is called a *cookie*. If you recall the discussion about HTTP headers, a header line is just a name-value pair. Not surprisingly, the header name used for sending a cookie is *cookie*. A sample HTTP request containing a cookie looks like this:

```
POST /servlet/testServlet HTTP/1.1
User-Agent= MOZILLA/1.0
cookie=jsessionId=61C4F23524521390E70993E5120263C6
Content-Type: application/x-www.formurlencoded
userid=john
```

← Header line for the cookie

The value of the cookie header shown above is

```
jsessionid=61C4F23524521390E70993E5120263C6
```

This technique was developed by Netscape and was adopted by all other browsers. Back in the early days of the Internet, cookies were only used to keep the session ID. But later on, companies started using cookies to store a lot of other information, such as user IDs, preferences, and so forth. They also started using cookies to track the browsing patterns of the users. Since the cookie management happens behind the scenes, very soon cookies became known as a potential security hazard, and many users started disliking them. Although most users still enable cookies in their browsers, some corporate policies now disable them. When cookies are disabled, the browser ignores any cookie header lines that are present in the HTTP responses, and consequently does not send any cookie header lines in the requests.

For some web sites, session support is extremely important, and so they cannot rely solely on cookies. In such cases, we need to use another technique that will work even if the users disable cookies. We examine this technique in the next section.

8.4.2 Supporting sessions using URL rewriting

In the absence of cookie support, we can attach the session ID to all of the URLs that are within an HTML page that is being sent as a response to the client. That way, when the user clicks on one of the URLs, the session ID is automatically sent back to the server as a part of the request line itself, instead of as a header line.

To better understand this, consider the following HTML page code returned by an imaginary servlet named `HomeServlet`:

```
<html>
<head></head>

<body>
A test page showing two URLs:<br>
<a href="/servlet/ReportServlet">First URL</a><br>
<a href="/servlet/AccountServlet">Second URL</a><br>
</body>
</html>
```


The above HTML page is a normal HTML page without any special code. However, if the cookies are disabled, the session ID will not be sent when the user clicks on the hyperlink displayed by this page. Now, let's see the same HTML code but with the URLs rewritten to include the session ID:

```
<html>
<head></head>
<body>
A test page showing two URLs:<br>
<a href=
  "/servlet/ReportServlet;jsessionid=C084B32241B2F8F060230440C0158114">
  View Report</a><br>
<a href=
  "/servlet/AccountServlet;jsessionid=C084B32241B2F8F060230440C0158114">
  View Account</a><br>
</body>
</html>
```

When the user clicks on the URLs displayed by the above page, the session ID will be sent as a part of the request line. We do not need cookies to do this. Although it is quite easy to attach the session ID with all the URLs, unlike cookies, it is not transparent to the servlet developer. The `HttpServletResponse` interface provides two methods for this purpose, as shown in table 8.8.

Table 8.8 `HttpServletResponse` methods for appending session IDs to the URLs

Method	Description
String <code>encodeURL(String url)</code>	This method returns the URL with the session ID attached. It is used for normal URLs emitted by a servlet.
String <code>encodeRedirectURL(String url)</code>	This method returns the URL with the session ID attached. It is used for encoding a URL that is to be used for the <code>HttpServletResponse.sendRedirect()</code> method.

Both methods first check to see if attaching the session ID is necessary. If the request contains a cookie header line, then cookies are enabled and the method need not rewrite the URL. In this case, the URL is returned without the session ID attached to it.

NOTE Observe that `jsessionid` is appended to the URL using a `;` and not a `?`. This is because `jsessionid` is a part of the path info of the request URL. It is not a request parameter and thus cannot be retrieved using the `getParameter("jsessionid")` method of `ServletRequest`.

Listing 8.8 illustrates how these methods can be used. `HomeServlet` generates the HTML page shown earlier.

Listing 8.8 Using URL rewriting to implement session support

```
import javax.servlet.*;
import javax.servlet.http.*;

public class HomeServlet extends HttpServlet
{
    public void doGet(HttpServletRequest req,
                      HttpServletResponse res)
    {
        HttpSession s = req.getSession();    ← Gets the session
        PrintWriter pw = res.getWriter();
        pw.println("<html>");
        pw.println("<head></head>");

        pw.println("<body>");
        pw.println("A test page showing two URLs:<br>");
        pw.println("<a href=\"\"
                  + res.encodeURL(\"/servlet/ReportServlet\")
                  + \"\">View Report</a><br>");
        pw.println("<a href=\"\"
                  + res.encodeURL(\"/servlet/AccountServlet\")
                  + \"\">View Account</a><br>");
        pw.println("</body>");
        pw.println("</html>");
    }
}
```

**Appends the
session ID**

Observe that the process of retrieving the session in the servlet remains the same. We can still safely call the `getSession()` methods to retrieve the session. The servlet container transparently parses the session ID attached to the requested URL and returns the appropriate session object.

In general, URL rewriting is a very robust way to support sessions. We should use this approach whenever we are uncertain about cookie support. However, it is important to keep the following points in mind:

- We should encode all the URLs, including all the hyperlinks and action attributes of the forms, in all the pages of the application.
- All the pages of the application should be dynamic. Because different users will have different session IDs, there is no way to attach proper session IDs to the URLs present in static HTML pages.
- All the static HTML pages must be run through a servlet, which would rewrite the URLs while sending the pages to the client. Obviously, this can be a serious performance bottleneck.

Quizlet

- Q:** You have developed your web application assuming that your clients support cookies. However, after deploying the application, you realize that most of your clients have disabled cookies. What will be the impact on your application? How can you fix it?
- A:** The impact will be drastic. The application will not be able to maintain the user's state. The servlet container will create a new session for each request from each user. The only way to fix this problem is to modify your servlet code to incorporate URL rewriting.

8.5 SUMMARY

A web application needs to impose state upon the inherently stateless HTTP protocol in order to keep track of a client's interactions with the server. This is done through session management using the `HttpSession` object. A session is a complete series of interactions between a client and the server; during a session, the server "remembers" the client and associates all the requests from that client with the client's unique session object.

We use listener interfaces to receive notifications when important events take place in the session and to initiate appropriate actions. These events include changes to the session attribute list and creating and destroying the session.

The server implements a session by assigning a unique identifier to it. Using either cookies or URL rewriting, this session ID is sent to the client in the response and returned to the server with each subsequent request. The session ends either when it times out or when the session is invalidated. The "session timeout" period is configurable through the deployment descriptor. This affects the timeout period of all the sessions. To change the timeout of a specific session, we can use `HttpSession.setMaxInactiveInterval(int seconds)` method.

At this point, you should be able to answer exam questions based on the semantics of `HttpSession` and the interfaces that are used to listen for changes in an `HttpSession`. You should also be able to answer questions based on session management using cookies and URL rewriting.

In the next chapter, we will look at another topic that is important from the perspective of the exam: security.

8.6 REVIEW QUESTIONS

1. Which of the following interfaces or classes is used to retrieve the session associated with a user? (Select one)
 - a `GenericServlet`
 - b `ServletConfig`
 - c `ServletContext`
 - d `HttpServlet`

- e HttpServletRequest
 - f HttpServletResponse
2. Which of the following code snippets, when inserted in the doGet () method, will correctly count the number of GET requests made by a user? (Select one)
 - a

```
HttpSession session = request.getSession();
int count = session.getAttribute("count");
session.setAttribute("count", count++);
```
 - b

```
HttpSession session = request.getSession();
int count = (int) session.getAttribute("count");
session.setAttribute("count", count++);
```
 - c

```
HttpSession session = request.getSession();
int count = ((Integer) session.getAttribute("count")).intValue();
session.setAttribute("count", count++);
```
 - d

```
HttpSession session = request.getSession();
int count = ((Integer) session.getAttribute("count")).intValue();
session.setAttribute("count", new Integer(++count));
```
 3. Which of the following methods will be invoked on a session attribute that implements HttpSessionBindingListener when the session is invalidated? (Select one)
 - a sessionDestroyed
 - b valueUnbound
 - c attributeRemoved
 - d sessionInvalidated
 4. Which of the following methods will be invoked on a session attribute that implements appropriate interfaces when the session is invalidated? (Select one)
 - a sessionDestroyed Of HttpSessionListener
 - b attributeRemoved Of HttpSessionAttributeListener
 - c valueUnbound Of HttpSessionBindingListener
 - d sessionWillPassivate Of HttpSessionActivationListener
 5. Which of the following methods will expunge a session object? (Select one)
 - a session.invalidate();
 - b session.expunge();
 - c session.destroy();
 - d session.end();
 - e session.close();
 6. Which of the following method calls will ensure that a session will never be expunged by the servlet container? (Select one)
 - a session.setTimeout(0);
 - b session.setTimeout(-1);

- c** `session.setTimeout(Integer.MAX_VALUE);`
 - d** `session.setTimeout(Integer.MIN_VALUE);`
 - e** None of these
7. How can you make sure that none of the sessions associated with a web application will ever be expunged by the servlet container? (Select one)
- a** `session.setMaxInactiveInterval(-1);`
 - b** Set the session timeout in the deployment descriptor to `-1`.
 - c** Set the session timeout in the deployment descriptor to `0` or `-1`.
 - d** Set the session timeout in the deployment descriptor to `65535`.
 - e** You have to change the timeout value of all the sessions explicitly as soon as they are created.
8. In which of the following situations will a session be invalidated? (Select two)
- a** No request is received from the client for longer than the session timeout period.
 - b** The client sends a `KILL_SESSION` request.
 - c** The servlet container decides to invalidate a session due to overload.
 - d** The servlet explicitly invalidates the session.
 - e** A user closes the active browser window.
 - f** A user closes all of the browser windows.
9. Which method is required for using the URL rewriting mechanism of implementing session support? (Select one)
- a** `HttpServletRequest.encodeURL()`
 - b** `HttpServletRequest.rewriteURL()`
 - c** `HttpServletResponse.encodeURL()`
 - d** `HttpServletResponse.rewriteURL()`
10. The users of your web application do not accept cookies. Which of the following statements are correct? (Select one)
- a** You cannot maintain client state.
 - b** URLs displayed by static HTML pages may not work properly.
 - c** You cannot use URL rewriting.
 - d** You cannot set session timeout explicitly.



CHAPTER 9

Developing secure web applications

- | | |
|---|--|
| 9.1 Basic concepts 140 | 9.4 Securing web applications programmatically 156 |
| 9.2 Understanding authentication mechanisms 142 | 9.5 Summary 158 |
| 9.3 Securing web applications declaratively 149 | 9.6 Review questions 159 |

EXAM OBJECTIVES

- 5.1** Based on the servlet specification, compare and contrast the following security mechanisms:
- authentication,
 - authorization,
 - data integrity, and
 - confidentiality.
- (Section 9.2)
- 5.2** In the deployment descriptor, declare
- A security constraint,
 - A Web resource,
 - The transport guarantee,
 - The login configuration, and
 - A security role.
- (Sections 9.2 and 9.3)

- 5.3** Compare and contrast the authentication types (BASIC, DIGEST, FORM, and CLIENT-CERT); describe how the type works; and given a scenario, select an appropriate type.
(Section 9.2)

INTRODUCTION

The utilization of the Internet as an essential business tool continues to grow, as more and more companies are web-enabling their operations. It is increasingly common for all types of business dealings to take place over the Internet. Currently, millions of people transmit personal information over the Internet as they shop at online stores. Many business transactions, such as banking, stock trading, and so forth, are conducted online each day. To support these applications, we need a robust security mechanism in place. It is not an overstatement to say that e-commerce is not possible without security.

In this chapter, we will learn about the various techniques that are used to make a web application secure.

9.1 BASIC CONCEPTS

The importance of web security will continue to increase as companies and individuals alike are paying more attention to ensuring that their resources are protected and their interactions are private. The Servlet specification provides methods and guidelines for implementing security in web applications, but before we go into the details of implementing those security features, let's look at some terms you need to know for the exam.

9.1.1 Authentication

The first fundamental requirement of security is to authenticate the user. *Authentication* is the process of identifying a person—or even a system, such as an application—and validating their credentials. It means verifying that the user is who she (or it) claims to be. For example, a traveler must show a passport before boarding a flight. This ID authenticates the traveler; it provides his credentials. In the Internet world, the basic credentials that authenticate a user are typically a username and a password.

9.1.2 Authorization

Once the user has been authenticated, she must be authorized. *Authorization* is the process of determining whether a user is permitted to access a particular resource that she has requested. For example, you will not be permitted to access a bank account that does not belong to you, even if you are a member of the bank. In short, you are not *authorized* to access anyone else's account. Authorization is usually enforced by maintaining an *access control list (ACL)*; this list specifies the users and the types of access they have to resources.

9.1.3 Data integrity

Data integrity is the process of ensuring that the data is not tampered with while in transit from the sender to the receiver. For example, if you send a request to transfer \$1000 from your account to another account, the bank should get a transfer request for \$1000 and not \$10,000. Data integrity is usually ensured by sending a hashcode or signature of the data along with the data. At the receiving end, the data and its hashcode are verified.

9.1.4 Confidentiality or data privacy

Confidentiality is the process of ensuring that no one except the intended user is able to access sensitive information. For example, sometimes when you send your user ID/password to log onto a web site, the information travels in plain text across the Internet. It is possible for hackers to access this information by sniffing the HTTP packets. In this case, the data is not confidential. Confidentiality is usually ensured by encrypting the information so that only the intended user can decrypt it. Today, most web sites use the HTTPS protocol to encrypt messages so that even if a hacker sniffs the data, he will not be able to decrypt it and hence cannot use it.

The difference between authorization and confidentiality is in the way the information is protected. Authorization prevents the information from reaching unintended parties in the first place, while confidentiality ensures that even if the information falls into the wrong hands, it remains unusable.

9.1.5 Auditing

Auditing is the process of recording security-related events taking place in the system in order to be able to hold users accountable for their actions. Auditing can help determine the cause of a breach, and is usually accomplished by maintaining the log files generated by the application.

9.1.6 Malicious code

A piece of code that is meant to cause harm to computer systems is called *malicious code*. This includes viruses, worms, and Trojan horses. Besides the threat from the outside, sometimes in-house developers leave a back door open into the software that they write, which provides a potential opportunity for misuse. Although we cannot prevent unknown programmers from writing malicious code, companies can definitely prevent malicious code from being written in-house by conducting peer-to-peer code reviews.

9.1.7 Web site attacks

Anything that is deemed valuable is a potential target for attacks and should be protected. Web sites are no exception. Their value lies in the information they contain or the services they provide to legitimate users. A web site may be attacked by different people for different reasons. For example, a hacker may attack for pleasure, a terminated

employee may attack for revenge, or a professional thief may attack for the purpose of stealing credit card numbers.

Broadly, there are three types of web site attacks:

- *Secrecy attacks*—Attempts to steal confidential information by sniffing the communications between two machines. Encrypting the data being transmitted can prevent such attacks. For example, it is a universal standard that financial institutions use HTTPS in online banking, stock trading, and so forth.
- *Integrity attacks*—Attempts to alter information in transit with malicious intent. If these attempts succeed, it will compromise the data integrity. IP spoofing is one of the common techniques used in integrity attacks. In this technique, the intruder sends messages to a server with an IP address indicating that the message is coming from a trusted machine. The server is thus fooled into giving access to the intruder. Such attacks can be prevented by using strong authentication techniques, such as public-key cryptography.
- *Denial-of-service attacks (or availability attacks)*—Attempts to flood a system with fake requests so that the system remains unavailable for legitimate requests. Creating network congestion by sending spurious data packets also comes under this category. Such attacks can be prevented by using firewalls that block network traffic on unintended ports.

Quizlet

- Q:** The process of showing your ID card to the security guard before entering a building is known as what?
- A:** *Authentication* as well as *authorization*. When you show the ID card, the security guard makes sure that you are indeed who you claim to be, possibly by looking at you and then the photograph on the ID card. This is *authentication*. Next, the guard makes sure you are allowed to go into the building, probably by verifying that your name appears on a list of approved individuals. This is *authorization*.

9.2 UNDERSTANDING AUTHENTICATION MECHANISMS

Now that you understand the basic terms regarding security in web applications, let's take a closer look at how authentication is implemented in Java servlets. The Servlet specification defines four mechanisms to authenticate users:

- HTTP Basic authentication
- HTTP Digest authentication
- HTTPS Client authentication
- HTTP FORM-based authentication

For the purpose of the exam, you will need to understand the basic features of each of these authentication mechanisms. They are all based on the *username/password* mechanism, in which the server maintains a list of all the usernames and passwords as well as a list of resources that have to be protected.

9.2.1 HTTP Basic authentication

HTTP Basic authentication, which is defined in the HTTP 1.1 specification, is the simplest and most commonly used mechanism to protect resources. When a browser requests any of the protected resources, the server asks for a username/password. If the user enters a valid username/password, the server sends the resource. Let's take a closer look at the sequence of the events:

- 1 A browser sends a request for a protected resource. At this time, the browser does not know that the resource is protected, so it sends a normal HTTP request. For example:

```
GET /servlet/SalesServlet HTTP/1.1
```

- 2 The server observes that the resource is protected, and so instead of sending the resource, it sends a 401 `Unauthorized` message back to the client. In the message, it also includes a header that tells the browser that the Basic authentication is needed to access the resource. The header also specifies the context in which the authentication would be valid. This context is called *realm*. It helps organize the access control lists on the server into different categories and, at the same time, tells users which user ID/password to use if they are allowed access in different realms. The following is a sample response sent by a server:

```
HTTP/1.1 401 Unauthorized
Server: Tomcat/5.0.25
WWW-Authenticate: Basic realm="sales"
Content-Length=500
Content-Type=text/html

<html>
...detailed message
</html>
```

← **Specifies authentication type and realm**

In the above response message, the `WWW-Authenticate` header specifies `Basic` and `sales` as the authentication type and the realm, respectively.

- 3 Upon receiving the above response, the browser opens a dialog box prompting for a username and password (see figure 9.1).
- 4 Once the user enters the username and password, the browser resends the request and passes the values in a header named `Authorization`:

```
GET /servlet/SalesServlet HTTP/1.1
Authorization: Basic am9objpqaamo=
```

← **Sends the Base64 encoded value**

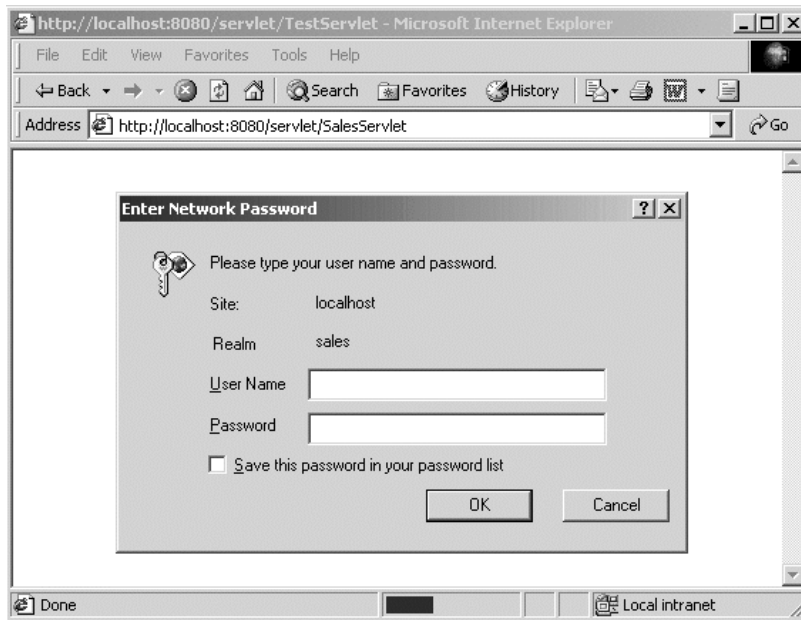


Figure 9.1 HTTP Basic authentication

The above request header includes the Base64 encoded value of the `username:password` string. The string, `am9objprqamo=`, is the encoded form of `john:jjj`.

- 5 When the server receives the request, it validates the username and the password. If they are valid, it sends the resource; otherwise, it sends the same 401 Unauthorized message again.
- 6 The browser displays the resource (or displays the username/password dialog box again).

Advantages

The advantages of HTTP Basic authentication are

- It is very easy to set up.
- All browsers support it.

Disadvantages

The disadvantages of HTTP Basic authentication are

- It is not secure because the username/password are not encrypted.
- You cannot customize the look and feel of the dialog box.

NOTE Base64 encoding is not an encryption method. Sun provides `sun.misc.Base64Encoder` and `sun.misc.Base64Decoder` classes that can encode and decode any string using this method. For more information, please refer to RFC 1521.

We will see how to use HTTP Basic authentication in section 9.2.5.

9.2.2 HTTP Digest authentication

The HTTP Digest authentication is the same as Basic except that in this case, the password¹ is sent in an encrypted format. This makes it more secure.

Advantage

The advantage of HTTP Digest authentication is

- It is more secure than Basic authentication.

Disadvantages

The disadvantages of HTTP Digest authentication are

- It is supported only by Microsoft Internet Explorer 5.
- It is not supported by many servlet containers since the specification does not mandate it.

9.2.3 HTTPS Client authentication

HTTPS is HTTP over SSL (Secure Socket Layer). SSL is a protocol developed by Netscape to ensure the privacy of sensitive data transmitted over the Internet. In this mechanism, authentication is performed when the SSL connection is established between the browser and the server. All the data is transmitted in the encrypted form using public-key cryptography, which is handled by the browser and the servlet container in a manner that is transparent to the servlet developers. The exam doesn't require you to know the details of this mechanism.

Advantages

The advantages of HTTPS Client authentication are

- It is the most secure of the four types.
- All the commonly used browsers support it.

¹ Actually, instead of the password, an MD5 digest of the password is sent. Please refer to RFC 1321 for more information.

Disadvantages

The disadvantages of HTTPS Client authentication are

- It requires a certificate from a certification authority, such as VeriSign.
- It is costly to implement and maintain.

9.2.4 FORM-based authentication

This mechanism is similar to Basic authentication. However, instead of using the browser's pop-up dialog box, it uses an HTML FORM to capture the username and password. Developers must create the HTML page containing the FORM, which allows them to customize its look and feel. The only requirement of the FORM is that its action attribute should be `j_security_check` and it must have two fields: `j_username` and `j_password`. Everything else is customizable.

Advantages

The advantages of FORM-based authentication are

- It is very easy to set up.
- All the browsers support it.
- You can customize the look and feel of the login screen.

Disadvantages

The disadvantages of FORM-based authentication are

- It is not secure, since the username/password are not encrypted.
- It should be used only when a session is maintained using cookies or HTTPS.

We will see how to use FORM-based authentication in the next section.

9.2.5 Defining authentication mechanisms for web applications

To ensure portability and ease of configuration at the deployment location, the authentication mechanism is defined in the deployment descriptor (`web.xml`) of the web application. However, before specifying which users should be authenticated, we have to configure their usernames and passwords. This step depends on the servlet container vendor. For Tomcat, it is quite easy.

Configuring users in Tomcat

Tomcat defines all the users in `<tomcat-root>\conf\tomcat-users.xml`. The following code snippet shows the default contents of this file:

```
<tomcat-users>
  <user name="tomcat" password="tomcat" roles="tomcat" />
  <user name="role1" password="tomcat" roles="role1" />
</tomcat-users>
```

```

    <user name="both" password="tomcat" roles="tomcat,role1" />
</tomcat-users>

```

This code defines three usernames: `tomcat`, `role1`, and `both`. The password is `tomcat` for all users.

An interesting piece of information in this file is the `roles` attribute. This attribute specifies the roles that the user plays. Permissions are assigned to roles instead of actual users. The concept of *role* comes straight from the real world; for example, a company may permit only a sales manager to access the sales data. It does not matter *who* the sales manager is. In fact, the sales manager may change over time. At any time, the sales manager is actually a user playing the sales manager role. Assigning permissions to roles instead of users gives us the flexibility to transfer permissions easily.

Let us add three more entries to the `tomcat-users.xml` file:

```

<tomcat-users>
  <user name="tomcat" password="tomcat" roles="tomcat" />
  <user name="role1" password="tomcat" roles="role1" />
  <user name="both" password="tomcat" roles="tomcat,role1" />

  <user name="john" password="jjj" roles="employee" />
  <user name="mary" password="mmm" roles="employee" />
  <user name="bob" password="bbb" roles="employee, supervisor" />
</tomcat-users>

```

We have added `john` and `mary` as employees and `bob` as a supervisor. Because a supervisor is also an employee, we have specified both roles for `bob`. We will employ these usernames later in the chapter.

Specifying the authentication mechanism

The authentication mechanism is specified in the deployment descriptor of the web application using the `<login-config>` element. The Servlet specification defines the `<login-config>` element as follows:

```

<!ELEMENT login-config (auth-method?, realm-name?, form-login-config?)>

```

Let's look at the subelements:

- `<auth-method>`. Specifies which of the four authentication methods should be used to validate the user: `BASIC`, `DIGEST`, `CLIENT-CERT`, or `FORM`.
- `<realm-name>`. Specifies the realm name to be used in HTTP Basic authorization only.
- `<form-login-config>`. Specifies the login page URL and the error page URL. This element is used only if `auth-method` is `FORM`; otherwise, it is ignored.

The following is a `web.xml` code snippet that shows an authentication mechanism configuration:

```

<web-app>
...
<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>sales</realm-name>
</login-config>
...
<web-app>

```

The above code uses the Basic mechanism to authenticate users. If we wanted to use the FORM mechanism, we'd need to write two HTML pages: one for capturing the username and password, and another to display an error message if the login fails. Finally, we'd need to specify these HTML files in the `<form-login-config>` element, as shown here:

```

<web-app>
...
<login-config>
  <auth-method>FORM</auth-method>
  <!--realm-name not required for FORM based authentication -->
  <form-login-config>
    <form-login-page>/formlogin.html</form-login-page>
    <form-error-page>/formerror.html</form-error-page>
  </form-login-config>
</login-config>
...
<web-app>

```

The `formlogin.html` file can be as simple as the following:

```

<html>
<body>
<h4>Please login:</h4>
  <form method="POST" action="j_security_check">
    <input type="text" name="j_username">
    <input type="password" name="j_password">
    <input type="submit" value="OK">
  </form>
</body>
</html>

```

The `formerror.html` file is even simpler:

```

<html>
<body>
  <h4>Sorry, your username and password do not match.</h4>
</body>
</html>

```

Observe that for the FORM method, we do not have to write any servlet to process the form. The action `j_security_check` triggers the servlet container to do the processing itself.

9.3 SECURING WEB APPLICATIONS DECLARATIVELY

It is very common for a web application to be developed by one group of individuals and then deployed by a very different group of people at another location. For example, many companies sell web applications as ready-made solutions for business needs. This means that the developer should be able to easily convey the security requirements of the application to the deployer. The deployer should also be able to customize certain aspects of the application's security without modifying the code. The servlet framework allows us to specify the detailed security requirements of the application in the deployment descriptor. This is called *declarative security*.

By default, all of the resources of a web application are accessible to everybody. To restrict access to the resources, we need to identify three things:

- *Web resource collection*—Identifies the resources of the application (that is, HTML files, servlets, and so forth) that must be protected from public access. A user must have appropriate authorization to access resources identified under a web resource collection.
- *Authorization constraint*—Identifies the roles that a user can be assigned. Instead of specifying permissions for individual users, permissions are assigned to roles. As discussed earlier, this reduces a tight coupling of permissions and the actual users. For example, an `AdminServlet` may be accessible to any user who is in the administrator role. At deployment time, any of the actual users may be configured as the administrator.
- *User data constraint*—Specifies the way the data must be transmitted between the sender and the receiver. In other words, this constraint specifies the transport layer requirement of the application. It formulates the policies for maintaining data integrity and confidentiality. For example, an application may require the use of HTTPS as a means of communication instead of plain HTTP.

We can configure all three of these items in the deployment descriptor of the web application by using the element `<security-constraint>`. This element, which falls directly under the `<web-app>` element of `web.xml`, is defined as follows:

```
<!ELEMENT security-constraint (display-name?, web-resource-collection+,  
auth-constraint?, user-data-constraint?)>
```

Let's look at the subelements one by one.

9.3.1 display-name

This is an optional element. It specifies a name for the security constraint that is easily identifiable.

9.3.2 web-resource-collection

As the name suggests, `web-resource-collection` specifies a collection of resources to which this security constraint applies. We can define one or more

web resource collections in the `<security-constraint>` element. It is defined as follows:

```
<!ELEMENT web-resource-collection (web-resource-name, description?,  
                                url-pattern*, http-method*)>
```

- `web-resource-name`—Specifies the name of the resource.
- `description`—Provides a description of the resource.
- `url-pattern`—Specifies the URL pattern through which the resource will be accessed. We can specify multiple URL patterns to group multiple resources together. Recall from chapter 5, “Structure and deployment,” that `<url-pattern>` is also used to specify the URL-to-servlet mapping.
- `http-method`—Provides a finer control over HTTP requests. This element specifies the HTTP methods to which this constraint will be applied. For example, we can use `http-method` to restrict POST requests only to authorized users while allowing GET requests for all the users.

Let’s look at a sample web resource collection:

```
<web-app>  
  ...  
  <security-constraint>  
    <web-resource-collection>  
      <web-resource-name>reports</web-resource-name>  
  
      <url-pattern>/servlet/SalesReportServlet/*</url-pattern>  
      <url-pattern>/servlet/FinanceReportServlet/*</url-pattern>  
      <url-pattern>/servlet/HRReportServlet/*</url-pattern>  
  
      <http-method>GET</http-method>  
      <http-method>POST</http-method>  
    </web-resource-collection>  
  ...  
</security-constraint>  
  ...  
</web-app>
```



In this collection, we specify three servlets to which we want to apply the security constraint. We have defined only GET and POST in the `<http-method>` section. This means that only these methods will have a restricted access; all other requests to these servlets will be open to all users.

If no `<http-method>` element is present, then the constraint applies to all of the HTTP methods.

9.3.3 auth-constraint

This element specifies the roles that can access the resources specified in the `web-resource-collection` section. It is defined as follows:

```
<!ELEMENT auth-constraint (description?, role-name*)>
```

- `description`—Describes the constraint.
- `role-name`—Specifies the role that can access the resources. It can be `*` (which means all the roles defined in the web application), or it must be a name that is defined in the `<security-role>` element of the deployment descriptor.

Here's an example:

```
<web-app>
  ...
  <security-role>
    <role-name>supervisor</role-name>
  </security-role>
  <security-role>
    <role-name>director</role-name>
  </security-role>
  <security-role>
    <role-name>employee</role-name>
  </security-role>
  ...
  <security-constraint>
    ...
    <auth-constraint>
      <description>accessible to all supervisors and
        directors</description>
      <role-name>supervisor</role-name>
      <role-name>director</role-name>
    </auth-constraint>
    ...
  </security-constraint>
  ...
</web-app>
```

This example specifies that the security constraint applies to all the users who are in the role of `supervisor` or `director`.

9.3.4 user-data-constraint

This element specifies how the data should be communicated between the client and the server. It is defined as follows:

```
<!ELEMENT user-data-constraint (description?, transport-guarantee)>
```

- `description`—Describes the constraint.
- `transport-guarantee`—Contains one of three values: `NONE`, `INTEGRAL`, or `CONFIDENTIAL`. `NONE` implies that the application does not need any guarantee about the integrity or confidentiality of the data transmitted, while `INTEGRAL` and `CONFIDENTIAL` imply that the application requires the data transmission to have data integrity and confidentiality, respectively. Usually, plain HTTP is used when `transport-guarantee` is set to `NONE`, and HTTPS is used when `transport-guarantee` is set to `INTEGRAL` or `CONFIDENTIAL`.

Here is an example of user-data-constraint:

```
<web-app>
...
  <security-constraint>
    ...
    <user-data-constraint>
      <description>requires the data transmission
        to be integral</description>
      <transport-guarantee>INTEGRAL</transport-guarantee>
    </user-data-constraint>
    ...
  </security-constraint>
  ...
</web-app>
```

9.3.5 Putting it all together

Now, let's build a simple web application containing just one servlet but with all of the bells and whistles for its security. As you work through the examples in this section, it's a good idea to restart Tomcat and to use a new browser window whenever you make changes to the code or configuration of the sample application.

The deployment descriptor

As we explained earlier, all of the security requirements of a web application can be specified in the deployment descriptor, as illustrated in listing 9.1.

Listing 9.1 web.xml showing declarative security configuration

```
<?xml version="1.1" encoding="ISO-8859-1"?>
<!DOCTYPE web-app
  PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
  "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>

  <servlet> ← Defines a servlet
    <servlet-name>SecureServlet</servlet-name>
    <servlet-class>SecureServlet</servlet-class>
  </servlet>

  <servlet-mapping>
    <servlet-name>SecureServlet</servlet-name>
    <url-pattern>/secure</url-pattern>
  </servlet-mapping>
  <security-constraint> ← Defines the security
    constraint for the servlet
    <web-resource-collection>
      <web-resource-name>declarative security test</web-resource-name>
      <url-pattern>/secure</url-pattern>
```

```

    <http-method>POST</http-method>
</web-resource-collection>

<auth-constraint>
    <role-name>supervisor</role-name>
</auth-constraint>

<user-data-constraint>
    <transport-guarantee>NONE</transport-guarantee>
</user-data-constraint>
</security-constraint>

<login-config> ← Defines the authentication mechanism
    <auth-method>FORM</auth-method>

    <form-login-config>
        <form-login-page>/formlogin.html</form-login-page>
        <form-error-page>/formerror.html</form-error-page>
    </form-login-config>
</login-config>
<security-role> ← Defines the security role
    <role-name>supervisor</role-name>
</security-role>
</web-app>

```

The `web.xml` file for our web application (listing 9.1) is straightforward. It defines a servlet followed by a security constraint for the servlet.

The resource to be protected is identified by the `<url-pattern>` element of `<web-resource-collection>`. Observe that in the `<web-resource-collection>` section we have specified only the POST method; this means that this security constraint applies only to the POST requests. All other HTTP methods are accessible to all the users. We could say that the word `resource` in `web-resource-collection` is a misnomer. With respect to a `web-resource-collection`, a resource is not just a servlet or a JSP page—it is an HTTP method sent to that servlet. In listing 9.1, the resource to which we are applying the constraint is the POST method sent to `SecureServlet`.

The `auth-constraint` section specifies that this resource should only be accessible to supervisors. The `role-name` that we use here must be defined in the `security-role` section.

The `transport-guarantee` is `NONE`, implying that HTTP will be used as the communication protocol.

The `<login-config>` section is exactly as we discussed in section 9.2.5. We can use either `BASIC` or `FORM` as the authentication mechanism.

The SecureServlet

Listing 9.2 contains the code for the SecureServlet, which was specified in the `<servlet-name>` element in the deployment descriptor in listing 9.1.

Listing 9.2 Code for SecureServlet

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class SecureServlet extends HttpServlet
{
    public void doGet(HttpServletRequest req,
                      HttpServletResponse res)
                      throws IOException
    {
        PrintWriter pw = res.getWriter();

        pw.println("<html><head>");
        pw.println("<title>Declarative Security Example</title>");
        pw.println("</head>");
        pw.println("<body>");
        pw.println("Hello! HTTP GET request is open to all
                                                           users.");
        pw.println("</body></html>");
    }

    public void doPost(HttpServletRequest req,
                      HttpServletResponse res)
                      throws IOException
    {
        PrintWriter pw = res.getWriter();

        pw.println("<html><head>");
        pw.println("<title>Declarative Security Example</title>");
        pw.println("</head>");
        pw.println("<body>");
        String name = req.getParameter("username");
        pw.println("Welcome, "+name+"!");
        pw.println("<br>You are seeing this page because you are
                                                           a supervisor.");
        pw.println("</body></html>");
    }
}
```

The servlet code in listing 9.2 is fairly simple and self-explanatory. We implemented the `doGet()` and `doPost()` methods for demonstration purposes only.

An important point to observe here is that the servlet does not have any security-related code. All of the security aspects are taken care of by the servlet container with the help of the deployment descriptor.

Running the example

You can access the complete working code from the Manning web site. Simply copy the `chapter09-declarative` directory to the `webapps` directory of your Tomcat installation and restart Tomcat.

- From your browser, go to `http://localhost:8080/chapter09-declarative/secure`. This sends a GET request to the servlet. Note that you are not asked for a username or password.
- To see the behavior of the POST method, go to `http://localhost:8080/chapter09-declarative/posttest.html`. This HTML file contains a FORM, which sends a POST request to the servlet. Observe that this time you get the login page (`formlogin.html`) because we have specified POST in `<http-method>` in the `<web-resource-collection>` section of the deployment descriptor (listing 9.1). The servlet's `doPost()`² method is executed only if you enter `bob` and `bbb` as the user ID and password, since `bob` is the only user we have defined in the `tomcat-users.xml` with the role of `super-visor`. For other values, you will get the error page (`formerror.html`). The code for `posttest.html` is simple and contains the following six lines:

```
<html><body>
<form action="/chapter09-declarative/secure" method="POST">
  Name: <input type="text" name="username">
  <input type="submit">
</form>
</body></html>
```

The combination of FORM-based authentication and POST works well in other containers.

² Unfortunately, Tomcat 5 incorrectly handles the combination of FORM-based authentication and HTTP POST requests for protected resources, and the example application reveals this flaw. Tomcat authenticates the user as expected but calls the `doGet()` method on the resource instead of `doPost()`. This means the page you see will be the same as that for the unsecured access. In order to see the correct page, switch the `<auth-method>` to BASIC and restart Tomcat. Use a new browser window to access `http://localhost:8080/chapter09-declarative/posttest.html`. After submitting the form, and authenticated with `bob` and `bbb` as the user ID and password, you should see the correct page.

9.4 SECURING WEB APPLICATIONS PROGRAMMATICALLY

In some cases, declarative security is not sufficient or fine-grained enough for the application. For example, suppose we want a servlet to be accessed by all employees. However, we want the server to generate a certain output for directors and a different output for other employees.

For such cases, the Servlet specification allows the servlet to have security-related code. This is called *programmatic security*. In this approach, a servlet identifies the role that a user is playing and then generates the output according to the role. As shown in table 9.1, the `HttpServletRequest` interface provides three methods for identifying the user and the role.

Table 9.1 `HttpServletRequest` methods for identifying a user

Method	Description
<code>String getRemoteUser()</code>	This method returns the login name of the user, if the user has been authenticated, or null if the user has not been authenticated.
<code>Principal getUserPrincipal()</code>	This method returns a <code>java.security.Principal</code> object containing the name of the current authenticated user. It returns null if the user is not authenticated.
<code>boolean isUserInRole(String rolename)</code>	This method returns a Boolean indicating whether the authenticated user is included in the specified logical role. It returns false if the user is not authenticated.

Let's modify the `SecureServlet` that we saw in section 9.3 to generate customized output. We will just change the `doPost()` method and leave the rest as is:

```
public void doPost(HttpServletRequest req,
                  HttpServletResponse res)
    throws IOException
{
    PrintWriter pw = res.getWriter();

    pw.println("<html><head>");
    pw.println("<title>Programatic Security Example</title>");
    pw.println("</head>");
    pw.println("<body>");

    String username = req.getRemoteUser();    ← Gets the username

    if(username != null)
        pw.println("<h4>Welcome, "+username+"!</h4>");

    if(req.isUserInRole("director"))    ← Determines if the user is a manager
    {
        pw.println("<b>Director's Page!</b>");
    }
    else
```

```

    {
        pw.println("<b>Employee's Page!</b>");
    }
    pw.println("</body></html>");
}

```

In this code, we retrieve the login name of the user using the `getRemoteUser()` method and determine whether the user is a director.

Obviously, this requires hard-coding the role name `director` in the servlet code. At the actual deployment location, however, users may be called supervisors instead of directors. To allow flexibility in defining the roles at deployment time, the servlet developer must convey the hard-coded values to the deployer. The deployer then maps these hard-coded values to the actual role values that are used in the deployment environment (as we'll see in the example that follows).

Now, let's modify the deployment descriptor of our previous example. Listing 9.3 contains the new deployment descriptor; the modified part appears in bold.

Listing 9.3 web.xml for programmatic security configuration

```

<?xml version="1.1" encoding="ISO-8859-1"?>
<!DOCTYPE web-app
    PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
        "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>
    <servlet>
        <servlet-name>SecureServlet</servlet-name>
        <servlet-class>SecureServlet</servlet-class>
        <security-role-ref>
            <role-name>director</role-name>
            <role-link>supervisor</role-link>
        </security-role-ref>
    </servlet>
    <security-constraint>
        <web-resource-collection>
            <web-resource-name>programmatic security test</web-resource-name>
            <url-pattern>/servlet/SecureServlet</url-pattern>
            <http-method>POST</http-method>
        </web-resource-collection>
        <auth-constraint>
            <role-name>employee</role-name>
        </auth-constraint>
        <user-data-constraint>
            <transport-guarantee>NONE</transport-guarantee>
        </user-data-constraint>
    </security-constraint>

```

Role name hard-coded
in the servlet

Role name defined in
the servlet container

Gives access to all
the employees


```

<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>sales</realm-name>
  <form-login-config>
    <form-login-page>/formlogin.html</form-login-page>
    <form-error-page>/formerror.html</form-error-page>
  </form-login-config>
</login-config>

<security-role>
  <role-name>supervisor</role-name>
</security-role>
<security-role>
  <role-name>employee</role-name>
</security-role>
</web-app>

```

In listing 9.3, the `<security-role-ref>` section is used to associate the hard-coded role name used by the servlet (`director`) to the actual role name (`supervisor`). Since we now want the servlet to be accessed by all employees, we have also changed the `<security-constraint>` to let all staff access this servlet. That's all there is to the programmatic security model.

9.5 SUMMARY

As companies conduct more of their business over the Internet, security issues will continue to increase in importance. In this chapter, we learned how to make a web application secure. We first introduced some important security-related terms. *Authentication* is validating who a user is, *authorization* is verifying what the user can do, and *auditing* holds the user accountable for her actions. We also discussed *confidentiality* and *data integrity*, and how these terms apply to web applications. The Servlet specification defines four mechanisms to authenticate users: BASIC, CLIENT-CERT, FORM, and DIGEST. The authentication mechanism is defined in the deployment descriptor (`web.xml`) of the web application.

We also learned how to secure a web application declaratively by configuring its security aspects in the deployment descriptor, and we learned how to implement security-related code in a servlet in order to secure a web application programmatically. Finally, we presented a complete web application that uses all the security features provided by the Servlet specification.

You should now be able to answer questions that require an understanding of authorization, authentication, data integrity, auditing, malicious code, and web site attacks. You should know how to specify the security requirements of an application in the deployment descriptor and how to identify incorrectly written security constraints.

In the next chapter, we will learn about multithreaded and single-threaded servlets. We'll also learn how to develop thread-safe servlets.

9.6 **REVIEW QUESTIONS**

1. Which of the following correctly defines data integrity? (Select one)
 - a It guarantees that information is accessible only to certain users.
 - b It guarantees that the information is kept in encrypted form on the server.
 - c It guarantees that unintended parties cannot read the information during transmission between the client and the server.
 - d It guarantees that the information is not altered during transmission between the client and the server.

2. What is the term for determining whether a user has access to a particular resource? (Select one)
 - a Authorization
 - b Authentication
 - c Confidentiality
 - d Secrecy

3. Which one of the following must be done before authorization takes place? (Select one)
 - a Data validation
 - b User authentication
 - c Data encryption
 - d Data compression

4. Which of the following actions would you take to prevent your web site from being attacked? (Select three)
 - a Block network traffic at all the ports except the HTTP port.
 - b Audit the usage pattern of your server.
 - c Audit the Servlet/JSP code.
 - d Use HTTPS instead of HTTP.
 - e Design and develop your web application using a software engineering methodology.
 - f Use design patterns.

5. Identify the authentication mechanisms that are built into the HTTP specification. (Select two)
 - a Basic
 - b Client-Cert
 - c FORM
 - d Digest
 - e Client-Digest
 - f HTTPS

6. Which of the following deployment descriptor elements is used for specifying the authentication mechanism for a web application? (Select one)
- a security-constraint
 - b auth-constraint
 - c login-config
 - d web-resource-collection
7. Which of the following elements are used for defining a security constraint? Choose only those elements that come directly under the security-constraint element. (Select three)
- a login-config
 - b role-name
 - c role
 - d transport-guarantee
 - e user-data-constraint
 - f auth-constraint
 - g authorization-constraint
 - h web-resource-collection
8. Which of the following *web.xml* snippets correctly identifies all HTML files under the sales directory? (Select two)
- a

```
<web-resource-collection>
  <web-resource-name>reports</web-resource-name>
  <url-pattern>/sales/*.html</url-pattern>
</web-resource-collection>
```
 - b

```
<resource-collection>
  <web-resource-name>reports</web-resource-name>
  <url-pattern>/sales/*.html</url-pattern>
</resource-collection>
```
 - c

```
<resource-collection>
  <resource-name>reports</resource-name>
  <url-pattern>/sales/*.html</url-pattern>
</resource-collection>
```
 - d

```
<web-resource-collection>
  <web-resource-name>reports</web-resource-name>
  <url-pattern>/sales/*.html</url-pattern>
  <http-method>GET</http-method>
</web-resource-collection>
```
9. You want your `PerformanceReportServlet` to be accessible only to managers. This servlet generates a performance report in the `doPost()` method

based on a FORM submitted by a user. Which of the following correctly defines a security constraint for this purpose? (Select one)

- a** <security-constraint>
<web-resource-collection>
 <web-resource-name>performance report</web-resource-name>
 <url-pattern>/servlet/PerformanceReportServlet</url-pattern>
 <http-method>GET</http-method>
</web-resource-collection>

<auth-constraint>
 <role-name>manager</role-name>
</auth-constraint>

<user-data-constraint>
 <transport-guarantee>NONE</transport-guarantee>
</user-data-constraint>
</security-constraint>
- b** <security-constraint>

<web-resource-collection>
 <web-resource-name>performance report</web-resource-name>
 <url-pattern>/servlet/PerformanceReportServlet</url-pattern>

 <http-method>*</http-method>
</web-resource-collection>

<accessibility>
 <role-name>manager</role-name>
</accessibility>

<user-data-constraint>
 <transport-guarantee>CONFIDENTIAL</transport-guarantee>
</user-data-constraint>
</security-constraint>
- c** <security-constraint>

<web-resource-collection>
 <web-resource-name>performance report</web-resource-name>
 <url-pattern>/servlet/PerformanceReportServlet</url-pattern>
 <http-method>POST</http-method>
</web-resource-collection>

<accessibility>
 <role-name>manager</role-name>
</accessibility>

<user-data-constraint>
 <transport-guarantee>CONFIDENTIAL</transport-guarantee>
</user-data-constraint>
</security-constraint>

```
d <security-constraint>
  <web-resource-collection>
    <web-resource-name>performance_report</web-resource-name>
    <url-pattern>/servlet/PerformanceReportServlet</url-pattern>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>manager</role-name>
  </auth-constraint>
</security-constraint>
```

10. Which of the following statements regarding authentication mechanisms are correct? (Select two)
- a The HTTP Basic mechanism transmits the username/password "in the open."
 - b The HTTP Basic mechanism uses HTML FORMs to collect usernames/passwords.
 - c The transmission method in the Basic and FORM mechanisms is the same.
 - d The method of capturing the usernames/passwords in the Basic and FORM mechanisms is the same.
11. Which of the following statements are correct for an unauthenticated user? (Select two)
- a `HttpServletRequest.getUserPrincipal()` returns null.
 - b `HttpServletRequest.getUserPrincipal()` throws `SecurityException`.
 - c `HttpServletRequest.isUserInRole(rolename)` returns false.
 - d `HttpServletRequest.getRemoteUser()` throws a `SecurityException`.

P A R T

3

JavaServer Pages and design patterns

In developing web components, we normally use JavaServer Pages (JSPs) for the presentation. In addition, we use design patterns to provide the theoretical framework for large-scale applications. In this final part of the book, we'll cover both topics.



CHAPTER 10

The JSP technology model—the basics

- 10.1 SP syntax elements 166
- 10.2 The JSP page life cycle 173
- 10.3 Understanding JSP page directive attributes 181
- 10.4 Summary 186
- 10.5 Review questions 186

EXAM OBJECTIVES

- 6.1** Identify, describe, or write the JSP code for the following expressions:
 - Template text;
 - Scripting elements (comments, directives, declarations, scriptlets, and expressions);
 - Standard and custom actions; and
 - Expression language elements.(Section 10.1)
- 6.2** Write JSP code that uses the directive:
 - ‘page’ (with attributes ‘import’, ‘session’, ‘contentType’, and ‘isELIgnored’)
 - ‘include’, and
 - ‘taglib’.(Section 10.3)
- 6.4** Describe the purpose and event sequence of the JSP life cycle:
 - JSP page translation
 - JSP page compilation

- Load class
- Create instance
- Call the `jspinit` method
- Call the `_jspService` method
- Call the `_jspDestroy` method

(Section 10.2)

INTRODUCTION

In the J2EE suite of specifications that includes servlets, JavaServer Pages (JSP), the Java Naming and Directory Interface (JNDI), Enterprise JavaBeans (EJB), and so forth, the JSP is a web-tier specification that supplements the Servlet specification and is useful in the development of web interfaces for enterprise applications. JSP is a technology that combines the HTML/XML markup languages and elements of the Java programming language to return dynamic content to a web client. For this reason, it is commonly used to handle the presentation logic of a web application, although the JSP pages may also contain business logic.

In this chapter, we will discuss the basic syntax of the JSP scripting language and the JSP page life cycle. This chapter gives you the basics you need to understand the JSP technology model and will help you grasp the more complex topics covered in the next chapter.

10.1 SP SYNTAX ELEMENTS

Just like any other language, the JSP scripting language has a well-defined grammar and includes syntax elements for performing various tasks, such as declaring variables and methods, writing expressions, and calling other JSP pages. At the top level, these syntax elements, also called JSP tags, are classified into six categories. Table 10.1 summarizes the element categories and their basic use.

Table 10.1 JSP element types

JSP tag type	Brief description	Tag syntax
Directive	Specifies translation time instructions to the JSP engine	<code><%@ Directives %></code>
Declaration	Declares and defines methods and variables	<code><%! Java Declarations %></code>
Scriptlet	Allows the developer to write free-form Java code in a JSP page	<code><% Some Java code %></code>
Expression	Used as a shortcut to print values in the output HTML of a JSP page	<code><%= An Expression %></code>
Action	Provides request-time instructions to the JSP engine	<code><jsp:actionName /></code>
Comment	Used for documentation and for commenting out parts of JSP code	<code><%-- Any Text --%></code>

The exam objectives covered in this chapter require you to know the syntax and purpose of the first four element types: directives, declarations, scriptlets, and expressions. We will briefly introduce actions in section 10.1.5, and explain them in detail in chapters 12 and 14. Although you don't need to be familiar with comments to do well on the exam, they are very useful when writing JSP pages, and we will discuss them briefly in section 10.1.6.

Listing 10.1 is a simple JSP page that counts the number of times it is visited. It demonstrates the use of the different elements, which we will explain in the sections following the listing.

Listing 10.1 counter.jsp

```
<html><body>
<%@ page language="java" %>   ← Directive
<%! int count = 0;           %> ← Declaration
<% count++;                  %> ← Scriptlet

Welcome! You are visitor number
<%= count                     %> ← Expression

</body></html>
```

When this file is accessed for the first time via the URL `http://localhost:8080/chapter10/counter.jsp`, it displays the following line in the browser window:

```
Welcome! You are visitor number 1
```

On subsequent requests, the counter is incremented by 1 before the message is printed.

10.1.1 Directives

Directives provide general information about the JSP page to the JSP engine. There are three types of directives: `page`, `include`, and `taglib`.

A `page` directive informs the engine about the overall properties of a JSP page. For example, the following `page` directive informs the JSP engine that we will be using Java as the scripting language in our JSP page:

```
<%@ page language="java" %>
```

An `include` directive tells the JSP engine to include the contents of another file (HTML, JSP, etc.) in the current page. Here is an example of an `include` directive:

```
<%@ include file="copyright.html" %>
```

A `taglib` directive is used for associating a prefix with a tag library. The following is an example of a `taglib` directive:

```
<%@ taglib prefix="test" uri="taglib.tld" %>
```

See section 10.3 for details on the page directive. In chapter 12, “Reusable web components,” we will take a close look at the `include` directive. Because the concept of a tag library is a vast topic in itself, the exam objectives devote two sections to it. We will learn about the `taglib` directive in detail in chapter 15, “Using custom tags,” and learn specific methods of tag development in chapters 16 and 17.

A directive always starts with `<%@` and ends with `%>`. The general syntax of the three directives is

```
<%@ page    attribute-list %>
<%@ include attribute-list %>
<%@ taglib  attribute-list %>
```

In the sample tags above, `attribute-list` represents one or more attribute-value pairs that are specific to the directive. Here are some important points to remember about the syntax of the directives:

- The tag names, their attributes, and their values are all case sensitive.
- The value must be enclosed within a pair of single or double quotes.
- A pair of single quotes is equivalent to a pair of double quotes.
- There must be no space between the equals sign (=) and the value.

10.1.2 Declarations

Declarations declare and define variables and methods that can be used in the JSP page.¹ The following is an example of a JSP declaration:

```
<%! int count = 0; %>
```

This declares a variable named `count` and initializes it to 0. The variable is initialized only once when the page is first loaded by the JSP engine, and retains its value in subsequent client requests. That is why the `count` variable in listing 10.1 is not reset to 0 each time we access the page.

A declaration always starts with `<%!` and ends with `%>`. It can contain any number of valid Java declaration statements. For example, the following tag declares a variable and a method in a single tag:

```
<%!
    String color[] = {"red", "green", "blue"};
    String getColor(int i)
    {
        return color[i];
    }
%>
```

¹ Theoretically, a JSP declaration can contain any valid Java declaration including inner classes and static code blocks. However, such declarations are rarely used.

We can also write the above two Java declaration statements in two JSP declaration tags:

```
<%! String color[] = {"red", "green", "blue"}; %>

<%!
    String getColor(int i)
    {
        return color[i];
    }
%>
```

Note that since the declarations contain Java declaration statements, each variable's declaration statement must be terminated with a semicolon.

10.1.3 Scriptlets

Scriptlets are Java code fragments that are embedded in the JSP page. For example, this line from the `counter.jsp` example (listing 10.1) is a JSP scriptlet:

```
<% count++; %>
```

The scriptlet is executed each time the page is accessed, and the `count` variable is incremented with each request.

Since scriptlets can contain any Java code, they are typically used for embedding computing logic within a JSP page. However, we can use scriptlets for printing HTML statements, too. The following is equivalent to the code in listing 10.1:

```
<%@ page language="java" %>
<%! int count = 0; %>

<%
    out.print("<html><body>");
    count++;
    out.print("Welcome! You are visitor number " + count);
    out.print("</body></html>");
%>
```

Instead of writing normal HTML code directly in the page, we are using a scriptlet to achieve the same effect. The variable `out` refers to an object of type `javax.servlet.jsp.JspWriter`. We will learn about `out` in chapter 11, “The JSP technology model—advanced topics.”

A scriptlet always starts with `<%` and ends with `%>`. Note, however, that unlike the other elements, the opening tag of a scriptlet does not have any special character following `<%`. The code within the scriptlet must be valid in the Java programming language. For example, this is an error because it does not terminate the `print` statement with a semicolon:

```
<% out.print(count) %>
```

10.1.4 Expressions

Expressions act as placeholders for Java language expressions. This is an example of a JSP expression:

```
<%= count %>
```

The expression is evaluated each time the page is accessed, and its value is then embedded in the output HTML. For instance, in the previous `counter.jsp` example (listing 10.1), instead of incrementing the `count` variable in a scriptlet, we could have incremented it in the expression itself:

```
<html><body>
<%@ page language="java" %>
<%! int count = 0; %>

Welcome! You are visitor number <%= ++count %>
</body></html>
```

← Evaluates the expression and prints it out

A JSP expression always starts with `<%=` and ends with `%>`. Unlike variable declarations, expressions must not be terminated with a semicolon. Thus, the following is not valid:

```
<%= count; %>
```

We can print the value of any object or any primitive data type (`int`, `boolean`, `char`, etc.) to the output stream using an expression. We can also print the value of any arithmetic or Boolean expression or a value returned by a method call. The exam may ask you to identify valid JSP expressions. Tables 10.2 and 10.3 contain some examples of valid and invalid JSP expressions based on the following declarations:

```
<%!
    int anInt = 3;
    boolean aBool = true;
    Integer anIntObj = new Integer(3);
    Float aFloatObj = new Float(12.6);

    String str = "some string";
    StringBuffer sBuff = new StringBuffer();

    char getChar(){ return 'A'; }
%>
```

Table 10.2 Valid JSP expressions

Expression	Explanation
<code><%= 500 %></code>	An integral literal
<code><%= anInt*3.5/100-500 %></code>	An arithmetic expression
<code><%= aBool %></code>	A Boolean variable
<code><%= false %></code>	A Boolean literal

continued on next page

Table 10.2 Valid JSP expressions (continued)

Expression	Explanation
<code><%= !false %></code>	A Boolean expression
<code><%= getChar() %></code>	A method returning a char
<code><%= Math.random() %></code>	A method returning a double
<code><%= aVector %></code>	A variable referring to a Vector object
<code><%= aFloatObj %></code>	A method returning a float
<code><%= aFloatObj.floatValue() %></code>	A method returning a float
<code><%= aFloatObj.toString() %></code>	A method that returns a String object

Table 10.3 Invalid JSP expressions

Expression	Explanation
<code><%= aBool; %></code>	You cannot use a semicolon in an expression.
<code><%= int i = 20 %></code>	You cannot define anything inside an expression.
<code><%= sBuff.setLength(12); %></code>	The method does not return any value. The return type is void.

10.1.5 Actions

Actions are commands given to the JSP engine. They direct the engine to perform certain tasks during the execution of a page. For example, the following line instructs the engine to include the output of another JSP page, `copyright.jsp`, in the output of the current JSP page:

```
<jsp:include page="copyright.jsp" />
```

There are six standard JSP actions:

- `jsp:include`
- `jsp:forward`
- `jsp:useBean`
- `jsp:setProperty`
- `jsp:getProperty`
- `jsp:plugin`

The first two, `jsp:include` and `jsp:forward`, enable a JSP page to reuse other web components. We will discuss these two actions in chapter 12, “Reusable web components.”

The next three, `jsp:useBean`, `jsp:setProperty`, and `jsp:getProperty`, are related to the use of JavaBeans in JSP pages. We will discuss these three actions in chapter 14, “Using JavaBeans.”

The last action, `jsp:plugin`, instructs the JSP engine to generate appropriate HTML code for embedding client-side components, such as applets. This action is not specified in the exam objectives, and its details are beyond the scope of this book.

In addition to the six standard actions, a JSP page can have user-defined actions. These are called custom tags. We will learn about custom tags in chapters 15 (“Using custom tags”), 16 (“Developing classic custom tag libraries”), and 17 (“Developing simple custom tag libraries”).

The general syntax of a JSP action is

```
<jsp:actionName attribute-list />
```

In this tag, `actionName` is one of the six actions mentioned and `attribute-list` represents one or more attribute-value pairs that are specific to the action. As with directives, you should keep in mind these points:

- The action names, their attributes, and their values are case sensitive.
- The value must be enclosed within a pair of single or double quotes.
- A pair of single quotes is equivalent to a pair of double quotes.
- There must be no space between the equals sign (=) and the value.

10.1.6 Comments

Comments do not affect the output of a JSP page in any way but are useful for documentation purposes. The syntax of a JSP comment is

```
<%-- Anything you want to be commented --%>
```

A JSP comment always starts with `<%--` and ends with `--%>`.

We can comment the Java code within scriptlets and declarations by using normal Java-style comments and the HTML portions of a page by using HTML-style comments, as shown here:

```
<html><body>
  Welcome!
  <%-- JSP comment --%>
  <% //Java comment %>
  <!-- HTML comment -->
</body></html>
```

As we mentioned earlier, the exam does not cover comments, but they can be quite useful when you’re debugging JSP pages. The JSP engine drops everything between `<%--` and `--%>`, so it is easy to comment out large parts of a JSP page—including nested HTML and other JSP tags. However, remember that you cannot nest JSP comments within other JSP comments.

Quizlet

Q: Which of the following page directives are valid?

- a** `<% page language="java" %>`
- b** `<%! page language="java" %>`
- c** `<%@ page language="java" %>`

- A:** Only option c is correct. Directives use an @ in the opening tag.
Q: What is wrong with the following code?

```
<!% int i = 5; %>
<!% int getI() { return i; } %>
```

- A:** The opening tag for a declaration is <%! and not <!%.
Q: Assuming that myObj refers to an object and m1 () is a valid method on that object, tell why each of the following are valid or invalid JSP constructs.
- a** <% myObj.m1 () %>
 - b** <%= myObj.m1 () %>
 - c** <% =myObj.m1 () %>
 - d** <% =myObj.m1 (); %>
- A:** The following table explains why an option is valid or invalid.

Construct	Explanation
<% myObj.m1 () %>	Invalid: It is not an expression because it does not have an = sign. It is an invalid scriptlet because a semicolon is missing at the end of the method call.
<%=myObj.m1 () %>	Depends: The = sign makes it an expression. But if the return type of the method m1 () is void, it is invalid. A method call inside an expression is valid if and only if the return type of the method is not void.
<% =myObj.m1 () %>	Invalid: There is a space between <% and =. Hence, it is not an expression but a scriptlet. However, the scriptlet construct is not valid because =myObj.m1 (), by itself, is not a valid Java statement.
<% =myObj.m1 (); %>	Invalid: Same as previous example except that it has a semicolon.

The valid way to write this as a scriptlet is

```
<% myObj.m1 (); %>
```

However, this will just call the method; it will not generate any output. If the method m1 () returns a value, then the correct way to write this as an expression is

```
<%= myObj.m1 () %>
```

This will print the return value of the method call to the output HTML.

10.2 THE JSP PAGE LIFE CYCLE

A JSP page goes through seven phases in its lifetime. These phases are called *life-cycle phases*. The exam requires you to know the sequence of the phases and the activity that takes place in each of the phases. But before we start discussing the life cycle of a JSP page, we need to understand the two important points regarding JSP pages that are explained in the following sections.

10.2.1 JSP pages are servlets

Although, structurally, a JSP page *looks* like an HTML page, it actually runs as a servlet. The JSP engine parses the JSP file and creates a corresponding Java file. This file declares a servlet class whose members map directly to the elements of the JSP file. The JSP engine then compiles the class, loads it into memory, and executes it as it would any other servlet. The output of this servlet is then sent to the client. Figure 10.1 illustrates this process.

10.2.2 Understanding translation units

Just as an HTML page can include the contents of other HTML pages (for example, when using frames), a JSP page can include the contents of other JSP pages and HTML

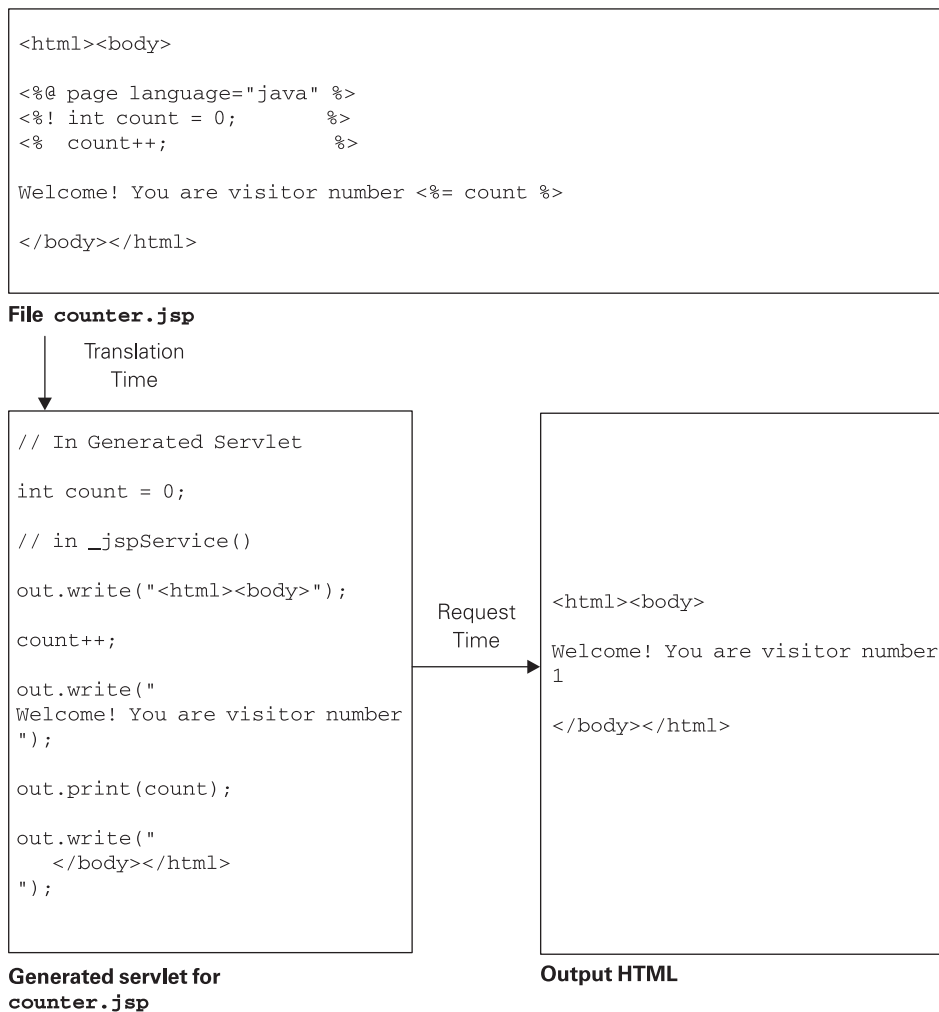


Figure 10.1 A JSP page as a servlet

pages. This is done with the help of the `include` directive (see chapter 12 for more information). But an important thing to remember here is that when the JSP engine generates the Java code for a JSP page, it also inserts the contents of the included pages into the servlet that it generates. The set of pages that is translated into a single servlet class is called a *translation unit*. Some of the JSP tags affect the whole translation unit and not just the page in which they are declared.

Keep in mind these other points regarding a translation unit:

- The page directives explained in section 10.3 affect the whole translation unit.
- A variable declaration cannot occur more than once in a single translation unit. For example, we cannot declare a variable in an included page using the `include` directive if it is already declared in the including page since the two pages constitute a single translation unit.
- The standard action `<jsp:useBean>` cannot declare the same bean twice in a single translation unit. We examine the `jsp:useBean` action further in chapter 14.

10.2.3 JSP life-cycle phases

You might have observed that when a JSP page is accessed for the first time, the server is slower in responding than it is in the second, third, and subsequent accesses. This is because, as we mentioned previously, every JSP page must be converted into an instance of a servlet class before it can be used to service client requests. For each request, the JSP engine checks the timestamps of the source JSP page and the corresponding servlet class file to determine if the JSP page is new or if it has already been converted into a class file. Therefore, if we modify a JSP page, the whole process of converting the JSP page into a servlet is performed again. This process consists of seven phases, and you need to understand their order and significance for the exam. Table 10.4 lists the phases in the order in which they occur.

Table 10.4 JSP page life-cycle phases

Phase name	Description
Page translation	The page is parsed and a Java file containing the corresponding servlet is created.
Page compilation	The Java file is compiled.
Load class	The compiled class is loaded.
Create instance	An instance of the servlet is created.
Call <code>jspInit()</code>	This method is called before any other method to allow initialization.
Call <code>_jspService()</code>	This method is called for each request.
Call <code>jspDestroy()</code>	This method is called when the servlet container decides to take the servlet out of service.

Creating the servlet instance

The first four life-cycle phases involve the process of converting the JSP page into an instance of a servlet class.

Translation

During the translation phase, the JSP engine reads a JSP page, parses it, and validates the syntax of the tags used. For example, the following directive is invalid since it uses an uppercase *P* in `Page` and will be caught during the translation phase:

```
<%@ Page language="java" %>
```

In addition to checking the syntax, the engine performs other validity checks, some of which involve verifying that

- The attribute-value pairs in the directives and standard actions are valid.
- The same JavaBean name is not used more than once in a translation unit.
- If we are using a custom tag library, the library is valid.
- The usage of custom tags is valid.

Once the validations are completed, the engine creates a Java file containing a public servlet class.

Compilation

In the compilation phase, the Java file generated in the previous step is compiled using the normal Java compiler `javac` (or using a vendor-provided compiler or even a user-specified compiler²). All the Java code that we write in declarations, scriptlets, and expressions is validated during this phase. For example, the following declaration tag is a valid JSP tag and will pass the translation phase, but the declaration statement is not a valid Java declaration statement because it does not end with a semicolon and will be caught during the compilation phase:

```
<%! int count = 0 %>
```

Scripting language errors (Java, in this case) are caught during the compilation phase. We can force a compilation of a JSP page without actually executing it by using the precompilation request parameter `jsp_precompile`. For example, if we want to compile the `counter.jsp` page without executing it, we must access the page as

```
http://localhost:8080/chapter10/counter.jsp?jsp_precompile=true
```

The engine will translate the JSP page and compile the generated servlet class without actually executing the servlet. This can be quite useful during the development phase if we have complex JSP pages that create database connections or access other J2EE

² This varies from container to container. Please consult the servlet container documentation for more information.

services. Also, it is always a good idea to precompile all the pages. In this way, we check for syntax errors and keep the pages ready to be served, thus reducing the response time for the first request to each page.

NOTE The parameter `jsp_precompile` takes a Boolean value, `true` or `false`. If the value is `false`, the precompilation will not occur. The parameter can also be specified without any value, in which case the default is `true`:

```
http://localhost:8080/chapter10/counter.jsp?jsp_precompile
```

In either case, `true` or `false`, the page will not be executed.

Also, this would be a good place to point out that all of the request parameter names that include the prefix `jsp` are reserved and must not be used for user-defined values. Thus, the following usage is not recommended and may result in unexpected behavior:

```
http://localhost:8080/chapter10/counter.jsp?jspTest=myTest
```

Loading and instantiation

After successful compilation, the container loads the servlet class into memory and instantiates it.

Calling the JSP life-cycle methods

The generated servlet class for a JSP page implements the `HttpJspPage` interface of the `javax.servlet.jsp` package. The `HttpJspPage` interface extends the `JspPage` interface of the same package, which in turn extends the `Servlet` interface of the `javax.servlet` package. The generated servlet class thus implements all the methods of these three interfaces and is also known as the page's implementation class.

The `JspPage` interface declares only two methods—`jspInit()` and `jspDestroy()`—that must be implemented by all JSP pages regardless of the client-server protocol. However, the JSP specification has provided the `HttpJspPage` interface specifically for JSP pages serving HTTP requests. This interface declares one method: `_jspService()`. Here are the signatures of the three JSP methods:

```
public void jspInit();

public void _jspService(HttpServletRequest request,
                        HttpServletResponse response)
                        throws
                        javax.servlet.ServletException,
                        java.io.IOException;

public void jspDestroy();
```

These methods are called the life-cycle methods of the JSP pages. The `jspInit()`, `_jspService()`, and `jspDestroy()` methods of a JSP page are equivalent to the `init()`, `service()`, and `destroy()` methods of a servlet, respectively.

Every JSP engine vendor provides a vendor-specific class that is used as a base class for the page's implementation class. This base class provides the default implementations of all the methods of the `Servlet` interface and the default implementations of both methods of the `JspPage` interface: `jspInit()` and `jspDestroy()`. During the translation phase, the engine adds the `_jspService()` method to the JSP page's implementation class, thus making the class a concrete subclass of the three interfaces.

jspInit()

The container calls `jspInit()` to initialize the servlet instance. It is called before any other method, and is called only once for a servlet instance. We normally define this method to do initial or one-time setup, such as acquiring resources and initializing the instance variables that have been declared in the JSP page using `<%! . . . %>` declarations.

_jspService()

The container calls the `_jspService()` for each request, passing it the request and the response objects. All of the HTML elements, the JSP scriptlets, and the JSP expressions become a part of this method during the translation phase. We discuss the details of this method in chapter 11.

jspDestroy()

When the container decides to take the instance out of service, it calls the `jspDestroy()` method. This is the last method that is called on the servlet instance, and it is used to clean up the resources acquired in the `jspInit()` method.

We are not required to implement the `jspInit()` and `jspDestroy()` methods, since they have already been implemented by the base class. If we need to override them, we can do so using the JSP declaration tag `<%! . . . %>`. However, we cannot define our own `_jspService()` method because the engine generates it automatically.

10.2.4 JSP life-cycle example

Let's modify our counter example to add persistence capabilities to it so that the counter does not start from 1 each time the server is shut down and restarted. Listing 10.2 illustrates how we can

- Use `jspInit()` to load the previous value of the counter from a file when the server starts.
- Use `jspDestroy()` to save the final value to the file when the server shuts down.

Listing 10.2 persistent_counter.jsp

```
<%@ page language="java" import="java.io.*" %>

<%!
    // A variable to maintain the number of visits.
    int count = 0;

    // Path to the file, counter.db, which stores the count
    // value in a serialized form. The file acts like a database.
    String dbPath;

    // This is the first method called by the container,
    // when the page is loaded. We open the db file,
    // read the integer value, and initialize the count variable.
    public void jspInit()
    {
        try
        {
            dbPath = getServletContext().getRealPath("/WEB-INF/counter.db");
            FileInputStream fis = new FileInputStream(dbPath);
            DataInputStream dis = new DataInputStream(fis);
            count = dis.readInt();
            dis.close();
        }
        catch(Exception e)
        {
            log("Error loading persistent counter", e);
        }
    }
%>

<%--
    The main content that goes to the browser.
    This will become a part of the generated _jspService() method
--%>
<html><body>
<% count++; %>
Welcome! You are visitor number
<%= count %>
</body></html>

<%!
    // This method is called by the container only once when the
    // page is about to be destroyed. We open the db file in this
    // method and save the value of the count variable as an integer.

    public void jspDestroy()
    {
        try
        {
            FileOutputStream fos = new FileOutputStream(dbPath);
            DataOutputStream dos = new DataOutputStream(fos);
```

```

        dos.writeInt(count);
        dos.close();
    }
    catch(Exception e)
    {
        log("Error storing persistent counter", e);
    }
}
%>

```

This example illustrates three things: the use of the `jspInit()` method, the use of the `jspDestroy()` method, and the use of the `getServletContext()` method.

When the page is first loaded into the servlet container, the engine will call the `jspInit()` method. In this method, we initialize the `count` variable to the value read in from the resource database file `"/WEB-INF/counter.db"`. During its lifetime, the JSP page may be accessed zero or more times, and each time the `_jspService()` method will be executed. Since the scriptlet `<% count++; %>` becomes a part of the `_jspService()` method, the expression `count++` is evaluated each time, increasing the counter by 1. Finally, when the page is about to be destroyed, the container will call the `jspDestroy()` method. In this method, we open the resource database file again, and save the latest value of the variable `count` into it.

Because the JSP page is converted into a servlet, we can call all the methods in a JSP page that we can call on a servlet. Hence we can get the `ServletContext` object via `getServletConfig().getServletContext()`. Also, the base class of the page's generated class extends `javax.servlet.http.HttpServlet`, which gives us access to the `log()` method. In Tomcat and many other containers, the base class of the page's generated class also implements the `ServletConfig` interface. Thus, in both methods, `jspInit()` and `jspDestroy()`, we get the `ServletContext` object by using the method `getServletContext()`, which is actually defined in the `javax.servlet.ServletConfig` interface. The returned `ServletContext` object can then be used in a JSP page exactly the way we use it in normal servlets. In our example, we are using the `ServletContext` object to convert the relative path of a resource into its real path. If the web application is installed in the directory `C:\jakarta-tomcat5.0.25\webapps\chapter10`, then a call to `getServletContext().getRealPath("/WEB-INF/counter.db")`; will return `C:\jakarta-tomcat-5.0.25\webapps\chapter10\WEB-INF\counter.db`.

When the server is started the very first time and the page is first accessed, the file `counter.db` does not exist and a `FileNotFoundException` is thrown. We can catch the exception and log the error message with this method.³ When the server is

³ Tomcat uses the `<CATALINA_HOME>\logs\` directory as the default directory to create log files.

shut down the first time, the `jspDestroy()` method creates a new file, and the current value of the variable is written into it. When the server is started the second time and the JSP page is loaded, the `jspInit()` method will find the file and initialize the count variable to its previously saved value.

The JSP technology thus combines the best of both worlds: the ease of use offered by the web scripting methodology and the object-oriented features of the servlet technology.

10.3 UNDERSTANDING JSP PAGE DIRECTIVE ATTRIBUTES

A page directive informs the JSP engine about the overall properties of a JSP page. This directive applies to the entire translation unit and not just to the page in which it is declared. Table 10.5 describes the 12 possible attributes for the page directive.

Table 10.5 Attributes for the page directive

Attribute name	Description	Default value/s
<code>import</code>	A comma-separated list of Java classes and packages that we want to use in the JSP page.	<code>java.lang.*;</code> <code>javax.servlet.*;</code> <code>javax.servlet.jsp.*;</code> <code>javax.servlet.http.*;</code>
<code>session</code>	A Boolean literal specifying whether the JSP page takes part in an HTTP session.	<code>true</code>
<code>errorPage</code>	Specifies a relative URL to another JSP page that is capable of handling errors on behalf of the current page.	<code>null</code>
<code>isErrorPage</code>	A Boolean literal specifying whether the current JSP page is capable of handling errors.	<code>false</code>
<code>language</code>	Any scripting language supported by the JSP engine.	<code>java</code>
<code>extends</code>	Any valid Java class that implements <code>javax.servlet.jsp.JspPage</code> .	Implementation dependent
<code>buffer</code>	Specifies the size of the output buffer. If a buffer size is specified, it must be in kilobytes (kb). If buffering is not required, specify the string <i>none</i> .	Implementation dependent
<code>autoFlush</code>	A Boolean literal indicating whether the buffer should be flushed when it is full.	<code>true</code>
<code>info</code>	Any informative text about the JSP page.	Implementation dependent
<code>contentType</code>	Specifies the MIME type and character encoding for the output.	<code>text/html; charset=ISO-8859-1</code>
<code>pageEncoding</code>	Specifies the character encoding of the JSP page.	<code>ISO-8859-1</code>

While the exam requires that you know all of the valid page directive attributes and their values, it focuses more on the usage of the first four: `import`, `session`, `errorPage`, and `isErrorPage`.

10.3.1 The import attribute

The `import` attribute of a page directive is similar to the `import` statement in a Java class. For example, if we want to use the `Date` class of the package `java.util`, then we have to either use the fully qualified class name in the code or import it using the page directive. At the time of translation, the JSP engine inserts an `import` statement into the generated servlet for each of the packages declared using this attribute.

We can import multiple packages in a single tag by using a comma-separated list of package names, as shown here:

```
<%@ page import="java.util.*, java.io.*, java.text.*,  
               com.mycom.*, com.mycom.util.MyClass " %>
```

We can also use multiple tags for readability. For example, the above page directive can also be written as:

```
<%@ page import="java.util.* " %>  
<%@ page import="java.io.* " %>  
<%@ page import="java.text.* " %>  
<%@ page import="com.mycom.*, com.mycom.util.MyClass " %>
```

Since the order of `import` statements in a Java class does not matter, the order of `import` tags shown here does not matter, either. A JSP engine always imports the `java.lang.*`, `javax.servlet.*`, `javax.servlet.jsp.*`, and `javax.servlet.http.*` packages, so we do not have to import them explicitly.

NOTE `import` is the only attribute of the page directive that can occur multiple times in a translation unit. Duplicate values are ignored.

10.3.2 The session attribute

The `session` attribute indicates whether the JSP page takes part in an HTTP session. The default value is `true`, in which case the JSP engine declares the implicit variable `session`. (We will learn more about implicit variables in chapter 11.) If we do not want the page to participate in a session, then we have to explicitly add the following line:

```
<%@ page session="false" %>
```

10.3.3 The errorPage and isErrorPage attributes

During the execution of a page, it is possible that the embedded Java code will throw exceptions. Just as in normal Java programs, we can handle the exceptions in JSP pages using `try-catch` blocks. However, the JSP specification defines a better approach, which separates the error-handling code from the main page and thus promotes reusability of the exception-handling mechanism. In this approach, a JSP page uses the

errorPage attribute to delegate the exception to another JSP page that has the error-handling code. In listing 10.3, errorHandler.jsp is specified as the error handler.

Listing 10.3 hello.jsp: Using errorPage to delegate exceptions

```
<%@ page errorPage="errorHandler.jsp" %>
<html>
<body>
  <%
    if (request.getParameter("name")==null)
    {
      throw new RuntimeException("Name not specified");
    }
  %>
  Hello, <%=request.getParameter("name")%>
</body>
</html>
```

The JSP page in listing 10.3 throws an exception if the parameter name is not supplied in the request, but it does not catch the exception itself. Instead, with the help of the errorPage attribute, it instructs the JSP engine to delegate the error handling to errorHandler.jsp.

The isErrorPage attribute conveys whether the current page can act as an error handler for any other JSP page. The default value of the isErrorPage attribute is false. For example, the errorHandler.jsp that we used in the previous example must explicitly set this attribute to true, as shown in listing 10.4. In this case, the JSP engine declares the implicit variable exception in the page's servlet class.

Listing 10.4 errorHandler.jsp: Handling exceptions

```
<%@ page isErrorPage="true" %>
<html>
<body>
  Unable to process your request: <%=exception.getMessage()%><br>
  Please try again.
</body>
</html>
```

Notice that this page only extracts the information from the exception and generates an appropriate error message. Because it does not implement any business logic, it can be reused for different JSP pages.

It is not necessary that the errorPage value be a JSP page. It can also be a static file, such as an HTML page:

```
<%@ page errorPage="errorHandler.html" %>
```

Obviously, we cannot write a scriptlet or an expression in the HTML file `errorHandler.html` to generate dynamic messages.

NOTE In general, it is always a good programming practice to specify an error page in all the JSP pages. This prevents unanticipated error messages from being displayed on the client's browser.

10.3.4 The language and extends attributes

The `language` attribute specifies the language used by a page in declarations, scriptlets, and expressions. The default value is `java`, which is also the only value allowed by the JSP Specification 2.0. Needless to say, adding the following line to a JSP page is redundant:

```
<%@ page language="java" %>
```

The `extends` attribute specifies that the supplied class be used as a base class of the generated servlet. This is useful only if we want to customize the behavior of the generated servlet class. The default base class is vendor specific and is designed to work efficiently with the rest of the framework. Consequently, this attribute is seldom used. The following line shows the syntax for this attribute:

```
<%@ page extends="mypackage.MySpecialBaseServlet" %>
```

10.3.5 The buffer and autoFlush attributes

The `buffer` attribute specifies the minimum size required by the output buffer that holds the generated content until it is sent to the client. The default size of the buffer is JSP engine implementation dependent, but the specification mandates it to be at least 8kb. The following line sets the buffer size to 32kb:

```
<%@ page buffer="32kb" %>
```

The value of the buffer is in kilobytes and the suffix `kb` is mandatory. To send the data directly to the client without any buffering, we can specify the value as `none`.

The `autoFlush` attribute specifies whether the data in the output buffer should be sent to the client automatically as soon as the buffer is full. The default value for `autoFlush` is `true`. If it is set to `false` and the buffer is full, an exception is raised when we attempt to add more data to the buffer. Here is the syntax for this attribute:

```
<%@ page autoFlush="false" %>
```

Obviously, the following combinations occurring in a JSP page are invalid and may either cause an error at translation time or have an unknown behavior at runtime:

```
<%@ page buffer="none" autoFlush="false" %>
<%@ page buffer="0kb" autoFlush="false" %>
```

10.3.6 The info attribute

The `info` attribute allows us to specify the value of the string returned by the `getServletInfo()` method of the generated servlet. The following line shows one possible use:

```
<%@ page info="This is a sample Page. " %>
```

The default value of this attribute is implementation dependent.

10.3.7 The contentType and pageEncoding attributes

The `contentType` attribute specifies the MIME type and character encoding of the output. The default value of the MIME type is `text/html`; the default value of the character encoding is `ISO-8859-1`. The MIME type and character encoding are separated by a semicolon, as shown here:

```
<%@ page contentType="text/html; charset=ISO-8859-1" %>
```

This is equivalent to writing the following line in a servlet:

```
response.setContentType("text/html; charset=ISO-8859-1");
```

The `pageEncoding` attribute specifies the character encoding of the JSP page. The default value is `ISO-8859-1`. The following line illustrates the syntax:

```
<%@ page pageEncoding="ISO-8859-1" %>
```

Quizlet

Q: Which of the following page directives are valid and which are invalid?

- a** `<%@ page import="java.util.* java.text.* " %>`
- b** `<%@ page import="java.util.*", "java.text.* " %>`
- c** `<%@ page buffer="8kb", session="false" %>`
- d** `<%@ page import="com.manning.servlets.* " %>`
`<%@ page session="true" %>`
`<%@ page import="java.text.*" %>`
- e** `<%@ page bgcolor="navy" %>`
- f** `<%@ page buffer="true" %>`
- g** `<%@ Page language='java' %>`

A: The following table explains why an option is valid or invalid.

Page directive	Valid/ invalid	Reasons
<code><%@ page import="java.util.* java.text.* " %></code>	Invalid:	A comma is required between the values. <code><%@ page import="java.util.*, java.text.* " %></code>
<code><%@ page import="java.util.*", "java.text.* " %></code>	Invalid:	Both packages must be specified in the same string.

continued on next page

Page directive	Valid/ invalid	Reasons
<code><%@ page buffer="8kb", session="false" %></code>	Invalid:	A comma is not allowed between attributes.
<code><%@ page import="com.managing.scwcd.servlets.* " %></code> <code><%@ page session="true" %></code> <code><%@ page import="java.text.*" %></code>	Valid:	The order and placement of page directives do not matter. The import attribute can occur multiple times.
<code><%@ page bgcolor="navy" %></code>	Invalid:	<code>bgcolor</code> is not a valid attribute.
<code><%@ page buffer="true" %></code>	Invalid:	<code>true</code> is not a valid value for the buffer attribute. The value must specify the size of the buffer in kb.
<code><%@ Page language='java' %></code>	Invalid	Directive names, attributes, and values are case sensitive. We must use <code>page</code> and not <code>Page</code> .

10.4 SUMMARY

In this chapter, we examined JavaServer Pages as a web scripting methodology. We learned the basic rules of the six JSP syntax elements—directives, declarations, scriptlets, expressions, actions, and comments—and we examined the first four in depth. We learned that JSP pages are translated into servlet instances before serving the client's requests, and we reviewed the seven phases of the JSP page life cycle. We then looked at the three life-cycle methods—`jspInit()`, `_jspService()`, and `jspDestroy()`—and how they are used in the initialization, servicing, and destruction of a JSP page.

Through its 12 attributes, a page directive provides information about the overall properties of a JSP page to the JSP engine. We need to understand all of the attributes for writing real-life JSP pages, but in preparing for the exam, it is especially important to understand `import`, `session`, `errorPage`, and `isErrorPage`.

In the next chapter, we will continue our discussion of JavaServer Pages as we examine some of the more advanced features that form a logical extension of the servlet technology.

10.5 REVIEW QUESTIONS

1. Consider the following code and select the correct statement about it from the options below. (Select one)

```
<html><body>
  <%! int aNum=5 %>
  The value of aNum is <%= aNum %>
</body></html>
```

- a It will print "The value of aNum is 5" to the output.
- b It will flag a compile-time error because of an incorrect declaration.

- c It will throw a runtime exception while executing the expression.
 - d It will not flag any compile-time or runtime errors and will not print anything to the output.
2. Which of the following tags can you use to print the value of an expression to the output stream? (Select two)
- a `<%@ %>`
 - b `<%! %>`
 - c `<% %>`
 - d `<%= %>`
 - e `<%-- --%>`
3. Which of the following methods is defined by the JSP engine? (Select one)
- a `jspInit()`
 - b `_jspService()`
 - c `_jspService(ServletRequest, ServletResponse)`
 - d `_jspService(HttpServletRequest, HttpServletResponse)`
 - e `jspDestroy()`
4. Which of the following exceptions may be thrown by the `_jspService()` method? (Select one)
- a `javax.servlet.ServletException`
 - b `javax.servlet.jsp.JSPException`
 - c `javax.servlet.ServletException` and `javax.servlet.jsp.JSPException`
 - d `javax.servlet.ServletException` and `java.io.IOException`
 - e `javax.servlet.jsp.JSPException` and `java.io.IOException`
5. Write the name of the method that you can use to initialize variables declared in a JSP declaration in the space provided. (Write only the name of the method. Do not write the return type, parameters, or parentheses.)
- a [_____]
6. Which of the following correctly declares that the current page is an error page and also enables it to take part in a session? (Select one)
- a `<%@ page pageType="errorPage" session="required" %>`
 - b `<%@ page isErrorPage="true" session="mandatory" %>`
 - c `<%@ page errorPage="true" session="true" %>`
 - d `<%@ page isErrorPage="true" session="true" %>`
 - e None of the above.



CHAPTER 11

The JSP technology model—advanced topics

- | | |
|--|--|
| 11.1 Understanding the translation process 189 | 11.3 Understanding JSP page scopes 207 |
| 11.2 Understanding JSP implicit variables and JSP implicit objects 198 | 11.4 JSP pages as XML documents 211 |
| | 11.5 Summary 215 |
| | 11.6 Review questions 216 |

EXAM OBJECTIVES

- 6.3** Write a JSP Document (XML-based document) that uses the correct syntax. (Section 11.4)
- 6.5** Given a design goal, write JSP code using the appropriate implicit objects: request, response, out, session, config, application, page, pageContext, and exception. (Sections 11.2 and 11.3)

INTRODUCTION

In chapter 10, “The JSP technology model—the basics,” we reviewed the basic elements of JSP pages. In this chapter, we continue our discussion of the JSP technology model by examining some of the more advanced features of the JSP framework.

11.1 UNDERSTANDING THE TRANSLATION PROCESS

As we discussed in chapter 10, the first phase of the life cycle of a JSP page is the translation phase, in which the JSP page is translated into a Java file containing the corresponding servlet. The JSP engine parses the JSP page and applies the following rules for translating the JSP elements to the servlet code:

- Some directives are used by the JSP engine during the translation phase to generate Java code. For example, the `import` attribute of the `page` directive aids in generating import statements, while the `info` attribute aids in implementing the `getServletInfo()` method of the generated servlet class. Some directives just inform the engine about the overall properties of the page; for instance, the `language` attribute informs the engine that we are using Java as the scripting language, and the `pageEncoding` attribute informs the engine of the character encoding of the page.
- All JSP declarations become a part of the generated servlet class. They are copied as is. Thus, variables declared in a JSP declaration become instance variables, and the methods declared in a JSP declaration become instance methods of the servlet.
- All JSP scriptlets become a part of the generated `_jspService()` method. They are copied as is. Thus, variables declared in a scriptlet become local variables of the `_jspService()` method. We cannot declare a method in a scriptlet since we cannot have methods declared inside other methods in Java.
- All JSP expressions become a part of the `_jspService()` method. They are wrapped inside `out.print()`.
- All JSP actions are replaced by calls to vendor-specific classes.
- All JSP comments are ignored.
- Any other text becomes part of the `_jspService()` method. It is wrapped inside `out.write()`. This text is also called *template text*.

In the following sections, we will look at some of the implications of these translation rules.

11.1.1 Using scripting elements

Since the declarations, scriptlets, and expressions allow us to write scripting language code in JSP pages, these elements are collectively referred to as the *scripting elements*. We use Java as the scripting language, and consequently, the rules of the Java programming language govern the compile-time and runtime behavior of the code in the scripting elements. Let's examine them one by one with examples.

Order of declarations

Because all the variables and methods defined in the declarations of a JSP page become members of the generated servlet class, their order of appearance in a page does not matter. The following example (listing 11.1) highlights this behavior.

Listing 11.1 area.jsp

```
<html>
<body>
    Using pi = <%=pi%>, the area of a circle<br>
    with a radius of 3 is <%=area(3)%>

<%!
    double area(double r)
    {
        return r*r*pi;
    }
%>

<%! final double pi=3.14159; %>
</body>
</html>
```

In this case, even though the constant `pi` and the method `area()` are used before they are defined, the page will translate, compile, and run just fine, printing the following output:

```
Using pi = 3.14159, the area of a circle
with a radius of 3 is 28.27431
```

Order of scriptlets

Since scriptlets become a part of the `_jspService()` method in the generated servlet, the variables declared in a scriptlet become local variables of the method; consequently, their order of appearance is important. The following code demonstrates this:

```
<html>
<body>
<% String s = s1+s2; %>
<%! String s1 = "hello"; %>
<% String s2 = "world"; %>
<% out.print(s); %>
</body>
</html>
```

Error: undefined variable s2
Member variable s1
Local variable s2

In this example, `s` and `s2` are declared in a scriptlet while `s1` is declared in a declaration. Because `s2` is used before it is declared, this code will not compile.

Initialization of the variables

In Java, instance variables are automatically initialized to their default values, while local variables must be initialized explicitly before they are used. Hence, the variables declared in JSP declarations are initialized to their default values, while the variables declared in JSP scriptlets must be initialized explicitly before they are used. Consider the following example:

```
<html>
<body>
  <%! int i; %>
  <% int j; %>
  The value of i is <%= i++ %> <br>
  The value of j is <%= j++ %> <br>
</body>
</html>
```

OK: i is 0 by default

Error: j not initialized

The variable `i`, declared using a declaration (`<%! . . . %>`), becomes an instance variable of the generated class and is initialized to 0. The variable `j`, declared using a scriptlet (`<% . . . %>`), becomes a local variable of the generated method `_jspService()` and remains uninitialized. Since Java requires local variables to be initialized explicitly before use, this code is invalid and will not compile.

Another important thing to remember is that the instance variables are created and initialized only once, when the JSP container instantiates the servlet. Thus, variables declared in JSP declarations retain their values across multiple requests. On the other hand, local variables are created and destroyed for every request. Thus, variables declared in a scriptlet do not retain their values across multiple requests and are reinitialized each time the JSP container calls `_jspService()`.

To make the above code compile, we have to initialize `j` as

```
<% int j=0; %>
```

Now, if we access the above page multiple times, the value of `i` will get incremented, printing a new value each time, while the value of `j` will always be printed as 0.

11.1.2 Using conditional and iterative statements

Scriptlets are used for embedding computational logic, and frequently this logic includes conditional and iterative statements. For example, the following scriptlet code uses a conditional statement to check a user's login status, and based on that status, it displays an appropriate message:

```
<%
  boolean isUserLoggedIn = ... //get login status
  if (isUserLoggedIn)
  {
    out.print("<h3>Welcome!</h3>");
  }
  else
  {
```

```

        out.println("Hi! Please log in to access the member's area.<br>");
        out.println("<A href='login.jsp'>Login</A>");
    }
%>

```

If we want to include a large amount of HTML within the body of a conditional statement, we can avoid writing multiple `out.println()` statements by spanning the conditional statement across multiple scriptlets in the JSP page, as shown in the following example:

```

<html><body>

<%
    boolean isUserLoggedIn = ... //get login status
    if (isUserLoggedIn)
    {
%>

    <h3>Welcome!</h3>
    A lot of HTML here...

<%
    }
    else
    {
%>

    Hi! Please log in to access the member's area.
    <A href="login.jsp">login</A>
    A lot of HTML here...

<%
    }
%>

</body></html>

```

In the above code snippet, the `if-else` statement is spread across three scriptlets. At runtime, the first scriptlet gets the login status of the user and assigns it to the boolean variable `isUserLoggedIn`. If the value of this variable is `true`, then the HTML code between the first scriptlet and the second scriptlet is included in the output stream. If the value is `false`, then the HTML code between the second scriptlet and the third scriptlet is included in the output stream.

Note the usage of the curly braces to mark the beginning and end of the Java programming language code blocks. Omitting the braces might cause an error at compile time or an undesired behavior at runtime. For example:

```

<% if (isUserLoggedIn) %>
Welcome, <%= userName %>!

```

will be translated to

```

if (isUserLoggedIn)
out.write("Welcome, ");
out.print(userName);

```

In this case, the statement `out.print(userName);` will be executed even if the value of `isUserLoggedIn` is `false`. The correct way to write this is

```
<% if (isUserLoggedIn)
{
%>
    Welcome, <%= userName %>!
<%
}
%>
```

Like conditional statements, iterative statements can also span across multiple scriptlets, with regular HTML code in between the scriptlets. Such constructs are commonly used for displaying long lists of values in a tabular format. The following example illustrates this usage:

```
<html><body>
    List of logged in users:
<table>

<tr>
    <th> Name </th>
    <th> email </th>
</tr>

<%
    User[] users = //get an array of logged in users

    for(int i=0; i< users.length; i++)
    {
%>

<tr>
    <td> <%= users[i].name %> </td>
    <td> <%= users[i].email %> </td>
</tr>

<%
    } // For loop ends
%>

</table>
</body></html>
```

The above code uses two scriptlets to enclose HTML code within a `for` loop; the first scriptlet opens the loop block and the second scriptlet closes the loop block. Notice that the HTML code between the two scriptlets contains only one row and that it embeds JSP expressions that use the loop variable `i` declared by the previous scriptlet.

At request time, the loop may be executed zero or more times based on the length of the `users` array. For each execution, the HTML code will insert one row into the table. Thus, if the length of the `users` array is 9, it will create a table with nine rows in the output stream. Also, since the scriptlet will increment the value of the variable

i each time after the loop is executed, the expressions within the HTML code will index into a different element in the `users` array with each iteration.

Thus, we can generate a variable-sized table containing dynamic rows and columns with the help of multiple scriptlets and expressions.

11.1.3 Using request-time attribute expressions

JSP expressions are not always written to the output stream of the JSP page; they can also be used to pass values to action attributes:

```
<% String pageURL = "copyright.html"; %>
<jsp:include page="<%= pageURL %%" />
```

In this case, the value of the JSP expression `<%= pageURL %>` does not go into the output stream. It is evaluated at request time, and its value is assigned to the page attribute of the `jsp:include` action. An expression used in this way to pass a value to an action attribute is called a *request-time attribute expression*.

An important point to remember here is that such a mechanism of providing request-time attribute values cannot be used in directives, because directives have translation time semantics; this means that the JSP engine uses the directives during the page translation time only. Thus, the two directives in the following example are not valid:

```
<%!
    String bSize = "32kb";
    String pageUrl = "copyright.html";
%>

<%@ page buffer="<%= bSize %%" %>
<%@ include file="<%= pageUrl %%" %>
```

11.1.4 Using escape sequences

Like any programming language, the JSP scripting language also has special characters that have a specific meaning from the JSP parser's point of view. These characters are the single quote, the double quote, the backslash, and the character sequences `<%@`, `<%!`, `<%=`, `<%`, `%>`, `<%--`, and `--%>`. To use them in a manner other than as special characters, we have to use a backslash and create an escape sequence in order to instruct the parser not to use them as special characters. Let's examine the different situations where we must use the escape sequences.

In template text

All scripting elements—declarations (`<%!`), scriptlets (`<%`), and expressions (`<%=`)—start with the characters `<%`. Hence, while parsing the JSP page, the parser looks for the character sequence `<%` to find the start of a tag. If we want to use the string literal `<%` in a normal template text as is, we have to escape the character `%` with a backslash (`\`), as shown in the following example:

```

<html><body>
  The opening tag of a scriptlet is <\%
  The closing tag of a scriptlet is %>
</body></html>

```

Note in this example that we can use the sequence %> without using an escape character, because the parser is not looking for that sequence while parsing the text.

In scripting elements

All scripting elements end with the tag %>. Therefore, after reading the opening tag, the parser looks for the character sequence %> to find the end of the tag. If we want to use the string literal %> within a scripting element as is, we have to escape the character > with a backslash, as shown in the following example:

```

<html><body>
  <%= "The opening tag of a scriptlet is <% " %>
  <%= "The closing tag of a scriptlet is %\>" %>
</body></html>

```

Note that we can use the sequence <% without an escape character, because the parser is not looking for that sequence; it is already in the middle of parsing the scripting element, which is an expression.

In attributes

For string literals used in attribute values, we have to escape the special characters with a backslash character. Consider the following snippet:

```

<%@ page info="A sample use of ', \", \\, <\%, and %\> characters. " %>
<html><body>
  <%= getServletInfo() %>
</body></html>

```

This code will generate the following into the output HTML:

```

A sample use of ', ", \, <%, and %> characters.

```

Note that we have used a backslash for the double quote but not for the single quote. This is because the value is enclosed in a pair of double quotes. If we use a pair of single quotes to enclose the entire value, then we have to escape the single quote appearing within the value as shown here:

```

<%@ page info='A sample use of \', ", \\, <\%, and %\> characters. ' %>

```

In case of request-time attribute expressions, we cannot use a pair of double quotes within a pair of double quotes or a pair of single quotes enclosed within a pair of single quotes. Thus, the following is invalid:

```

<jsp:include page="<%= "copyright.html" %>" />

```

It can be rectified either by using the escape sequence `\"` or by using a pair of single quotes for the entire value and a pair of double quotes for the string literal in the JSP expression:

```
<jsp:include page='<%= "copyright.html" %>' />
<jsp:include page="<%= \"copyright.html\" %>" />
```

NOTE Some browsers will not render content between tags it doesn't understand. For example, Internet Explorer 6.0 on Windows XP SP2 (and possibly earlier versions) does not render the page as expected. To ensure you get a `<%` and `%>` on your web page, you can replace the `<` and `>` characters with their HTML-escaped equivalents:

```
<html><body>
    The opening tag of a scriptlet is &ltt;%
    The closing tag of a scriptlet is %&gtt;
</body></html>
```

As you can see, this also eliminates the problem of `<%` and `%>` being parsed as special characters.

Quizlet

Q: Explain whether the following are valid or invalid JSP constructs.

- a** `<%=myObj.m1(); %>`
- b** `<% int x=4, y=5; %>`
`<%=x=y%>`
- c** `<% myObj.m1(); %>`

A: **a** Invalid: The `=` sign makes it a JSP expression. However, JSP expressions are not terminated with a semicolon. The generated servlet code will cause a syntax error:

```
out.print(myObj.m1());
```

b Valid: `<%=x=y%>` will be translated to:

```
out.print(x=y);
```

The value of `y` is assigned to `x`, and the new value of `x` is then printed out. The output will be 5.

c Valid: It is a valid scriptlet because a semicolon ends the method call statement. It would be valid even if the method returns a value, because the value would be ignored. Let's look at a similar example:

```
Welcome! You are visitor number <% ++count; %>
```

This is a very common mistake, and it is often difficult to debug. The above code will compile and run without any errors, but it will not print the desired output. This is how the JSP engine will translate the code:

```
out.write("Welcome! You are visitor number ");
++count;
```

It increments the count variable, but does not use `out.print()` to print its value. To print the value, we have to make it an expression by inserting an equals sign (=) after the opening tag (`<%`) and removing the semicolon (;) from the end, as shown here:

```
Welcome! You are visitor number <%= ++count %>
```

Q: What is wrong with the following code?

```
<%@ page language='java' %>
<%
    int x = 0;
    int incr() { return ++x; }
%>
The value of x is <%=incr()%>
```

A: We cannot define methods in a scriptlet. Upon translation of the above code into a servlet, the `_jspService()` method will look like this:

```
public void _jspService(...)
{
    ...other code

    int x = 0;
    int incr() { return ++x; }
    out.write("The value of x is ");
    out.print(incr());
}
```

Since the `incr()` method is declared inside the `_jspService()` method, the code will not compile.

Q: Will the following code compile?

```
<% int x = 3; %>
<%! int x = 5; %>
<%! int y = 6; %>
The sum of x and y is <%=x+y%>
```

A: Yes. It will compile and print

```
The sum of x and y is 9
```

Upon translation of the above code into a servlet, the variable `x` will be declared twice: once global to the class because of the declaration `<%! int x = 5; %>` and once local to the `_jspService()` method because of the scriptlet `<% int x = 3; %>`:

```
public class xyz ...
{
    ...other code
```



```

int x = 5;
int y = 6;

public void _jspService(...)
{
    ...other code

    int x = 3;
    out.write("The sum of x and y is ");
    out.print(x+y);
}
}

```

Since local variables have precedence over global variables, the expression `x+y` evaluates as `3+6`.

Q: What is the output of the following code?

```

<% int i; %>
<%
    for(i=0; i<3; i++)
%>
The value of i is <%=i%>

```

A: This code will translate into

```

int i;
for (int i=0; i<3; i++)
out.write("The value of i is ");
out.print(i);

```

Since we have not enclosed the body of the loop in a block `{ ... }`, it will print

```
The value of i is The value of i is The value of i is 3
```

11.2 UNDERSTANDING JSP IMPLICIT VARIABLES AND JSP IMPLICIT OBJECTS

During the translation phase, the JSP engine declares and initializes nine commonly used variables in the `_jspService()` method. We have already seen the use of one of them, `out`, in some of our previous examples:

```

<html><body>
    <%
        out.print("Hello World! ");
    %>
</body></html>

```

Even though we have not defined the variable `out` in this example, the code will translate, compile, and execute without errors. This is because `out` is one of the nine variables that the JSP engine implicitly makes available to the JSP page. Table 11.1 describes these variables.

Table 11.1 Implicit variables available to JSP pages

Identifier name	Class or interface	Description
<code>application</code>	interface <code>javax.servlet.ServletContext</code>	Refers to the web application's environment
<code>session</code>	interface <code>javax.servlet.http.HttpSession</code>	Refers to the user's session
<code>request</code>	interface <code>javax.servlet.http.HttpServletRequest</code>	Refers to the current request to the page
<code>response</code>	interface <code>javax.servlet.http.HttpServletResponse</code>	Used for sending a response to the client
<code>out</code>	class <code>javax.servlet.jsp.JspWriter</code>	Refers to the output stream for the page
<code>page</code>	class <code>java.lang.Object</code>	Refers to the page's servlet instance
<code>pageContext</code>	class <code>javax.servlet.jsp.PageContext</code>	Refers to the page's environment
<code>config</code>	interface <code>javax.servlet.ServletConfig</code>	Refers to the servlet's configuration
<code>exception</code>	class <code>java.lang.Throwable</code>	Used for error handling

Let's take a look at the way Tomcat 5.0.25 declares these variables in the generated servlet for a JSP file. Follow these steps:

- 1 Create a blank file in the `C:\jakarta-tomcat-5.0.25\webapps\chapter11` directory and name it `implicit.jsp`.
- 2 Start Tomcat.
- 3 From your browser, navigate to the URL `http://localhost:8080/chapter11/implicit.jsp`.

Although you will not see any content in the browser window, the JSP engine will create a Java file named `implicit_jsp.java` in the `C:\Jakarta-tomcat-5.0.25\work\Catalina\localhost\chapter11\org\apache\jsp` directory.¹ This file contains the servlet that corresponds to the JSP file, `implicit.jsp`. The `_jspService()` method of the servlet is shown in Listing 11.2.

¹ The exact name and location of the generated Java source file, and the exact contents of that file, are not the same between various versions of Tomcat.

Listing 11.2 implicit.jsp.java

```
public void _jspService(  
    HttpServletRequest request,  
    HttpServletResponse response)  
    throws java.io.IOException,  
        ServletException  
{  
    ...other code  
  
    PageContext pageContext = null;  
    HttpSession session = null;  
    ServletContext application = null;  
    ServletConfig config = null;  
    JspWriter out = null;  
    Object page = this;  
  
    ...other code  
  
    pageContext = ...//get it from somewhere  
    session = pageContext.getSession();  
    application = pageContext.getServletContext();  
    config = pageContext.getServletConfig();  
    out = pageContext.getOut();  
  
    ...other code  
}
```

As you can see, eight variables are already declared and available within the `_jspService()` method. To declare the ninth one, open the `implicit.jsp` file and add the following line:

```
<%@ page isErrorPage="true" %>
```

Save the file and go to the same URL again. Now you should see the following line added to the generated `implicit.jsp.java` file:

```
Throwable exception =  
    (Throwable) request.getAttribute("javax.servlet.jsp.jspException");
```

Because the page author does not (and cannot) declare these variables explicitly, they are called *implicit variables*. The objects that these variables refer to are created by the servlet container and are called *implicit objects*. The exam may ask you to state their types, scopes, and uses. We will discuss the use of these nine implicit variables first, and then talk about their scopes in section 11.3.

11.2.1 application

The `application` variable is of type `javax.servlet.ServletContext`, and it refers to the environment of the web application to which the JSP page belongs. (We discussed the `ServletContext` class at length in chapter 6, “The servlet container model.”) Thus, the following two scriptlets are equivalent:

```

<%
    String path = application.getRealPath("/WEB-INF/counter.db");
    application.log("Using: "+path);
%>

<%
    String path = getServletContext().getRealPath("/WEB-INF/counter.db");
    getServletContext().log("Using: "+path);
%>

```

11.2.2 session

Before we discuss the `session` implicit variable, let's clarify that the word *session* refers to four different but related things in JSP:

- Session, as in an HTTP session, is a concept that logically groups multiple requests from the same client as part of one conversation. We discussed HTTP sessions in chapter 8, "Session management."
- `session`, as used in a page directive, refers to the attribute named *session*.

```
<%@ page session="true" %>
```

- Its value, which is `true` or `false`, determines whether or not the JSP page participates in an HTTP session.
- `session`, as an implicit object (which we will talk about in this section), refers to the variable `session` of type `javax.servlet.http.HttpSession`.
- Session, as a scope of an object, refers to the lifetime and availability of the object. A session-scoped object persists throughout the life of an HTTP session. We will discuss scope in the next section.

The implicit variable `session` is declared if the value of the `session` attribute of the page directive is `true`. Since by default, the value of the `session` attribute is `true`, this variable is declared and is made available to the page even if we do not specify the page directive. However, if we explicitly set the `session` attribute to `false`, the JSP engine does not declare this variable, and any use of the variable results in an error. The following example demonstrates this:

```

<html>
<body>
    <%@ page session="false" %>    ← session is not used

    Session ID = <%=session.getId()%>    ← Error: undefined
</body>                                   symbol session
</html>

```

In the above example, the page directive sets the `session` attribute to `false` in order to indicate that the current page will not participate in an HTTP session. This makes the implicit variable `session` unavailable in the page. So the line `session.getId()` will generate a compile-time error.

11.2.3 request and response

The request and response implicit variables are of type `javax.servlet.http.HttpServletRequest` and `javax.servlet.http.HttpServletResponse`, respectively. They are passed in as parameters to the `_jspService()` method when the page's servlet is executed upon a client request. We use them in JSP pages in exactly the same way we use them in servlets—that is, to analyze the request and send a response:

```
<%
    String remoteAddr = request.getRemoteAddr();
    response.setContentType("text/html; charset=ISO-8859-1");
%>

<html><body>

Hi! Your IP address is <%=remoteAddr%>

</body></html>
```

11.2.4 page

The implicit variable `page` is of class `java.lang.Object`, and it refers to the instance of the generated servlet. It is declared as

```
Object page = this; //this refers to the instance of this servlet.
```

This variable is rarely used. In fact, since it is a variable of type `Object`, it cannot be used to directly call the servlet methods:

```
<%= page.getServletInfo()           %>  ← Error
<%= ((Servlet)page).getServletInfo() %>  ← OK: typecast
<%= this.getServletInfo()           %>  ← OK
```

The first expression will generate a compile-time error indicating that `getServletInfo()` is not a method of `java.lang.Object`.

In the second expression, we have typecast the `page` reference to `Servlet`. Since `page` refers to the generated class and the class implements the `Servlet` interface, it is a valid cast. Also, because `getServletInfo()` is a method of the `Servlet` interface, the expression will compile and execute without errors. Note that, in this case, the `page` variable could also be cast to `JspPage` or `HttpJspPage` since these two interfaces are derived from the `Servlet` interface and are implemented by the generated servlet class.

In the third expression, we are using the Java keyword `this` to refer to the generated servlet. Therefore, it will also compile and execute without errors.

11.2.5 pageContext

The `pageContext` variable is of type `javax.servlet.jsp.PageContext`. The `PageContext` class is an abstract class, and the JSP engine vendor provides its concrete subclass. It does three things:

- Stores references to the implicit objects. If you look at the generated servlet code for the `implicit.jsp` file (listing 11.2), you will see that the `session`, `application`, `config`, and `out` implicit variables are initialized using the objects retrieved from `pageContext`. The `pageContext` object acts as a one-stop place for managing all the other objects, both user-defined and implicit, used by the JSP page, and it provides the getter methods to retrieve them.
- Provides convenience methods to get and set attributes in different scopes. These are explained in section 11.3.4.
- Provides convenience methods, described in table 11.2, for transferring requests to other resources in the web application.

Table 11.2 Convenience methods of `javax.servlet.jsp.PageContext` for transferring requests to other resources

Method	Description
<code>void include(String relativeURL)</code>	Includes the output of another resource in the output of the current page. Same as <code>ServletRequest.getRequestDispatcher().include();</code>
<code>void forward(String relativeURL)</code>	Forwards the request to another resource. Same as <code>ServletRequest.getRequestDispatcher().forward();</code>

For example, to forward a request to another resource from a servlet, we have to write the following two lines:

```
RequestDispatcher rd = request.getRequestDispatcher("other.jsp");
rd.forward(request, response);
```

In a JSP page, we can do that in just one line by using the `pageContext` variable:

```
pageContext.forward("other.jsp");
```

For a complete list of all the methods of the `PageContext` class, please refer to the JSP API.

11.2.6 out

The implicit variable `out` is of type `javax.servlet.jsp.JspWriter`. This variable is the workhorse of JSP pages. We use it directly in scriptlets and indirectly in expressions to generate HTML code:

```
<% out.print("Hello 1"); %>
<%= "Hello 2" %>
```

For both of the above lines, the generated servlet code will use the `out` variable to print the values:

```
public void _jspService(...)
{
```

```

//other code
out.print("Hello 1");
out.print("Hello 2");
}

```

The `JspWriter` class extends `java.io.Writer` and inherits all the overloaded `write()` methods. On top of these methods, `JspWriter` adds its own set of overloaded `print()` and `println()` methods for printing out all the primitive data types, Strings, and the user-defined objects in the output stream. The following example prints out different data types using the `out` variable:

```

<%
  int anInt = 3;
  Float aFloatObj = new Float(11.6);

  out.print(anInt); //int
  out.print(anInt > 0); //boolean
  out.print(anInt*3.5/100-500); //float expression
  out.print(aFloatObj); //object
  out.print(aFloatObj.floatValue()); //float method
  out.print(aFloatObj.toString()); //String method
%>

```

11.2.7 config

The implicit variable `config` is of type `javax.servlet.ServletConfig`. As we saw in chapter 5 (“Structure and deployment”), each servlet can be passed a separate set of configuration parameters in the deployment descriptor, and the servlet can then retrieve this information using its own copy of the `ServletConfig` object.

Similarly, we can also pass configuration parameters that are specific to a JSP page, which the page can retrieve using the implicit variable `config`. To achieve this, we have to first declare a servlet with a `<servlet-name>` in the deployment descriptor `web.xml`. Then, instead of providing a `<servlet-class>`, we associate the named servlet with the JSP file, using the element `<jsp-file>`. All of the initialization parameters for this named servlet will then be available to the JSP page via the page’s `ServletConfig` implicit object. The `web.xml` file shown in listing 11.3 illustrates this.

Listing 11.3 Configuring `InitTestServlet` and mapping it to a JSP page in `web.xml`

```

<?xml version="1.0" encoding="ISO-8859-1"?>

<!DOCTYPE web-app
  PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
  "http://java.sun.com/dtd/web-app_2_3.dtd">

<web-app>

  <servlet>
    <servlet-name>InitTestServlet</servlet-name>
    <jsp-file>/initTest.jsp</jsp-file>
  
```

```
<init-param>
  <param-name>region</param-name>
  <param-value>North America</param-value>
</init-param>
</servlet>
</web-app>
```

The deployment descriptor in listing 11.3 declares a servlet named `InitTestServlet` and maps it to the JSP file `<document root>/initTest.jsp`. Then it specifies an initialization parameter named `region` with the value of `North America` for the servlet. This information can be retrieved by `initTest.jsp` using the implicit variable `config`, as shown in listing 11.4.

Listing 11.4 `initTest.jsp`

```
<html><body>
  Servlet Name = <%=config.getServletName()%><br>
  Parameter region = <%=config.getInitParameter("region")%>
</body></html>
```

When the JSP page `initTest.jsp` is accessed as a servlet, using the URL `http://localhost:8080/chapter11/servlet/InitTestServlet`, it will print the following output:

```
Servlet Name = InitTestServlet
Parameter region = North America
```

However, if this page is accessed directly using the `initTest.jsp` page's actual URL, `http://localhost:8080/chapter11/initTest.jsp`, then the configuration in listing 11.4 for `InitTestServlet` is not used. This is because the JSP engine creates two different instances of the generated servlet class—one for accessing it as a named servlet and one for accessing it as a JSP page—and will pass each servlet instance a different `ServletConfig` object. In order to be able to use the same servlet instance—and hence the same configuration—when using either of the URLs mentioned above, we have to explicitly map the JSP page's URL in the deployment descriptor file using the `<servlet-mapping>` element:

```
<servlet-mapping>
  <servlet-name>InitTestServlet</servlet-name>
  <url-pattern>/initTest.jsp</url-pattern>
</servlet-mapping>
```

When mapped this way, the container will create only one instance of the generated servlet class and requests for both the URLs will be served by the same instance, thus guaranteeing the same configuration.

The rules for the `<servlet-mapping>` for JSP pages are the same as in the case of normal servlets. We can provide just about any URL pattern, and each time the pattern matches the client's request URL, the container will execute the specified JSP page (actually, the generated servlet for the JSP page).

11.2.8 exception

This implicit variable is of type `java.lang.Throwable`. The exception variable is available to the pages that act as error handlers for other pages. Recall from chapter 11 that error-handler pages have the page directive attribute `isErrorPage` set to `true`. Consider the following two JSP code examples:

Example 1

```
<html><body>
  <%@ page isErrorPage='true' %>
  Msg: <%=exception.toString() %>
</body></html>
```

OK: exception defined implicitly

Example 2

```
<html><body>
  Msg: <%=exception.toString() %>
</body></html>
```

Error: exception not defined

In example 1, the engine defines the exception variable implicitly because the attribute `isErrorPage` is set to `true`. The exception variable refers to the uncaught `java.lang.Throwable` object thrown by a page that uses this page as its error handler.

In example 2, the engine does not define the variable `exception` implicitly because the attribute `isErrorPage` has a default value of `false`.

Quizlet

Q: What is wrong with the following code?

```
<%!
  public void jspInit() {
    application.getInitParameter("Region");
  }
%>
```

A: Recall that all the implicit variables are automatically declared by the JSP engine in the generated `_jspService()` method, which means they are available only in scriptlets and expressions. In a declaration, we have to explicitly declare the variables:

```
<%!
  public void jspInit() {
    ServletContext application = this.getServletContext();
    application.getInitParameter("Region");
  }
%>
```

11.3 UNDERSTANDING JSP PAGE SCOPES

In chapter 4, “The servlet model,” we introduced the concept of *scope*, which is the way that data is shared between servlets using the three container objects, `ServletContext`, `HttpSession`, and `ServletRequest`. The scopes associated with these three container objects are, respectively, the application scope, the session scope, and the request scope. The JSP technology, since it is based on servlet technology, also uses the three scopes, which are referred to in JSP pages as *application*, *session*, and *request* scopes. In addition, JSP pages have a fourth scope, the *page* scope, which is maintained by the container object `PageContext`.

All of the implicit objects as well as the user-defined objects in a JSP page exist in one of these four scopes, described in table 11.3. These scopes define the existence and accessibility of objects from within the JSP pages and servlets.

Table 11.3 Scopes of objects in JSP pages

Scope Name	Existence and Accessibility
Application	Limited to a single web application
Session	Limited to a single user session
Request	Limited to a single request
Page	Limited to a single page (translation unit) and a single request

As shown in table 11.3, objects in the *application* scope are the most accessible and objects in the *page* scope are the least accessible. Let’s take a closer look at these scopes.

11.3.1 Application scope

Application-scoped objects are shared across all the components of the web application and are accessible for the life of the application. These objects are maintained as attribute-value pairs by an instance of the `ServletContext` class. In a JSP page, this instance is available in the form of the implicit object `application`. Thus, to share objects at the application level, we use the `setAttribute()` and `getAttribute()` methods of the `ServletContext` interface.

11.3.2 Session scope

Objects in the session scope are shared across all the requests that belong to a single-user session and are accessible only while the session is valid. These objects are maintained as attribute-value pairs by an instance of the `HttpSession` class. In a JSP page, this instance is available in the form of the implicit object `session`. Thus, to share objects at the session level, we can use the `session.setAttribute()` and `session.getAttribute()` methods.

In the following example, the `login.jsp` page adds the user ID to the session scope so that the `userProfile.jsp` page can retrieve it:

```

<%--
    Add the userId to the session
--%>
<%
    String userId = // getUserLoggedIn
    session.setAttribute("userId", userId);
%>

```

In the file `userProfile.jsp`:

```

<%--
    Retrieve the userId from the session
--%>
<%
    String userId = (String) session.getAttribute("userId")

    //use the userId to retrieve user details.
    String name = getNameById(userId);
%>

    User Name is: <%=name%>

```

Here, the session scope is used to make the username and ID available to all the requests in the session.

11.3.3 Request scope

Objects in the request scope are shared across all the components that process the same request and are accessible only while that request is being serviced. These objects are maintained as attribute-value pairs by an implementation instance of the interface `HttpServletRequest`. In a JSP page, this instance is available in the form of the implicit object `request`. Thus, we can add attributes to the request in one page and forward the request to another page. The second page can then retrieve these attributes to generate a response.

In this example, the file `login.jsp` creates a user object and adds it to the request, and then forwards the request to `authenticate.jsp`:

```

<%
    //Get login and password information from the request object
    //and file it in a User Object.

    User user = new User();
    user.setLogin(request.getParameter("login"));
    user.setPassword(request.getParameter("password"));

    //Set the user object in the request scope for now
    request.setAttribute("user", user);

    //Forward the request to authenticate.jsp
    pageContext.forward("authenticate.jsp");
    return;
%>

```

In the file `authenticate.jsp`:

```

<%
    //Get user from the forwarding page
    User user = (User) request.getAttribute("user");

    //Check against the database.
    if (isValid(user))
    {
        //remove the user object from request scope
        //and maintain it in the session scope
        request.removeAttribute("user");
        session.setAttribute("user",user);

        pageContext.forward("account.jsp");
    }
    else
    {
        pageContext.forward("loginError.jsp");
    }
    return;
%>

```

Here, the page (`login.jsp`) adds a `User` object to the request scope and forwards the request to `authenticate.jsp`. At that point, `authenticate.jsp` validates the user information against the database and, depending on the outcome of the authentication process, either transfers the object into the session scope and forwards the request to `account.jsp`, or forwards the request to `loginError.jsp`, which generates an appropriate response by using the `User` object. We call `return;` after forwarding the request to prevent writing anything to the output stream after the request is forwarded because doing so would throw an `IllegalStateException`.

11.3.4 Page scope

Objects in the page scope are accessible only in the translation unit in which they are defined. They do not exist outside the processing of a single request within a single translation unit. These objects are maintained as attribute-value pairs by an instance of a concrete subclass of the abstract class `PageContext`. In a JSP page, this instance is available in the form of the implicit object `pageContext`.

The use of the page scope and the `pageContext` container object may not be obvious right now, but it will become clearer when we will learn about the use of JavaBeans in chapter 14, “Using JavaBeans,” and custom tags in chapters 16, “Developing classic custom tag libraries,” and 17, “Developing ‘Simple’ custom tag libraries.” The only way for actions (standard JSP actions and user-defined custom tags) to share data and JavaBean objects with other actions or custom tags appearing in the same JSP page (translation unit) and in the same request thread is to use the `pageContext` implicit object and the page scope.

To share objects in the page scope, we can use the two methods defined by `PageContext`, shown in table 11.4.

Table 11.4 Convenience methods of `javax.servlet.jsp.PageContext`

Method	Description
<code>void setAttribute(String name, Object attribute)</code>	Adds an attribute to the page scope
<code>java.lang.Object getAttribute(String name)</code>	Returns the object associated with the name in the page scope or null if not found

The `PageContext` object also provides a common and convenient way to handle all of the objects in all of the scopes. Table 11.5 describes the defined constants and methods used for this purpose.

Table 11.5 Convenience scope-handling constants and methods of `javax.servlet.jsp.PageContext`

Member	Description
Integer constants that work with scopes	
<code>static final int APPLICATION_SCOPE</code>	Indicates application scope
<code>static final int SESSION_SCOPE</code>	Indicates session scope
<code>static final int REQUEST_SCOPE</code>	Indicates request scope
<code>static final int PAGE_SCOPE</code>	Indicates page scope
Methods that accept scope constants	
<code>void setAttribute(String name, Object object, int scope);</code>	Sets the attribute in the specified scope
<code>java.lang.Object getAttribute(String name, int scope);</code>	Returns the object associated with the name in the specified scope or null if not found
<code>void removeAttribute(String name, int scope)</code>	Removes the object associated with the specified name from the given scope
<code>java.util.Enumeration getAttributeNamesIn Scope(int scope)</code>	Enumerates all the attributes in a given scope
Convenience scope search methods	
<code>Object findAttribute(java.lang.String name)</code>	Searches for the named attribute in page, request, session (if valid), and application scope(s) in this order and returns the associated value
<code>int getAttributesScope(String name)</code>	Gets the scope in which a given attribute is defined

The six implicit objects—`response`, `out`, `page`, `pageContext`, `config`, and `exception`—are also considered to have page scope by the JSP specification. They are all maintained by the `pageContext` container object and can be obtained using their respective getter methods.

However, their existence and accessibility is somewhat different from the user-defined objects. Although the JSP specification defines these implicit objects as being in the page scope, they are not logically restricted to that scope. For example, the `page` implicit object refers to the generated servlet instance. For the same servlet instance, if there are multiple threads that are serving multiple requests, possibly even in multiple sessions, then the same implicit object `page` is shared by all those threads and is thus accessible by all the requests and in all the sessions. Similarly, the `Servlet-Config` object, referred to as the implicit variable `config`, is also shared by multiple threads that are serving multiple requests, possibly in multiple sessions.

Quizlet

- Q:** Objects in which scope are accessible to all of the web applications of a servlet container?
- A:** There is no scope that can share objects across multiple web applications. To do that, we have to either use `ServletContext.getContext()` as explained in chapter 4 or use other mechanisms, such as an external database.

11.4 JSP PAGES AS XML DOCUMENTS

The JSP specification defines two sets of syntax for authoring JSP pages: standard JSP syntax format and XML syntax format. JSP files that use the standard syntax are called JSP pages, while the JSP files that use the XML syntax are called JSP documents.

We have already seen the standard syntax of the elements of a JSP page in the previous chapter. In this section, we will learn about the XML-style tags that make up a JSP document. The exam does not require you to know all the details of the XML format, but you are expected to know the XML-based tags for writing the directives and the scripting elements.

The best way to learn and remember the XML-style tags is to compare them with their JSP-style counterparts. So let's first look at the code for the JSP document `counter_xml.jsp` (listing 11.5), which is an XML equivalent of the code for the JSP page `counter.jsp` that we saw in chapter 11 (listing 11.1).

Listing 11.5 `counter_xml.jsp`

```
<jsp:root
  xmlns:jsp="http://java.sun.com/JSP/Page"
  version="2.0">

  <html><body>

  <jsp:directive.page language="java" />

  <jsp:declaration>
    int count = 0;
  </jsp:declaration>
```

```

<jsp:scriptlet>
    count++;
</jsp:scriptlet>

<jsp:text>
    Welcome! You are visitor number
</jsp:text>

<jsp:expression>
    count
</jsp:expression>

</body></html>

</jsp:root>

```

The JSP document in listing 11.5 counts the number of times it is visited during a server session. We will discuss its XML elements and tags in the following sections.

You should remember two important points about using the XML and JSP syntax together:

- The standard JSP tags and XML-based tags cannot be mixed within a single JSP page.
- A page written in one syntax format can, however, include or forward to a page in the other syntax format by using either directives or actions.

As with directives and actions in the JSP syntax format, the following rules apply to all of the elements in XML syntax format:

- The tag names, attribute names, and attribute values are all case sensitive.
- The value must be surrounded by a pair of single or double quotes.
- A pair of single quotes is equivalent to a pair of double quotes.
- There must be no space between the equals sign (=) and the value.

11.4.1 The root element

As seen in listing 11.5, the XML syntax requires that the entire JSP page be enclosed in a single root element:

```

<jsp:root
    xmlns:jsp="http://java.sun.com/JSP/Page"
    version="2.0" >

    Rest of the page

</jsp:root>

```

The attribute-value pair `xmlns:jsp="http://java.sun.com/JSP/Page"` tells the JSP engine that the prefix `jsp` is used to identify the tags in the library specified by the URI `http://java.sun.com/JSP/Page`. This library contains the standard elements defined by the JSP specification. Thus, all the JSP tags in listing 11.5

are of the form `<jsp: . . . >`. The attribute `version` informs the engine about the version of the JSP specification used in the page. Both of the attributes are mandatory.

`xmlns` stands for *XML Name Space*. Functionally, it is the same as the `prefix` attribute of a `taglib` directive in the JSP syntax. Hence, it is also used to specify the use of the custom tag libraries in the page:

```
<jsp:root
    xmlns:jsp="http://java.sun.com/JSP/Page"
    xmlns:myLib="www.someserver.com/someLib"
    version="2.0" >
```

Rest of the page

```
</jsp:root>
```

The attribute-value pair `xmlns:myLib="www.someserver.com/someLib"` tells the JSP engine that the page uses custom tags of the form `<myLib: . . . >` and that the details of these custom tags are located in the library indicated by the URI `www.someserver.com/someLib`.

Tag libraries will be discussed in chapter 15, “Using custom tags,” chapter 16, “Developing classic custom tag libraries,” and chapter 17, “Developing ‘Simple’ custom tag libraries.”

NOTE There is no equivalent to `<jsp:root>` in the JSP syntax.

11.4.2 Directives and scripting elements

There are only two directives in the XML format: `page` and `include`:

```
<jsp:directive.page ...attributeList... />
<jsp:directive.include ...attributeList... />
```

There is no `taglib` directive in the XML format since the tag library information is specified in the root element. The attributes and use of the `page` and `include` directives are the same in both XML syntax and JSP syntax.

The three scripting elements—declarations, scriptlets, and expressions—use the following syntax, respectively:

```
<jsp:declaration>
    Any valid Java declaration statements
</jsp:declaration>

<jsp:scriptlet>
    Any valid Java code
</jsp:scriptlet >

<jsp:expression>
    Any valid Java expression
</jsp:expression >
```

This syntax for `expression` is used when we want to write the expression values to the output HTML. For a request-time attribute expression, we have to use `%= . . . %`, as in the following example:


```

<jsp:scriptlet>
    String pageURL = "copyright.html";
</jsp:scriptlet>

<jsp:include page="%=pageURL%" />

```

In this case, the value of the `pageURL` variable is sent as a parameter to the `include` action.

Compare this with the JSP style, where the same syntax, `<%= ... %>`, is used for both purposes.

11.4.3 Text, comments, and actions

A major difference between the JSP and XML syntax is the placement of normal text. The JSP syntax allows us to incorporate text into the page without using any special tags. However, the XML syntax requires us to embed the text between `<jsp:text>` and `</jsp:text>`, as shown here:

```

<html><body>
    <jsp:text>Have a nice day!</jsp:text>
</body></html>

```

The JSP specification does not stipulate any special tag for writing comments in XML format. However, since XML-based JSP pages are treated as XML documents, we should use the standard XML-style comments:

```

<!-- comment here -->

```

The standard JSP actions use the same tags in both the JSP syntax and XML syntax. This is an example of an `include` action in either syntax:

```

<jsp:include page="someOtherPage.jsp" />

```

Quizlet

Q: What is wrong with the following code?

```

<jsp:root
    xmlns:jsp=http://java.sun.com/JSP/Page
    version="2.0" >

    2 + 3 = <jsp:expression>=2+3</jsp:expression>

</jsp:root>

```

A: `xmlns:jsp` is the attribute and `"http://java.sun.com/JSP/Page"` is the value. Thus the quotes must be around the value as:

```

xmlns:jsp="http://java.sun.com/JSP/Page"

```

Unlike JSP syntax elements, XML syntax elements do not use any extra characters in their tags.

```

<%=2+3%>    ← Valid JSP expression
<jsp:expression>=2+3</jsp:expression>    ← Error: = not required

```

```
<jsp:expression>2+3</jsp:expression> ← Valid XML expression
```

Also, the text must be enclosed in `<jsp:text>`:

```
<jsp:text>2 + 3 = </jsp:text><jsp:expression>2+3</jsp:expression>
```

Q: What is wrong with the following code?

```
<jsp:root
    xmlns:jsp="http://java.sun.com/JSP/Page"
    version="2.0" >

    <jsp:Text>2 + 3 = </jsp:Text>
    <jsp:Expression>2+3</jsp:Expression>

</jsp:root>
```

A: All the tags are case sensitive. We must use `jsp:text` and `jsp:expression` instead of `jsp:Text` and `jsp:Expression`, respectively.

11.5 SUMMARY

In this chapter, we continued our discussion on the JSP technology by looking more closely at the translation phase rules that a JSP engine applies during the page's conversion into a servlet. We started by discussing the various traps and pitfalls involved in using the scripting elements (declarations, scriptlets, and expressions) to declare and initialize variables and to write conditional and iterative statements. A good grasp of these issues is essential when combining the Java programming language with HTML to develop error-free and effective JSP pages.

The JSP engine declares and initializes nine commonly used variables and makes them available to the JSP page's generated servlet code. These are called implicit variables, and they refer to the implicit objects that are created by the servlet container. We discussed the use of these nine implicit variables with examples, and we also saw how they relate to the four scopes: page, request, session, and application.

The JSP specification also supports the use of XML syntax to create JSP pages. Although much of the exam focuses on standard JSP syntax, some knowledge of the XML syntax is expected. We reviewed the XML tags for writing JSP documents, including directives and scripting elements, and compared them with their JSP counterparts.

With the end of this chapter, you should be ready to answer questions about the structure of a JSP page in both the standard and the XML format, the mapping of its elements into a servlet, and the implicitly provided objects with their scopes.

In the next chapter, we will learn different mechanisms provided by the JSP specification for reusing the JSP pages.

11.6 REVIEW QUESTIONS

1. What will be the output of the following code? (Select one)

```
<html><body>
  <% x=3;      %>
  <% int x=5; %>
  <%! int x=7; %>
  x = <%=x%>, <%=this.x%>
</body></html>
```

- a x = 3, 5
- b x = 3, 7
- c x = 5, 3
- d x = 5, 7
- e Compilation error

2. What will be the output of the following code? (Select one)

```
<html><body>
  The value is <%= "" %>
</body></html>
```

- a Compilation error
- b Runtime error
- c The value is
- d The value is null

3. Which of the following implicit objects is not available to a JSP page by default? (Select one)

- a application
- b session
- c exception
- d config

4. Which of the following implicit objects can you use to store attributes that need to be accessed from all the sessions of a web application? (Select two)

- a application
- b session
- c request
- d page
- e pageContext

5. The implicit variable `config` in a JSP page refers to an object of type: (Select one)

- a `javax.servlet.PageConfig`
- b `javax.servlet.jsp.PageConfig`

- c** `javax.servlet.ServletConfig`
 - d** `javax.servlet.ServletContext`
6. A JSP page can receive context initialization parameters through the deployment descriptor of the web application.
- a** True
 - b** False
7. Which of the following will evaluate to `true`? (Select two)
- a** `page == this`
 - b** `pageContext == this`
 - c** `out instanceof ServletOutputStream`
 - d** `application instanceof ServletContext`
8. Select the correct statement about the following code. (Select one)
- ```
<%@ page language="java" %>
<html><body>
 out.print("Hello ");
 out.print("World ");
</body></html>
```
- a** It will print `Hello world` in the output.
  - b** It will generate compile-time errors.
  - c** It will throw runtime exceptions.
  - d** It will only print `Hello`.
  - e** None of above.
9. Select the correct statement about the following code. (Select one)
- ```
<%@ page language="java" %>
<html><body>
<%
    response.getOutputStream().print("Hello ");
    out.print("World");
%>
</body></html>
```
- a** It will print `Hello world` in the output.
 - b** It will generate compile-time errors.
 - c** It will throw runtime exceptions.
 - d** It will only print `Hello`.
 - e** None of above.
10. Which of the following implicit objects does not represent a scope container? (Select one)
- a** `application`
 - b** `session`

- c request
- d page
- e pageContext

11. What is the output of the following code? (Select one)

```
<html><body>
  <% int i = 10;%>
  <% while(--i>=0) { %>
    out.print(i);
  <% } %>
</body></html>
```

- a 9876543210
- b 9
- c 0
- d None of above

12. Which of the following is not a valid XML-based JSP tag? (Select one)

- a <jsp:directive.page />
- b <jsp:directive.include />
- c <jsp:directive.taglib />
- d <jsp:declaration></jsp:declaration>
- e <jsp:scriptlet></jsp:scriptlet>
- f <jsp:expression></jsp:expression>

13. Which of the following XML syntax format tags do not have an equivalent in JSP syntax format? (Select two)

- a <jsp:directive.page/>
- b <jsp:directive.include/>
- c <jsp:text></jsp:text>
- d <jsp:root></jsp:root>
- e </jsp:param>

14. Which of the following is a valid construct to declare that the implicit variable session should be made available to the JSP page? (Select one)

- a <jsp:session>>true</jsp:session>
- b <jsp:session required="true" />
- c <jsp:directive.page>
 <jsp:attribute name="session" value="true" />
</jsp:directive.page>
- d <jsp:directive.page session="true" />
- e <jsp:directive.page attribute="session" value="true" />



CHAPTER 12

Reusable web components

- 12.1 Static inclusion 220
- 12.2 Dynamic inclusion 223
- 12.3 Summary 232
- 12.4 Review questions 232

EXAM OBJECTIVES

- 6.7** Given a specific design goal for including a JSP segment in another page, write the JSP code that uses the most appropriate inclusion mechanism (the include directive or the `jsp:include` standard action).
(Sections 12.1 and 12.2)
- 8.2** Given a design goal, create a code snippet using the following standard actions:
 - `jsp:include`,
 - `jsp:forward`, and
 - `jsp:param`(Section 12.2)

INTRODUCTION

Instead of building a case for reusing software components, in this chapter we will reiterate the well-acknowledged fact that reusable components enhance the productivity and maintainability of applications. In this respect, the JSP specification defines mechanisms that allow us to reuse web components.

Our aim in this chapter is to understand how web components can be reused. Although the exam objective corresponding to this topic looks narrow, it requires that you know a lot more than just syntax in order to answer the questions in the exam.

In the JSP world, reusing web components essentially means including the content or the output of another web component in a JSP page. This can be done in one of two ways: statically or dynamically. Static inclusion involves including the contents of the web component in a JSP file at the time the JSP file is translated, while in dynamic inclusion, the output of another component is included within the output of the JSP page when the JSP page is requested.

12.1 **STATIC INCLUSION**

In static inclusion, the contents of another file are included with the current JSP file at translation time to produce a single servlet. We use the JSP `include` directive to accomplish this. We have already seen the JSP syntax of the `include` directive in chapter 10, “The JSP technology model—the basics,” and the XML syntax in chapter 11, “The JSP technology model—advanced topics.” Here is a review of that syntax:

```
<%@ include file="relativeURL" %>
<jsp:directive.include file="relativeURL" />
```

The `file` attribute is the only attribute of the `include` directive, and it is mandatory. It refers to the file that contains the static text or code that will be inserted into the including JSP page. It can refer to any text-based file—HTML, JSP, XML, or even a simple `.txt` file—using a relative URL. A relative URL means that it cannot have a protocol, a hostname, or a port number. It can either be a path relative to the current JSP file—that is, it does not start with a `/`—or it can be a path relative to the document root of the web application—that is, it starts with a `/`. Figure 12.1 illustrates the way the `include` directive works.

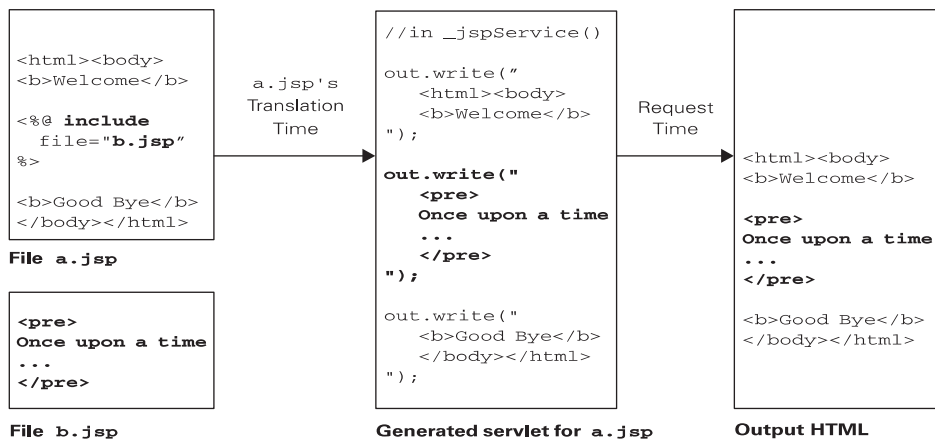


Figure 12.1 Static inclusion using the `include` directive

Figure 12.1 shows two JSP files: `a.jsp` and `b.jsp`. The `a.jsp` file contains an `include` directive that refers to the `b.jsp` file.

While generating the servlet code for `a.jsp`, the JSP engine includes all the elements of `b.jsp`. The resulting code, which is a combination of the including page and the included page, is then compiled as a single translation unit.

When a request is made for `a.jsp`, it will be processed by the servlet generated for `a.jsp`. However, because this servlet also contains the code from `b.jsp`, the resulting HTML page will contain the generated output from `a.jsp` as well as `b.jsp`.

12.1.1 Accessing variables from the included page

Since the code of the included JSP page becomes a part of the including JSP page, each page can access the variables and methods defined in the other page. They also share all of the implicit objects, as shown in listing 12.1.

Listing 12.1 `productsSearch.jsp`

```
<html><body>
  <%
    //Get the search criteria from the request.
    String criteria = request.getParameter("criteria");

    //Search the product database and get the product IDs.
    String productId[] = getMatchingProducts(criteria);
  %>

  The following products were found that match your criteria:<br>

  <!--
    Let productDescription.jsp generate the description
    for each of the products
  -->

  <%@ include file="productDescription.jsp" %>

  New Search:
  <!--
    FORM for another search
  -->
  <form>...</form>

</body></html>
```

In listing 12.1, the `productsSearch.jsp` file processes the search criteria entered by the user and retrieves the matching products from the database. It then includes the `productDescription.jsp` file to generate the product description.

The code for `productDescription.jsp` is shown in listing 12.2.

Listing 12.2 productDescription.jsp

```
<%
// The implicit variable request used here is
// actually that of the including page.
String sortBy = request.getParameter("sortBy");

// Use the productId array defined by productsSearch.jsp
// to sort and generate the description of the products
productId = sort(productId, sortBy);

for(int i=0; i<productId.length; i++)
{
    // Generate a tabular description
    // for the products.
}
%>
```

The `productDescription.jsp` file uses the implicit `request` object and the `productId` array defined in `productsSearch.jsp` to generate a tabular display of the products.

12.1.2 Implications of static inclusion

When an `include` directive includes a file, the following rules apply:

- No processing can be done at translation time, which means the `file` attribute value cannot be an expression. Therefore, the following use of the `include` directive is invalid:

```
<% String myURL ="copyright.html"; %>
<%@ include file="<%= myURL %>" %>
```

- Because request parameters are a property of the requests and do not make any sense at translation time, the `file` attribute value cannot pass any parameters to the included page. Thus, the value of the `file` attribute in the following example is invalid:

```
<%@ include file="other.jsp?abc=pqr" %>
```

- The included page may or may not be able to compile independently. If you look at listing 12.2, the `productDescription.jsp` file cannot be compiled, since it does not define the variable `productId`. In general, it is better to avoid such dependencies and use the implicit variable `pageContext` to share objects across statically included pages by using the `pageContext.setAttribute()` and `pageContext.getAttribute()` methods.

12.2 DYNAMIC INCLUSION

In dynamic inclusion, when the JSP page is requested, it sends a request to another object, and the output from that object is included in the requested JSP page. We use the standard JSP actions `<jsp:include>` and `<jsp:forward>` to implement dynamic inclusion. Their syntax is as follows:

```
<jsp:include page="relativeURL" flush="true" />
<jsp:forward page="relativeURL" />
```

The `page` attribute is mandatory. It must be a relative URL, and it can refer to any static or dynamic web component, including a servlet. It can also be a request-time expression, such as

```
<% String pageURL = "other.jsp"; %>
<jsp:include page="<%= pageURL %>" />
```

The `flush` attribute is only valid for `<jsp:include>` and not for `<jsp:forward>`. It is optional and specifies that if the output of the current JSP page is buffered, then the buffer should be flushed before passing the output stream to the included component. The default value for the `flush` attribute is `false`.

Functionally, the `<jsp:include>` and `<jsp:forward>` actions are equivalent to the `RequestDispatcher.include()` and `RequestDispatcher.forward()` methods that are used in servlets to include and forward the requests to other components.

12.2.1 Using `jsp:include`

The `<jsp:include>` action delegates the control of the request processing to the included component temporarily. Once the included component finishes its processing, the control is transferred back to the including page. Figure 12.2 illustrates this process.

Figure 12.2 shows two JSP files, `a.jsp` and `b.jsp`. The `a.jsp` file contains an `include` action that refers to `b.jsp` file. While generating the servlet code for `a.jsp`, the JSP engine includes a request-time call to `b.jsp`, and also generates the servlet code for `b.jsp` if it does not already exist. When a request is made for `a.jsp`, it will be received and processed by the servlet generated for `a.jsp`. However, since this servlet contains a call to the `b.jsp` file, the resulting HTML page will contain the output from the servlet generated for `b.jsp` as well as the servlet generated for `a.jsp`.

Because the semantics of `<jsp:include>` are the same as those of `RequestDispatcher.include()`, the following three constructs are equivalent:

Construct 1

```
<%
    RequestDispatcher rd =
        request.getRequestDispatcher("other.jsp");
```

```
rd.include(request, response);
%>
```

Construct 2

```
<%
pageContext.include("other.jsp");
%>
```

Construct 3

```
<jsp:include page="other.jsp" flush="true"/>
```

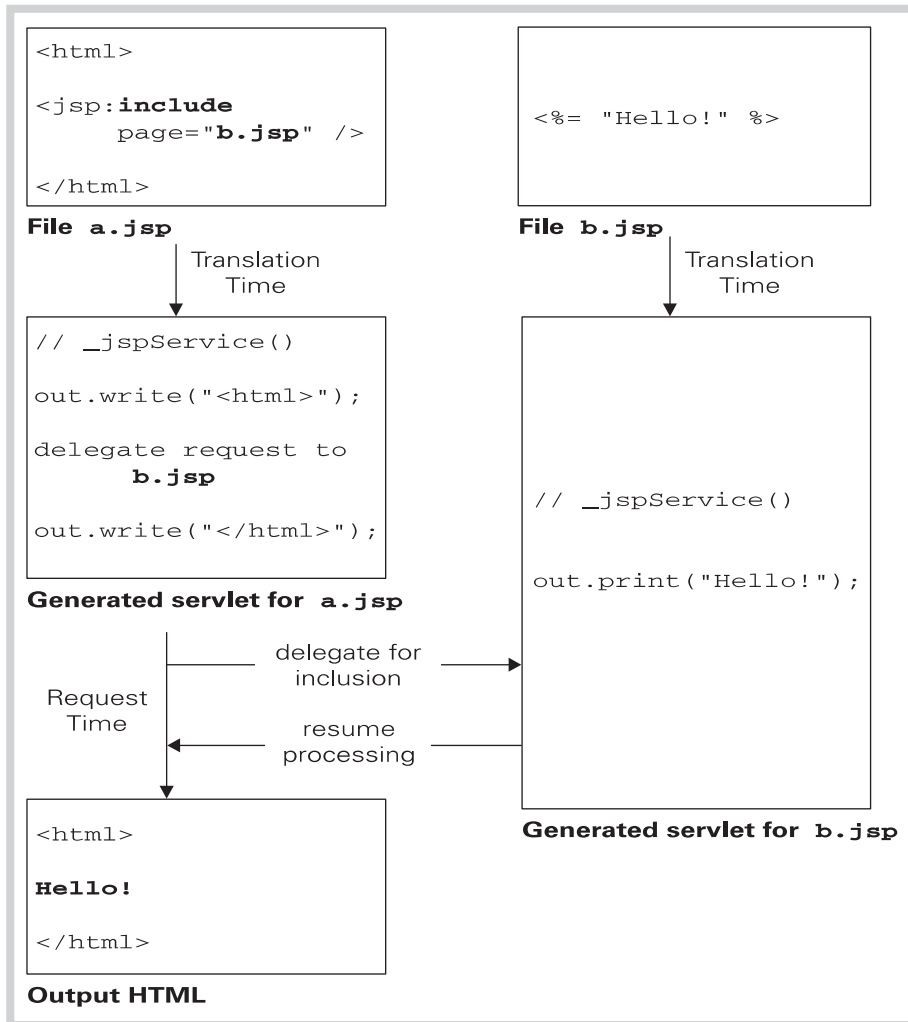


Figure 12.2 Dynamic inclusion using the include action

12.2.2 Using `jsp:forward`

The `<jsp:forward>` action delegates the request processing to the forwarded component. The forwarded component then sends the reply to the client. Figure 12.3 illustrates this process.

Figure 12.3 shows two JSP files, `a.jsp` and `b.jsp`. The `a.jsp` file contains a forward action that refers to `b.jsp`. While generating the servlet code for `a.jsp`, the JSP engine includes a request-time call to `b.jsp`, and also generates the servlet code for `b.jsp` if it does not already exist. At request time, `a.jsp` partially handles the request and delegates it to `b.jsp`, which then completes the request processing and sends the reply to the client.

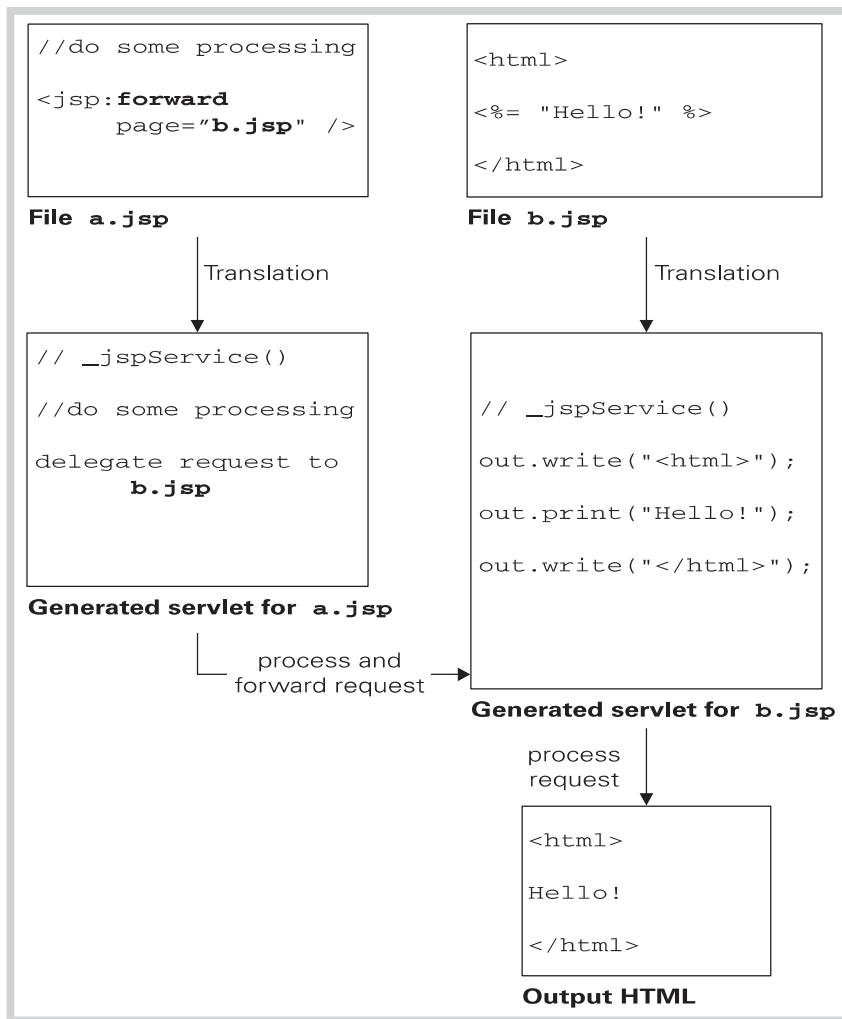


Figure 12.3 Dynamic inclusion using the forward action

Since the semantics of `<jsp:forward>` are the same as those of the `RequestDispatcher.forward()`, the following three constructs are equivalent:

Construct 1

```
<%
    RequestDispatcher rd =
        request.getRequestDispatcher("other.jsp");
    rd.forward(request, response);
%>
```

Construct 2

```
<%
    pageContext.forward("other.jsp");
%>
```

Construct 3

```
<jsp:forward page="other.jsp" />
```

In all three cases, if the output is buffered, it is first cleared, and then the request is forwarded to the other resource. However, if the output is not buffered and/or if the response is already committed by the forwarding resource, then a `java.lang.IllegalStateException` is raised when we attempt to forward the request.

12.2.3 Passing parameters to dynamically included components

We can pass parameters to the dynamically included components by using the `<jsp:param />` tags. The following examples illustrate the use of the `<jsp:param>` tag to pass two parameters to the included page:

```
<jsp:include page="somePage.jsp">
    <jsp:param name="name1" value="value1" />
    <jsp:param name="name2" value="value2" />
</jsp:include>
```

There can be any number of `<jsp:param>` elements nested within the `<jsp:include>` or `<jsp:forward>` element. The value of the `value` attribute can also be specified using a request-time attribute expression in the following way:

```
<jsp:include page="somePage.jsp">
    <jsp:param name="name1" value="<%= someExpr1 %>" />
    <jsp:param name="name2" value="<%= someExpr2 %>" />
</jsp:include>
```

In addition to parameters that are explicitly passed using the above methods, the included components have access to parameters that were originally present in the request to the including component. However, if the original parameter names are repeated in the explicitly passed parameters, the new values take precedence over the old values. For example, consider the two files `paramTest1.jsp` and `paramTest2.jsp` (shown in listings 12.3 and 12.4).

Listing 12.3 The file paramTest1.jsp

```
<html><body><pre>
In paramTest1:
First name is <%= request.getParameter("firstname") %>
Last name is <%= request.getParameter("lastname") %>

<jsp:include page="paramTest2.jsp" >
  <jsp:param name="firstname" value="mary" />
</jsp:include>

</pre></body></html>
```

Listing 12.4 The file paramTest2.jsp

```
In paramTest2:
First name is <%= request.getParameter("firstname") %>
Last name is <%= request.getParameter("lastname") %>

Looping through all the first names
<%
    String first[] = request.getParameterValues("firstname");
    for (int i=0; i<first.length; i++)
    {
        out.println(first[i]);
    }
%>
```

If you access the `paramTest1.jsp` file with the URL

```
http://localhost:8080/chapter12/paramTest1.jsp?firstname=john&lastname=smith
```

the output to the browser will be

```
In paramTest1:
First name is john
Last name is smith

In paramTest2:
First name is mary
Last name is smith

Looping through all the first names
mary
john
```

This is because when we call `paramTest1.jsp` (listing 12.3) using the URL given above, it receives the request parameters as `firstname=john&lastname=smith`. It prints out these values and then passes a new name-value pair, `firstname=mary`, to the included page `paramTest2.jsp` (listing 12.4) using the `<jsp:param>` element.

In the included page, the new value of the `firstname` parameter takes precedence over the original value and therefore receives the parameters as `firstname=`

mary&firstname=john&lastname=smith. Thus, a call to `request.getParameter("firstname")` will return "mary", while a call to `request.getParameterValues("firstname")` will return an array of `Strings` containing the values "mary" and "john". Since no new value for `lastname` was supplied, `paramTest2.jsp` uses the original value, "smith".

The name-value pairs passed in via the `<jsp:param>` tag exist within the request object and are available only for the included component. After the included component has finished processing, the engine removes these values from the request object. Thus, if the file `paramTest1.jsp` calls `request.getParameterValues("firstname")` after `paramTest2.jsp` returns, the call will return an array of `Strings` containing only one value: "john".

All of the above examples used `<jsp:include>`, but this discussion is equally applicable to `<jsp:forward>`.

12.2.4 Sharing objects with dynamically included components

The dynamically included pages execute separately, so they do not share the variables and methods defined by the including page. However, they process the same request and thus share all the objects present in the request scope, as shown in listing 12.5.

Listing 12.5 `productsSearch.jsp`

```
<html><body>

  <%
    //Get the search criteria from the request.
    String criteria = request.getParameter("criteria");

    //Search the product database and get the product IDs.
    String productId[] = getMatchingProducts(criteria);

    request.setAttribute("productIds", productId);
  %>

  The following products were found that match your criteria:<br>

  <!--
    Let productDescription.jsp generate the description
    for each of the products
  -->

  <jsp:include page="productDescription.jsp" />

  New Search:
  <!--
    FORM for another search
  -->
  <form>...</form>

</body></html>
```

Listing 12.5 produces the same results as listing 12.1, but it uses dynamic inclusion instead of static inclusion. The including file, `productSearch.jsp`, adds the `productId` object into the request scope by calling the `request.setAttribute()` method. The included file, `productDescription.jsp` (listing 12.6), then retrieves this object by calling the `request.getAttribute()` method.

Listing 12.6 `productDescription.jsp`

```
<%
//The implicit variable request used here is
//not the same as that of the including page.
//But the objects in the request scope are shared.

String sortBy = request.getParameter("sortBy");

String[] productIds = (String[])
    request.getAttribute("productIds");

//Use the productId array here
for(int i=0; i<productIds.length; i++)
{
    // Generate a tabular description
    // for the products.
}
%>
```

Here, the implicit variable `request` in `productsDescription.jsp` is accessing the same request scope objects as the implicit variable `request` in the `productSearch.jsp` file. The same mechanism can also be used with the `<jsp:forward>` action.

Note that in addition to `request`, we could use the implicit variables `session` and `application` to share objects with included and forwarded pages, but they are not meant for sharing request-dependent values. For example, if we use `application` instead of `request` in listings 12.5 and 12.6, then the product IDs generated for a search request from one client could affect the result of a search for another client because they share the value of the attribute `productIds` in the application scope.

Quizlet

Q: Explain whether each of the following are valid or invalid.

- a** `<jsp:include url="catalog.jsp" />`
- b** `<jsp:include page="http://myserver/catalog.jsp" />`
- c** `<jsp:include flush="true" />`
- d** `<jsp:forward flush="false" page="catalog.jsp" />`
- e** `<jsp:include page="/servlets/catalogServlet" />`
- f** `<%@ include page="catalog.jsp" %>`
- g** `<%@ include file="/servlets/catalogServlet" %>`
- h** `<%@ include file="catalog.jsp?category=gifts" %>`


```

i <% String fileURL = "catalog.jsp"; %>
j <%@ include file="<%= fileURL %>" %>

```

A: The following table explains why an option is valid or invalid.

Example	Construct	Valid/ invalid	Explanation
a	<code><jsp:include url="catalog.jsp" /></code>	Invalid	There is no attribute named <i>url</i> in <code>jsp:include</code> or <code>jsp:forward</code> . Use <i>page</i> .
b	<code><jsp:include page="http://myserver/catalog.jsp" /></code>	Invalid	The value of <i>page</i> has to be a relative URL. We cannot specify a protocol, hostname, or port number.
c	<code><jsp:include flush="true" /></code>	Invalid	The mandatory attribute <i>page</i> is missing.
d	<code><jsp:forward flush="false" page="catalog.jsp" /></code>	Invalid	The attribute <i>flush</i> is only for <code>jsp:include</code> , not for <code>jsp:forward</code> .
e	<code><jsp:include page="/servlets/catalogServlet" /></code>	Valid	The attribute <i>page</i> can point to a servlet.
f	<code><%@ include page="catalog.jsp" %></code>	Invalid	The include directive uses <i>file</i> , not <i>page</i> .
g	<code><%@ include file="/servlets/catalogServlet" %></code>	Invalid	When using the include directive, you can include XML, text, HTML, or even JSP files, but not a servlet.
h	<code><%@ include file="catalog.jsp?category=gifts" %></code>	Invalid	Query strings cannot be used with the include directive.
i	<code><% String fileURL = "catalog.jsp"; %> <%@ include file="<%=fileURL%>" %></code>	Invalid	The include directive cannot use a request-time attribute expression.

Q: What is the difference between the following constructs?

```

<% pageContext.include("other.jsp"); %>
<jsp:include page="other.jsp" />

```

A: They are functionally similar but with a minor difference. The `pageContext.include()` method always flushes the output of the current page before including the other components, while `<jsp:include>` flushes the output of the current page only if the value of `flush` is explicitly set to `true`, as in this example:

```

<jsp:include page="other.jsp" flush="true" />

```

Q: What is wrong with the following JSP document?

```
<jsp:root
    xmlns:jsp=http://java.sun.com/JSP/Page
    version="2.0">

    <jsp:scriptlet>
        String pageURL = "other.jsp";
    </jsp:scriptlet>

    <jsp:include page="<%= pageURL %>" />

</jsp:root>
```

A: The expression `<%= pageURL %>` is in the JSP syntax format. Recall from the previous chapter that the correct syntax for a request-time expression in XML syntax format is `%= expr %`. Thus, the `<jsp:include>` action in the above code must be written as

```
<jsp:include page="%=pageURL%" />
```

Q: The following code does not work. Identify the problem and rectify it. In the `main.jsp` file:

```
<html><body>

    <%
        Integer oneHundred = new Integer(100);
        Integer twoHundred = new Integer(200);
    %>

    <jsp:include page="display.jsp" >
        <jsp:param name="one" value="<%= oneHundred %>" />
        <jsp:param name="two" value="<%= twoHundred %>" />
    </jsp:include>

</body></html>
```

In the `display.jsp` file:

```
<%@ page import="java.lang.* " %>

<%
    Integer oneHundred = (Integer) request.getParameter("one");
    Integer twoHundred = (Integer) request.getParameter("two");
%>
```

A: We can use only `String` to pass and retrieve parameters using the `<jsp:param>` and `request.getParameter()` mechanisms. To pass any other type of object, we have to use `request.setAttribute()` in the including component and `request.getAttribute()` in the included component.

12.3 SUMMARY

JSP technology reuses web components by including the content or output of the components in a JSP page. There are two ways to do this: static inclusion and dynamic inclusion.

Static inclusion happens during the translation process, and it uses the `include` directive: `<%@ include %>`. The included file gets translated together with the including JSP file into a single servlet class. In this way, the pages are able to share all variables and methods.

Dynamic inclusion occurs at the time the including JSP page is requested, and it is accomplished by the standard actions, `<jsp:include>` and `<jsp:forward>`. In this chapter, we reviewed the techniques for passing parameters to the included components. Even though a dynamically included component does not share variables and methods with the including JSP page, they both process the same request and therefore share all the objects present in the request scope.

At this point, you should be able to answer questions based on the concepts of static and dynamic inclusion of web components in JSP pages.

In the next chapter, we will examine a new way of accessing and displaying variables in JSPs—the Expression Language.

12.4 REVIEW QUESTIONS

1. Which of the following JSP tags can be used to include the output of another JSP page into the output of the current page at request time? (Select one)

- a `<jsp:insert>`
- b `<jsp:include>`
- c `<jsp:directive.include>`
- d `<jsp:directive:include>`
- e `<%@ include %>`

2. Consider the contents of the following two JSP files:

File 1: `test1.jsp`

```
<html><body>
    <% String message = "Hello"; %>

    //1 Insert LOC here.

    The message is <%= message %>
</body></html>
```

File 2: `test2.jsp`

```
<% message = message + " world!"; %>
```

Which of the following lines can be inserted at `//1` in `test1.jsp` so that it prints "The message is Hello world!" when requested? (Select one)

- a `<%@ include page="test2.jsp" %>`
- b `<%@ include file="test2.jsp" %>`
- c `<jsp:include page="test2.jsp" />`
- d `<jsp:include file="test2.jsp" />`

3. Which of the following is a correct way to pass a parameter equivalent to the query string `user=mary` at request time to an included component? (Select one)

- a `<jsp:include page="other.jsp" >`
`<jsp:param paramName="user" paramValue="mary" />`
`</jsp:include>`
- b `<jsp:include page="other.jsp" >`
`<jsp:param name="mary" value="user" />`
`</jsp:include>`
- c `<jsp:include page="other.jsp" >`
`<jsp:param value="mary" name="user" />`
`</jsp:include>`
- d `<jsp:include page="other.jsp" >`
`<jsp:param param="user" value="mary"/>`
`</jsp:include>`
- e `<jsp:include page="other.jsp" >`
`<jsp:param user="mary" />`
`</jsp:include>`

4. Identify the JSP equivalent of the following code written in a servlet. (Select one)

```
RequestDispatcher rd = request.getRequestDispatcher("world.jsp");
rd.forward(request, response);
```

- a `<jsp:forward page="world.jsp"/>`
- b `<jsp:action.forward page="world.jsp"/>`
- c `<jsp:directive.forward page="world.jsp"/>`
- d `<%@ forward file="world.jsp"%>`
- e `<%@ forward page="world.jsp"%>`

5. Consider the contents of two JSP files:

File 1: `test1.jsp`

```
<html><body>
    <% pageContext.setAttribute("ninetyNine", new Integer(99)); %>

    //1

</body></html>
```

File 2: `test2.jsp`

```
The number is <%= pageContext.getAttribute("ninetyNine") %>
```

Which of the following, when placed at line //1 in the `test1.jsp` file, will allow the `test2.jsp` file to print the value of the attribute when `test1.jsp` is requested? (Select one)

- a `<jsp:include page="test2.jsp" />`
- b `<jsp:forward page="test2.jsp" />`
- c `<%@ include file="test2.jsp" %>`
- d None of the above because objects placed in `pageContext` have the `page` scope and cannot be shared with other components.

6. Consider the contents of two JSP files:

File 1: `this.jsp`

```
<html><body><pre>
  <jsp:include page="that.jsp" >
    <jsp:param name="color" value="red" />
    <jsp:param name="color" value="green" />
  </jsp:include>
</pre></body></html>
```

File 2: `that.jsp`

```
<%
  String colors[] = request.getParameterValues("color");
  for (int i=0; i<colors.length; i++)
  {
    out.print(colors[i] + " ");
  }
%>
```

What will be the output of accessing the `this.jsp` file via the following URL? (Select one)

```
http://localhost:8080/chapter12/this.jsp?color=blue
```

- a blue
- b red green
- c red green blue
- d blue red green
- e blue green red

7. Consider the contents of two JSP files:

File 1: `this.jsp`

```
<html><body>
  <%= request.getParameter("color") %>
  <jsp:include page="that.jsp" >
    <jsp:param name="color" value="red" />
```

```
</jsp:include>
<%= request.getParameter("color") %>
</body></html>
```

File 2: `that.jsp`

```
<%= request.getParameter("color") %>
```

What will be the output of accessing the `this.jsp` file via the following URL?
(Select one)

`http://localhost:8080/chapter12/this.jsp?color=blue`

- a** blue red blue
- b** blue red red
- c** blue blue red
- d** blue red null

8. Consider the contents of three JSP files:

File 1: `one.jsp`

```
<html><body><pre>
  <jsp:include page="two.jsp" >
    <jsp:param name="color" value="red" />
  </jsp:include>
</pre></body></html>
```

File 2: `two.jsp`

```
<jsp:include page="three.jsp" >
  <jsp:param name="color" value="green" />
</jsp:include>
```

File 3: `three.jsp`

```
<%= request.getParameter("color") %>
```

What will be the output of accessing the `one.jsp` file via the following URL?
(Select one)

`http://localhost:8080/chapter12/one.jsp?color=blue`

- a** red
- b** green
- c** blue
- d** The answer cannot be determined.



CHAPTER 13

Creating JSPs with the Expression Language (EL)

- 13.1 Understanding the Expression Language 237
- 13.2 Using EL operators 241
- 13.3 Incorporating functions with EL 244
- 13.4 Summary 249
- 13.5 Review questions 249

EXAM OBJECTIVES

- 7.1** Given a scenario, write EL code that accesses the following implicit variables:
 - `pageScope`, `requestScope`, `sessionScope`, and `applicationScope`
 - `param` and `paramValues`, `header` and `headerValues`
 - `cookie`, `initParam` and `pageContext`(Section 13.1)
- 7.2** Given a scenario, write EL code that uses the following operators:
 - property access (the `.` operator) and collection access (the `[]` operator)
 - arithmetic operators, relational operators, and logical operators(Section 13.2)
- 7.3** Given a scenario:
 - write EL code that uses a function
 - write code for an EL function
 - configure the EL function in a tag library descriptor(Section 13.3)

INTRODUCTION

So far, we've examined JSP declarations, expressions, and scriptlets, along with the tags needed to identify them. These are useful capabilities, but they present problems for complex JSPs. First, the clutter of different tags can make reading and debugging code a painful task. Second, these scripts still incorporate a great deal of business logic (Java) within presentation logic (HTML).

The developers of the Java Standard Tag Library (JSTL) responded to this with an Expression Language (EL) that depends less on Java and doesn't use tags at all. Sun added this language to the new JSP standard as an alternative to scripting. Since the EL objective is the only new one in the 310-081 exam, and since there are nine more questions than before, it's a safe bet that you'll see plenty of EL expressions on your SCWCD exam.

We'll begin by describing the role of EL and how it functions. Then, we'll explain the different operators available for manipulating variables. Finally, we'll describe functions in EL and explain how they relate to standard Java methods. We'll present code samples throughout and conclude the chapter with a set of review questions.

13.1 UNDERSTANDING THE EXPRESSION LANGUAGE

The Expression Language isn't a new set of XML tags or Java classes; it's a self-contained programming language complete with operators, syntax, and reserved words. Our job as JSP developers is to create EL expressions that can be added to the servlet's response. Let's start with the basics.

First, we'll compare EL expressions with material we've already covered: JSP's regular scripting expressions. We'll discuss the similarities between the two, and how they differ with regard to variables. Then, we'll explore the many implicit objects available in EL and how they enable us to access information outside the page.

13.1.1 EL expressions and JSP script expressions

An easy way to understand EL expressions is to compare them with traditional JSP script expressions. Both enable you to insert dynamic information within a static presentation. For example, if you want to describe a changing quantity in your page, you can use a JSP script expression like

```
The outside temperature is <%= temp %> degrees.
```

or, you can use EL

```
The outside temperature is ${temp} degrees.
```

Both statements produce the same output and the web container processes them in the same way. That is, once it receives a request, it evaluates the expression, converts it to a `String`, and inserts it into the response output stream.

These expressions also allow you to update attributes of standard or custom tags. To set the font of a given statement using a JSP expression, you can use

```
<FONT FACE=<%= font %>>This sentence uses the <%= font %> font.</FONT>
```

or, with EL,

```
<FONT FACE=${font}>This sentence uses the ${font} font.</FONT>.
```

But there are two important differences that you need to keep in mind. The first is obvious: all EL expressions begin with “`#{`” and end with “`}`”, while script expressions are enclosed within tags. The second difference is subtle but important. It deals with the variables that can be presented inside expressions.

With traditional scripting, it’s easy to declare a variable in a JSP. All you need are `<%!` and `%>` tags:

```
<%! int JSPvariable = 100; %>.
```

The problem is that this code uses Java, and a major goal of EL is to remove Java from JSPs. Therefore, *EL expressions can’t use variables declared in scripts*. For example, after the above declaration, the statement

```
The JSPvariable is <%= JSPvariable %>
```

will tell you that the variable is 100, while

```
The JSPvariable is ${JSPvariable}
```

will return an undefined value. So, since EL can’t declare variables by itself, we need other ways to create these placeholders. You can use tag libraries and JavaBean members, but the easiest variables to access are the implicit variables provided by the JSP itself.

13.1.2 Using implicit variables in EL expressions

When writing servlets, you can use a number of methods, such as `getServletContext()` and `getSession()`, to obtain information about the application. But these methods are inappropriate for JSPs, where Java programming should be kept to a minimum.

Instead, JSP designers retrieve information through implicit variables. This way, you can directly access the results of the servlet methods instead of coding them yourself. Implicit variables are very important to know—both for the exam and for practical JSP development. Table 13.1 lists them along with their descriptions.

Table 13.1 Implicit variables usable inside EL expressions

Name	Description
<code>pageContext</code>	Accesses the JSP’s regular implicit objects
<code>pageScope</code>	A Map containing the page scope attributes

continued on next page

Table 13.1 Implicit variables usable inside EL expressions (continued)

Name	Description
<code>requestScope</code>	A Map containing the request scope attributes
<code>sessionScope</code>	A Map containing the session scope attributes
<code>applicationScope</code>	A Map containing the application scope attributes
<code>param</code>	A Map containing a request parameter String
<code>paramValues</code>	A Map containing a request parameter String []
<code>header</code>	A Map containing a request header String
<code>headerValues</code>	A Map containing a request header String []
<code>cookie</code>	A Map matching Cookie fields to a single object

The `pageContext` variable gives you access to the implicit objects mentioned in chapter 12, “Reusable web components,” such as `application`, `session`, and `request`. To display the `bufferSize` of the page’s `JSPWriter`, use the expression

```
${pageContext.out.bufferSize}
```

or, to retrieve the request’s HTTP method, use this line of code:

```
${pageContext.request.method}
```

But because EL restrains you from invoking Java methods, you *can’t* use an expression like

```
${pageContext.request.getMethod() }
```

However, the following script expression will continue to work fine:

```
<%= request.getMethod() %>
```

The next four entries are easy to understand. They don’t give you direct access to the actual page, `ServletRequest`, `HttpSession`, or `ServletContext`, but instead return Maps that relate names of scope attributes to their values. For example, if you add a `totalPrice` attribute to the session to handle a user’s multiple purchases, you can display the value with the following expression:

```
${sessionScope.totalPrice}
```

Remember that you access `ServletContext` attributes through the `applicationScope` variable, not the `pageContext` variable.

The `param` and `paramValues` variables allow you to retrieve input values from the `ServletRequest`. The `param` variable is the result of invoking `getParameter (String name)` with the name of the parameter, and it is displayed by

```
${param.name}
```

using EL. Similarly, `paramValues` uses the `getParameterValues (String [] name)` method to return an array of values for a given name. Don’t worry; we’ll discuss arrays in EL shortly.

The `header` and `headerValues` variables work like `param` and `paramValues`, except that they retrieve values from the request header. The following line of code displays the `accept` field of an incoming header with EL:

```
${header.accept}
```

The final implicit variable, `cookie`, returns the result of the servlet's `getCookies()` method.

NOTE When trying to resolve a variable in an expression, such as `x` in `${x}`, the resolver will first examine the implicit variables. If it fails to find `x`, it will look through attributes of the page, request, session, and application scopes. If it still doesn't find `x`, the resolver will return `null`.

We can now finish our discussion of variables with a short example. Listing 13.1 presents a JSP that combines scripts, forms, implicit variables, and EL expressions.

Listing 13.1 EL_Variables.jsp

```
<html><body>
  <b>Expression Language Variables</b>
  <%! int x=4; %>
  <p>The script expression for x = <%= x %>.
  <p>The EL expression for x = ${x}.
  <form action="EL_Variables.jsp" method="GET">
    <p>What is x? <input type="text" size=2 name="num">
    <p><input type="submit">
  </form>
  <p>That's ${param.num == 4}!
</body></html>
```

The result is shown in figure 13.1. As you can see, the JSP receives an input parameter and displays it as an EL variable. But for real processing, we need to do something with these variables—transform them or combine them. For this, we need operators.

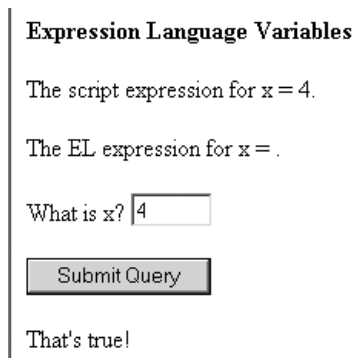


Figure 13.1
HTML output from
EL_Variables.jsp

Quizlet

- Q:** How would you display a request's URI with EL?
- A:** Insert `${pageContext.request.requestURI}` into your JSP. Although the `requestScope` variable allows you to access page attributes, you need the `pageContext` variable to access the request itself.
- Q:** If `result` is a valid parameter name, what will make `${paramValues.result}` a valid expression?
- A:** Since the `paramValues` variable returns a `Map` containing a `String[]`, you need to access an individual element of the array. Both `${paramValues.result[0]}` and `${paramValues.result["0"]}` will work, but `${paramValues.result.0}` won't. Remember that both `paramValues` and `headerValues` return `Maps` with `String[]`.

13.2 USING EL OPERATORS

Now that we've shown how EL uses variables, we can examine the operators that combine variables together. EL operators can be divided into four categories: property/collection access operators, arithmetic operators, relational operators, and logical operators. These should look familiar to those of you experienced in regular Java development. But there are a few odd details that you should keep in mind.

13.2.1 EL operators for property and collection access

Property access operators allow you to access an object's members, while collection access operators retrieve elements of a `Map`, `List`, or `Array`. These operators are particularly useful for dealing with implicit variables that contain or collect information. In EL, these operators are described by

- `a.b`—Returns the property of `a` associated with the identifier, `b`
- `a[b]`—Returns the value of `a` associated with the key or index, `b`

You've probably used these operators in your Java programming. They function similarly in EL, but EL treats them interchangeably if `b` is a `String`. This means the following expressions produce the same result:

```
${header["host"]}  
${header['host']}  
${header.host}
```

In this case, `header` is a `Map`, so the `header.get("host")` method is invoked to display the expression's result. Similarly, `headerValues.host` is an `Array` whose 0th element is displayed with one of the following expressions:

```
${headerValues.host["0"]}  
${headerValues.host['0']}  
${headerValues.host[0]}
```

But because `Arrays` are accessed by integer, the “`[]`” operator cannot be interchanged with the “`.`” operator. Both of the following expressions will cause a compilation error:

```
#{headerValues.host.0}
#{headerValues.host."0"}
```

So, EL doesn't quite treat property and collection access in exactly the same manner as traditional Java. However, its arithmetic operators are very standard and shouldn't cause any surprises.

13.2.2 EL arithmetic operators

EL allows you to use numerical values with data types similar to those provided by the `java.math` package. In particular, you can use `Integer` and `BigInteger` values for fixed-point numbers, and `Double` and `BigDecimal` values for floating-point numbers. The arithmetic operators available for these values are as follows:

- Addition: `+`
- Subtraction: `-`
- Multiplication: `*`
- Division: `div` and `/`
- Modulo division: `mod` and `%`

These operations invoke corresponding methods from `java.lang.Math` and behave as you'd expect. But it's important to remember the data type that results from an operation. For example, the result of an operation between a fixed-point and floating-point number is always a floating-point value. Similarly, an operation between a low-precision value and a high-precision value, such as an `Integer` added to a `BigInteger`, will always result in a high-precision value.

Here are some examples of EL's arithmetic operator usage. Note that “`e`” can be used in floating-point values to represent exponential notation:

```
#{2 * 3.14159} evaluates to 6.28318.
#{6.80 + -12} evaluates to -5.2.
#{24 mod 5} and #{24 % 5} evaluate to 4.
#{25 div 5} and #{25/5} evaluate to 5.0.
#{-30.0/5} evaluates to -6.0.
#{1.5e6/1000000} evaluates to 1.5.
#{1e6 * 1} evaluates to 1000000.0.
```

Along with numbers, `Strings` can be used in arithmetic operators as long as they can be converted into numbers:

```
#{ "16" * 4 } evaluates to 64.
#{ a div 4 } evaluates to 0.0.
#{ "a" div 4 } produces a compilation error.
```

We've seen how EL uses `Strings` and numbers, but there's one data type that we haven't encountered yet. When comparing variables with relational or logical operators, EL produces boolean results that can be `true` or `false`.

13.2.3 EL relational and logical operators

EL's relational operators are identical to those used in normal Java code. They include the following:

- Equality: `==` and `eq`
- Non-equality: `!=` and `ne`
- Less than: `<` and `lt`
- Greater than: `>` and `gt`
- Less than or equal: `<=` and `le`
- Greater than or equal: `>=` and `ge`

Relational expressions produce boolean variables that can be combined with the following logical operators:

- Logical conjunction: `&&` and `and`
- Logical disjunction: `||` and `or`
- Logical inversion: `!` and `not`

Because EL doesn't allow use of Java control statements like `if`, `for`, and `while`, there are only two ways we can use logical expressions. The first involves directly displaying the expression's boolean value. For example,

```
#{ 8.5 gt 4 } evaluates to true.
```

and

```
#{ (4 >= 9.2) || (1e2 <= 63) } evaluates to false.
```

EL's conditional operator gives us a second way. In this case, the displayed expression is based on a variable's boolean value. The syntax is `A ? B : C`. If `A` has a `true` value, then `B` will be the expression's result. If `A` is false, then `C` will be used. Here are a few examples of this operator:

```
#{ (5 * 5) == 25 ? 1 : 0 } evaluates to 1.
```

```
#{ (3 gt 2) && !(12 gt 6) ? "Right" : "Wrong" } evaluates to "Wrong".
```

```
#{ ("14" eq 14.0) && (14 le 16) ? "Yes" : "No" } evaluates to "Yes".
```

```
#{ (4.0 ne 4) || (100 <= 10) ? 1 : 0 } evaluates to 0.
```

EL provides many operators for comparing and combining variables, but it rejects any Java method within an expression. This makes it difficult to build custom operators such as traditional `log` or `pow` methods. But EL does allow you to invoke Java

methods in another file as long as you reference them within a tag library. These method references are called *functions*.

Quizlet

- Q:** What is the result of `#{ (10 le 10) && ! (24+1 lt 24) ? "Yes" : "No" }`?
- A:** The expression will evaluate to "Yes". The first expression is true since 10 equals 10, and the second expression is true since 25 isn't less than 24. Therefore, the conjunction of the two statements is true, and the first result, "Yes", will be displayed.

13.3 INCORPORATING FUNCTIONS WITH EL

Although EL functions can be complex to work with, they provide JSPs with complete separation of business and presentation logic. Instead of calling Java methods with scriptlets, these functions invoke methods by accessing their corresponding XML tags. As you'll see, this means that the page designer only needs the function names and the tag descriptor's URI in order to access the function within the JSP.

The process of inserting an EL function into a JSP involves creating or modifying four files:

- 1 *Method class (*.java)*—Contains the Java methods that you want to use in your JSP.
- 2 *Tag library descriptor (*.tld)*—Matches each Java method to an XML function name.
- 3 *Deployment descriptor (web.xml)*—Matches the TLD to a tag library URI. (Note: Changing this file is optional, but recommended.)
- 4 *JavaServer Page (*.jsp)*—Uses the tag library URI and function name to invoke the method.

The best way to show how these files interact is to create our own EL functions. Since these functions need to be matched to Java methods, we'll begin by defining these methods in a class.

13.3.1 Creating the static methods

The process of EL function development begins with creating the Java methods that will be indirectly invoked by the JSP. Our example, shown in listing 13.2, contains two simple methods, `upper()` and `length()`, that perform common `String` operations.

Listing 13.2 `myFunc.StrMethods.java`

```
package myFunc;

public class StrMethods {
    public static String upper( String x )
```

```

    {
        return x.toUpperCase();
    }

    public static int length( String x )
    {
        return x.length();
    }
}

```

Remember these important points when creating methods for EL:

- 1 The methods need to be declared `public` and `static`. The class needs to be `public`. This way, the servlet can access the class and its methods without creating a new object.
- 2 The method's arguments and return value must be valid within EL. Otherwise, the web container won't recognize the method signatures.
- 3 The class file needs to be saved in the `/WEB-INF/classes` directory.

Creating a method class is straightforward—it's regular Java. But getting the JSP to access this code requires more work. In this case, you need to speak the JSP's language—XML. The XML file that makes this possible is called the *tag library descriptor*.

13.3.2 Creating a tag library descriptor (TLD)

You'll have to wait for chapter 15, "Using custom tags," for a full treatment on tag libraries and their descriptor files, but we'll present a brief introduction here. The goal of the TLD is to map static methods into function names that you can use in your JSP. This is necessary because EL won't let you invoke Java methods.

In our example, the descriptor matches the `upper()` method to the function name, `upper`, and matches the `length()` method to the function name, `length`. This is shown in listing 13.3.

Listing 13.3 Functions.tld

```

<taglib 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
    web-jsptaglibrary_2_0.xsd"
  version="2.0">
  <tlib-version>1.0</tlib-version>
  <function>
    <name>upper</name>
    <function-class>myFunc.StrMethods</function-class>
    <function-signature>
      java.lang.String upper(java.lang.String)
    </function-signature>
  </function>

```



```

<function>
  <name>length</name>
  <function-class>myFunc.StrMethods</function-class>
  <function-signature>
    java.lang.int length(java.lang.String)
  </function-signature>
</function>
</taglib>

```

The first two tags in listing 13.3, `<taglib>` and `<tlib-version>`, identify how the rest of the file should be processed. We'll cover these in greater depth in chapter 15.

The tags contained within `<function>` and `</function>` do the main work of matching the functions to the Java methods. For each `<function>`, `</function>` pair, you need to provide three pieces of information:

- 1 The `<name>` and `</name>` tags contain the function names that will be used in the JSP. In the example, these names are `upper` and `length`.
- 2 The `<function-class>` and `</function-class>` tags contain the fully qualified name of the method class, which is `myFunc.StrMethods`.
- 3 The `<function-signature>` and `</function-signature>` tags identify a static method and the full data types of its arguments and return value. These data types must be resolvable inside an EL expression.

This `*.tld` file is usually placed inside the `/WEB-INF` directory or a subdirectory inside `/WEB-INF`. But since the container needs to know where the TLD is located, we will modify the application's deployment descriptor.

13.3.3 Modifying the deployment descriptor

One of the many functions of the deployment descriptor is to tell the web container where to find TLDs. The `web.xml` file does this by matching the TLD's actual location with a unique URI that can be used throughout the application. This centralized mapping means that we can move our TLDs from place to place, and only have to make a single alteration to `web.xml`.

You don't need to completely rewrite the deployment descriptor for this purpose, but listing 13.4 shows where the TLD locator tags should be added.

Listing 13.4 Modified `web.xml`

```

<web-app>
  ...
  <servlet>
    ...
  </servlet>
  ...
  <taglib>
    <taglib-uri>
      http://myFunc/Functions

```

```

        </taglib-uri>
        <taglib-location>
            /WEB-INF/myFunc/Functions.tld
        </taglib-location>
    </taglib>
    ...
</web-app>

```

The goal of this modification is to tell the web container that any servlet or JSP accessing the tag library with a URI of `http://myFunc/Functions` should be directed to `Functions.tld` located at `/WEB-INF/myFunc`. This information is contained within `<taglib>` and `</taglib>` tags, located directly beneath the descriptor's `<web-app>` and `</web-app>` tags. Two pieces of information need to be included:

- 1 The `<taglib-uri>` and `</taglib-uri>` tags contain a URI that will be used in servlets and JSPs to access the library. This URI can be absolute (`http://...`) or relative (`/...`).
- 2 The `<taglib-location>` and `</taglib-location>` tags contain the context path of the tag library descriptor. Since our TLD is in the `myFunc` directory inside `WEB-INF`, the location is given as `/WEB-INF/myFunc/Functions.tld`.

So far, you've created a Java class with two static methods. You've created a TLD that matches the methods to function names, and you've updated `web.xml` to match the TLD's location to a URI. With this accomplished, it's time to finish the process by building the JSP.

13.3.4 Accessing EL functions within a JSP

Once you've taken care of setting up the function names and TLD URI, calling the function inside your JSP is very simple. The process has two steps:

- 1 Use the `taglib` directive to access the TLD and assign a prefix to represent the tag library.
- 2 Create an EL expression using the TLD prefix and function name. Be sure to use proper argument types.

This process is shown in listing 13.5. First, we assign the prefix, `myString`, to represent our TLD, whose location is determined by its URI. Then, to call the functions, we use the expressions `${myString:upper() }` and `${myString:length() }` with the `String` parameter received from the request.

Listing 13.5 `Stringfun.jsp`

```

<%@ taglib prefix="myString"
uri="http://myFunc/Functions"%>
<html><body>

```

```

<b>Enter text:</b>
<form action="Stringfun.jsp" method="GET">
  <input type="text" name="x">
    <p><input type="submit">
</form>
<table border="1">
  <tr>
    <td>Uppercase:</td>
    <td>${myString:upper(param.x)}</td>
  </tr>
  <tr>
    <td>String length:</td>
    <td>${myString:length(param.x)}</td>
  </tr>
</table>
</body></html>

```

Figure 13.2 shows how a browser displays the JSP. As you can see, our custom function works just like a regular JSP action.

Before we finish our discussion on EL functions, we need to mention one last item. If you set the URI attribute of the `taglib` directive equal to the exact location of the TLD, you don't need to modify the deployment descriptor. This saves work for simple web pages, but creates complications for large enterprise applications. When you have many JSPs accessing multiple TLDs, it helps to have a central location for assigning URIs to file locations. Therefore, we've updated `web.xml` to assign a URI to the TLD's file location.

Quizlet

Q: What is wrong with the following `web.xml` assignment?

```

<taglib>
  <taglib-uri>
    http://myDir/File
  </taglib-uri>
  <taglib-directory>
    /WEB-INF/myDir/File.tlf
  </taglib-directory>
</taglib>

```

A: Two things. First, the TLD file should have a `*.tld` suffix. Second, the tags surrounding the TLD should be `<taglib-location>` and `</taglib-location>`, not `<taglib-directory>` and `</taglib-directory>`.

Figure 13.2 HTML output from `Stringfun.jsp`

13.4 SUMMARY

To understand the many rules and constraints of the Expression Language, it is important to remember its primary goal: to remove Java from JSP development. To meet this goal, it has its own operators, syntax, and methodology of function calling. It serves as a complete alternative to traditional declarations, expressions, and scriptlets.

If you are writing your own JSPs, you may want to continue using scripts. But, for the new SCWCD exam, EL's theory and implementation will be a very important topic. Therefore, we strongly recommend that you become very familiar with its operators and rules. Since the exam will focus on the tricks and trivia behind EL, we further recommend that you build and execute your own code to better see how its different aspects work.

EL operators and implicit variables won't present any problems for an experienced programmer. The Expression Language provides essentially the same constructs for property access, collection access, arithmetic, logic, and relational comparisons as C or Java. It is important to remember that property access and collection access are essentially the same thing in EL, and that EL numbers must be of the `Integer`, `BigInteger`, `Double`, or `BigDecimal` data types.

Function calling in EL, however, is a very new capability. You can invoke Java methods from a JSP, but only through custom tags in a TLD. The process begins with `public static` Java methods, which are mapped to function names in a tag library descriptor (TLD). Then, the JSP accesses this TLD using its URI or actual file location, and uses an assigned prefix to access the functions. The application's deployment descriptor can (and probably should) be used to centrally assign a URI to the TLD.

Although this chapter has touched on the process of tag library development, we've only scratched the surface. As we will see, there is much more that we can do with custom tags and TLDs that call EL functions.

13.5 REVIEW QUESTIONS

1. Consider the following code and select the correct statement from the options below.

```
<html><body>
    ${ (5 + 3 + a > 0) ? 1 : 2 }
</body></html>
```

- a It will print 1 because the statement is valid.
 - b It will print 2 because the statement is valid.
 - c It will throw an exception because `a` is undefined.
 - d It will throw an exception because the expression's syntax is invalid.
2. Which statement best expresses the purpose of a tag library descriptor (TLD) *in an EL function*?
 - a It contains the Java code that will be compiled.
 - b It invokes the Java method as part of the JSP.

- c** It matches the tag library with a URI.
 - d** It matches function names to tags that can be used in the JSP.
- 3. Which of the following variables is not available for use in EL expressions?
 - a** param
 - b** cookie
 - c** header
 - d** pageContext
 - e** contextScope
- 4. Which tags tell the web container where to find your TLD file in your filesystem?
 - a** <taglib-directory></taglib-directory>
 - b** <taglib-uri></taglib-uri>
 - c** <taglib-location></taglib-location>
 - d** <tld-directory></tld-directory>
 - e** <taglib-name></taglib-name>
- 5. Which two of the following expressions won't return the header's accept field?
 - a** `${header.accept}`
 - b** `${header[accept]}`
 - c** `${header['accept']}`
 - d** `${header["accept"]}`
 - e** `${header.'accept'}`
- 6. When writing a TLD, which tags would you use to surround `fnName(int num)`, a Java method declared in a separate class?
 - a** <function-signature></function-signature>
 - b** <function-name></function-name>
 - c** <method-class></method-class>
 - d** <method-signature></method-signature>
 - e** <function-class></function-class>
- 7. Which of the following method signatures is usable in EL functions?
 - a** `public static expFun(void)`
 - b** `expFun(void)`
 - c** `private expFun(void)`
 - d** `public expFun(void)`
 - e** `public native expFun(void)`



CHAPTER 14

Using JavaBeans

- | | |
|---|--|
| 14.1 JavaBeans: a brief overview 252 | 14.5 More about properties in
JavaBeans 276 |
| 14.2 Using JavaBeans with JSP
actions 258 | 14.6 Summary 280 |
| 14.3 JavaBeans in servlets 271 | 14.7 Review questions 281 |
| 14.4 Accessing JavaBeans from
scripting elements 274 | |

EXAM OBJECTIVES

8.1 Given a design goal, create a code snippet using the following standard actions:

- `jsp:useBean` (with attributes: 'id', 'scope', 'type', and 'class');
- `jsp:getProperty`; and
- `jsp:setProperty` (with all attribute contributions)

(Sections 14.2 – 14.4)

INTRODUCTION

JavaBeans are independent software components that we can use to assemble other components and applications. JSP technology uses standard tags to access JavaBeans components, which allow us to encapsulate code, perform complex operations, and leverage existing components to save time. In this chapter, we will give you a brief overview of JavaBeans from the JSP perspective, and show you how they are used in JSP pages.

14.1 JAVABEANS: A BRIEF OVERVIEW

The JavaBeans component model architecture is both a specification and a framework of APIs that supports a set of features that includes component introspection, properties, events, and persistence. Since it is platform independent, it enables us to write portable and reusable components.

Components developed according to this specification are called *beans*. From a developer's perspective, a bean is a Java class object that encapsulates data in the form of instance variables. These variables are referred to as *properties* of the bean. The class then provides a set of methods for accessing and mutating its properties. The actual strength of a bean as a reusable component lies in its ability to allow programmatic introspection of its properties. This ability facilitates automated support for bean customization using software programs called *bean containers*.

14.1.1 JavaBeans from the JSP perspective

In the JSP technology, the JSP engine acts as a bean container. Any class that follows these two conventions can be used as a JavaBean in JSP pages:

- The class must have a public constructor with no arguments. This allows the class to be instantiated as needed by the JSP engine.
- For every property, the class must have two publicly accessible methods, referred to as the *getter* and the *setter*, that allow the JSP engine to access or mutate the bean's properties.

The name of the method that accesses the property should be `getXXX()` and the name of the method that mutates the property should be `setXXX()`, where `XXX` is the name of the property with the first character capitalized. Here are the signatures of the methods:

```
public property-type getXXX();
public void setXXX(property-type);
```

In the following getter and setter methods, the name of the property is `color` and its data type is `String`:

```
public String getColor();
public void setColor(String);
```

Let's look at a simple example of a Java class that can be used as a JavaBean in a JSP page. In listing 14.1, the class `AddressBean` encapsulates the address information in four private attributes and provides access to them via the corresponding setter and getter methods.

Listing 14.1 A simple JavaBean class named `AddressBean`

```
public class chapter14.AddressBean
{
    //properties
    private String street;
    private String city;
```

```

private String state;
private String zip;

//setters
public void setStreet(String street){ this.street = street; }
public void setCity(String city)    { this.city   = city;   }
public void setState(String state)  { this.state  = state;  }
public void setZip(String zip)      { this.zip    = zip;    }

//getters
public String getStreet(){ return this.street; }
public String getCity()  { return this.city;   }
public String getState() { return this.state;  }
public String getZip()   { return this.zip;    }
}

```

Note that the name of our class `AddressBean` ends with the word *Bean*. Although this is not a requirement, many developers like to follow this convention (`UserBean`, `AccountBean`, etc.) to differentiate between JavaBean classes and ordinary classes, thus making their intent clear to co-developers.

The rules for placing the bean classes are the same as for any other class, such as servlets, utility classes, or third-party tools. They must be present in the classpath of the web application—which means we can keep them directly in the `/WEB-INF/classes` directory, or in a JAR file under the `/WEB-INF/lib` directory. Then, to use these classes within the JSP pages, we have to import them via the `import` attribute of the `page` directive.

14.1.2 The JavaBean advantage

Let's look at an example of using the `AddressBean` class in a JSP page. In this example, we want to capture the address information of visitors to our web site and maintain it during the lifetime of a session. Listing 14.2 shows an HTML page code with an input form that will collect this information.

Listing 14.2 addressForm.html

```

<html>
<body>
Please give your address:<br>
  <form action="address.jsp">
    Street: <input type="text" name="street"><br>
    City:   <input type="text" name="city"><br>
    State:  <input type="text" name="state"><br>
    Zip:    <input type="text" name="zip"><br>
           <input type="submit"><br>
  </form>
</body>
</html>

```


When the user fills out the form and submits the page, we need to perform the following tasks on the server:

- 1 Check if an `AddressBean` object already exists in the session.
- 2 If not, create a new `AddressBean` object and add it to the session.
- 3 Call `request.getParameter()` for all the HTML FORM fields.
- 4 Set the respective values into the `AddressBean` object.

If there were no support from the JSP engine, we would have to code the above steps in a scriptlet, as shown here:

```
<%@ page import="chapter14.AddressBean" %>
<%
    AddressBean address = null;

    synchronized(session)
    {
        //Get an existing instance
        address = (AddressBean) session.getAttribute("address");

        //Create a new instance if required
        if (address==null)
        {
            address = new AddressBean();
            session.setAttribute("address", address);
        }

        //Get the parameters and fill up the address object
        address.setStreet(request.getParameter("street"));
        address.setCity(request.getParameter("city"));
        address.setState(request.getParameter("state"));
        address.setZip(request.getParameter("zip"));
    }
%>
```

However, the JSP specification defines standard actions that provide a convenient means of handling HTML FORM input and sharing information across the JSP pages using JavaBeans. The scriptlet code shown above can be replaced with the following lines using the `JavaBean` and the standard JSP actions:

```
<%@ page import="chapter14.AddressBean" %>
<jsp:useBean id="address" class="AddressBean" scope="session" />
<jsp:setProperty name="address" property="*" />
```

Shorter code is not the only incentive to use JavaBeans in JSP pages; JavaBeans also help to increase code reusability. Suppose that after setting all the fields of the `AddressBean` to the values retrieved from the request parameters we want to persist this information into a database. We could write a scriptlet and include the logic of opening the database connection and saving the bean's properties in that scriptlet.

But what if this functionality is required by more than one JSP page? We have to repeat the same scriptlet code in all the JSP pages. And then, what if the logic to access the database changes? We have to modify all the affected pages to adapt to the new database logic. However, if we build that logic into a method in the `AddressBean` class itself, then all the pages can use that method. Furthermore, if the database access logic changes, only the `AddressBean` class and its method change; the JSP pages remain unaffected.

Another advantage of using beans is that they are Java programming language objects, which means we can fully utilize the object-oriented features provided by the language. Let's suppose the application needs to maintain two different types of addresses: one for businesses and one for residences. We can have two separate beans, `BusinessAddressBean` and `ResidentialAddressBean`, both derived from a common base class, `AddressBean`. The base class can implement all the logic common to both beans, while the derived classes can handle the logic that is specific to each bean independently.

14.1.3 Serialized JavaBeans

Usually, in an enterprise application we use some form of a database to store and retrieve persistent data. For example, we can have a relational database with a table for storing address information captured in our `AddressBean` when a visitor first goes to the site. The same address information can be retrieved from the database when we want to display the address at a later date. Storing JavaBean properties in a database is one way of persisting JavaBeans. The JavaBeans specification also allows us to persist JavaBeans in the file system as serialized objects.

A *serialized bean* is a bean instance that is converted into a data stream and stored in a file so that its attributes and values are saved permanently and can be retrieved later as required. The process of serialization is achieved using the standard *Object Serialization* mechanism of Java. First, we make our bean class capable of being serialized by implementing the `java.io.Serializable` interface.¹ Then we can serialize the individual bean instances by using the `java.io.ObjectOutputStream` class.

Serialized beans are considered to be resources and have the following requirements:

- The file that stores a serialized bean must have the extension `.ser`. For example, we can serialize instances of the `AddressBean` class in files and name the files after the person to which the address belongs, such as `John.ser` or `Mary.ser`.
- The file that stores the bean must be present in the classpath of the web application. Since the `/WEB-INF/classes` directory is always present in the classpath of a web application, we can create and save all the beans in either

¹ It can also be done by implementing the `java.io.Externalizable` interface. Please see the JDK API for more details.

the `/WEB-INF/classes` directory or a subdirectory of the `/WEB-INF/classes` directory. For example:

```
/WEB-INF/classes/John.ser  
/WEB-INF/classes/businessData/visitorAddresses/Mary.ser
```

- The combination of the path and the filename that stores the bean is treated as the *name* of the bean, similar to the way classes and packages are treated.

Thus, in the examples shown above, the names of the two serialized beans are `John` and `businessData.visitorAddresses.Mary`. This is because `/WEB-INF/classes/` is in the classpath of the web application. If we also add the `businessData` directory to the classpath so that any file under `<WEB-INF>/classes/businessData` is available as a resource, we can refer to Mary's bean as `visitorAddresses.Mary`.

Once saved as serialized beans, these objects can be loaded in any Java program using the `java.beans.Beans.instantiate()`² method.

In the following example, we will assume that we have a directory structure called `/WEB-INF/classes/businessData/visitorAddresses/`, and we will create the serialized beans in that directory.

First, let's make our `AddressBean` class serializable as follows:

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

Listing 14.3, `beanSaver.jsp`, accepts the user's name and address information in the request parameter, creates an instance of `AddressBean`, and serializes it to a file that is given the name of the user.

Listing 14.3 beanSaver.jsp

```
<%@ page import="chapter14.AddressBean, java.io.* " %>  
  
<%  
    String message="";  
  
    try  
    {  
  
        //Create an instance. Set the properties  
        AddressBean address = new AddressBean();  
        address.setCity(request.getParameter("city"));  
        address.setState(request.getParameter("state"));  
  
        //Get the user's name to build the file path
```

² See the JDK documentation of the package `java.beans` for more details

```

String name = request.getParameter("name");

String appRelativePath =
    "/WEB-INF/classes/businessData/visitorAddresses/"
    + name
    + ".ser";

String realPath = application.getRealPath(appRelativePath);

//Serialize the object into the file
FileOutputStream fos = new FileOutputStream(realPath);
ObjectOutputStream oos = new ObjectOutputStream(fos);
oos.writeObject(address);
oos.close();

message = "Successfully saved the bean as " + realPath;
}
catch(Exception e){
    message = "Error: Could not save the bean";
}
}
%>
<html><body>
    <h3><%= message %></h3>
</body></html>

```

In this example, we first import the `AddressBean` class and the `java.io` package via the `import` attribute of the `page` directive. Then, we create an instance of the `AddressBean` class and set its `city` and `state` properties as specified in the `request` parameter. Next, using the value of the `request` parameter `name`, we build the path to a filename. Finally, using the `FileOutputStream` and `ObjectOutputStream` classes of the `java.io` package, we serialize the bean into the file.

We can access the above page via the following URL:

```

http://localhost:8080/chapter14/
    beanSaver.jsp?name=John&city=Topeka&state=Kansas

```

If all goes well, it will create a file named `John.ser` under the directory `/WEB-INF/classes/businessData/visitorAddresses` and will print the message with the filename on the browser window.

Thus, we can create as many serialized beans as we want with different sets of property values as long as the names of the bean files are different. We can then use these serialized beans in other components of the application.

In the following sections, we will explain the standard JSP actions that help us use JavaBeans in JSP pages effectively.

14.2 USING JAVABEANS WITH JSP ACTIONS

Table 14.1 summarizes the three standard actions for using JavaBeans in a JSP page.

Table 14.1 Standard JSP actions for using JavaBeans

Action	Description
<code><jsp:useBean></code>	Declares the use of a JavaBean instance in a JSP page
<code><jsp:setProperty></code>	Sets new values to the bean's properties
<code><jsp:getProperty></code>	Gets the current value of the bean's properties

In the sections that follow, we will describe each of these actions in details.

14.2.1 Declaring JavaBeans using `<jsp:useBean>`

The `jsp:useBean` action declares a variable in the JSP page and associates an instance of a JavaBean with it. The association is a two-step process. First, the action tries to find an existing instance of the bean. If an instance is not found, a new instance is created and associated with the declared variable. We can customize the behavior of this action using the five attributes shown in table 14.2.

Table 14.2 `<jsp:useBean>` attributes

Attribute Name	Description	Examples
<code>id</code>	The name by which the bean is identified in the JSP page.	<code>id="address"</code>
<code>scope</code>	The scope of the bean's instance: <code>page</code> , <code>request</code> , <code>session</code> , or <code>application</code> . The default value is <code>page</code> .	<code>scope="session"</code>
<code>class</code>	The Java class of the bean.	<code>class="BusinessAddressBean"</code>
<code>type</code>	Specifies the type of the variable to be used to refer to the bean.	<code>type="AddressBean"</code>
<code>beanName</code>	The name of a serialized bean if we are loading from a file or the name of a class if we are creating a new instance.	<code>beanName="businessData.John"</code> <code>beanName="AddressBean"</code>

Of the five attributes in table 14.2:

- The `id` attribute is mandatory.
- The `scope` attribute is optional.
- The three attributes `class`, `type`, and `beanName` can only be used in one of the following four combinations, and at least one of these attributes or combinations of attributes must be present in a `useBean` action:
 - `class`
 - `type`

- `class` and `type`
- `beanName` and `type`

To make it easier to understand and remember the usage of these attributes, let's first examine their meaning in the following sections. Then, we will look at some examples that will demonstrate how we can use the different combinations.

The id attribute

The `id` attribute uniquely identifies a particular instance of a bean. It is mandatory because its value is required by the other JSP actions, `<jsp:setProperty>` and `<jsp:getProperty>`, to identify the particular instance of the bean. In the generated Java servlet, the value of `id` is treated as a Java language variable; therefore, we can use this variable name in expressions and scriptlets in the JSP page. Note that since the value of the `id` attribute uniquely identifies a particular instance of a bean, we cannot use the same value for the `id` attribute in more than one `<jsp:useBean>` action within a single translation unit.

The scope attribute

The `scope` attribute specifies the scope in which the bean instance resides. Like implicit objects, the existence and accessibility of JavaBeans from JSP pages are determined by the four JSP scopes: *page*, *request*, *session*, and *application*. This attribute is optional, and if not specified, the *page* scope is used by default.

We cannot use the *session* scope for a bean in a JSP page if we set the value of the page directive attribute `session` to `false`.

The class attribute

The `class` attribute specifies the Java class of the bean instance. If the `<jsp:useBean>` action cannot find an existing bean in the specified scope, it creates a new instance of the bean's class as specified by the value of the `class` attribute using the class's publicly defined no-argument constructor. Therefore, the class specified by the `class` attribute must be a `public` non-abstract class and must have a `public` no-argument constructor. If the class is part of a package, then the fully qualified class name must be specified as `mypackage.MyClass`.

The type attribute

The `type` attribute specifies the type of the variable declared by the `id` attribute. Since the declared variable refers to the actual bean instance at request time, its type must be the same as the bean's class, or a superclass of the bean's class, or an interface implemented by the bean's class. Again, if the class or the interface is part of a package, then the fully qualified name must be specified as `mypackage.MyClass`.

The beanName attribute

The `beanName` attribute specifies the name of a bean as expected by the `instantiate()` method of the `java.beans.Beans` class. For this reason, `beanName` can refer either to a serialized bean or to the name of a class whose instance is to be created.

If the `beanName` attribute refers to a serialized bean, then the bean is loaded from the file that holds the bean. For example, if the attribute is specified as `beanName="businessData.visitorAddresses.John"`, the bean is loaded from the file `businessData/visitorAddresses/John.ser`. Note that we do not use the extension `.ser` in the value for the `beanName` attribute.

If the `beanName` attribute refers to a class, the class is loaded into memory, an instance of the class is created, and the instance is used as a bean. For example, if the attribute is specified as `beanName="chapter14.AddressBean"`, then the class is loaded from the file `chapter14/AddressBean.class`. Note that we do not use the extension `.class` in the value for the `beanName` attribute.

Both `class` and `beanName` can be used to create new instances from the specified class. However, the advantage of using `beanName` over `class` is that the value of the `beanName` attribute can also be specified as a request-time attribute expression, so it can be decided at request time and need not be specified at translation time.

Using the combinations of attributes

Now that we have introduced the five attributes of the `<jsp:useBean>` action, let's take a look at the way they are used in JSP pages.

A simple useBean declaration

This action uses three attributes—`id`, `class`, and `scope`—to declare the use of a `JavaBean`:

```
<jsp:useBean id="address" class="chapter14.AddressBean" scope="session" />
```

This informs the JSP engine that the variable that refers to the bean should be named `address` and that the bean should be an instance of the `AddressBean` class. The `scope` attribute specifies the bean's scope as `session`.

At request time, if an object named `address` is already present in the session scope, it is assigned to the variable `address`. Otherwise, a new object of class `AddressBean` is created, which is then assigned to the variable and added to the session. It is equivalent to the following code:

```
chapter14.AddressBean address = (chapter14.AddressBean)
session.getAttribute("address");
if (address == null)
{
    address = new chapter14.AddressBean ();
    session.setAttribute("address", address);
}
```

The default scope

The following declaration uses only two attributes—`id` and `class`:

```
<jsp:useBean id="address" class="chapter14.AddressBean" />
```

This is similar to the previous example, except that we have not specified the `scope` attribute. In this case, the page scope is used by default. Thus, this bean is available only in the JSP page in which it is defined and only for the request for which it is created. It is equivalent to the following code:

```
chapter14.AddressBean address = (chapter14.AddressBean)
    pageContext.getAttribute("address");

if (address == null)
{
    address = new chapter14.AddressBean();
    pageContext.setAttribute("address", address);
}
```

The typecast problem

Suppose we have two classes, `BusinessAddressBean` and `ResidentialAddressBean`, both derived from a common base class, `AddressBean`. We have two JSP pages, both declaring a bean with the same `id` value, but each with a different `class` value, as shown here:

In `residential.jsp`:

```
<jsp:useBean id="address"
    scope="session"
    class="chapter14.ResidentialAddressBean" />
```

In `business.jsp`:

```
<jsp:useBean id="address"
    scope="session"
    class="chapter14.BusinessAddressBean" />
```

If the page `residential.jsp` is accessed first, its `<jsp:useBean>` action will add an object of the `ResidentialAddressBean` class into the session scope with the name `address`. Now, if the `business.jsp` page is accessed within the same session, the `<jsp:useBean>` action of `business.jsp` will locate the `address` object in the session scope and will try to cast it to the `BusinessAddressBean` class. Since the two classes do not have a class-subclass relationship, it will raise a `java.lang.ClassCastException`. Similarly, if the `business.jsp` page is accessed first, then the `<jsp:useBean>` action in the `residential.jsp` page will raise a `java.lang.ClassCastException`.

Using the class and type attributes

The following declaration uses the `class` attribute as well as the `type` attribute:

```
<jsp:useBean id="address"
             type="AddressBean"
             class="chapter14.BusinessAddressBean"
             scope="session" />
```

In this action, the variable named `address` that refers to the bean is declared of type `AddressBean`. As we mentioned earlier, before creating an instance using the `class` attribute, the engine looks for an existing bean by the name `address` in the session scope. If an existing bean is found, then it is assigned to the `address` variable. Note that, in this case, since the `address` variable is declared of type `AddressBean`, the actual class of the existing instance may be `AddressBean` or any subclass of `AddressBean`. Therefore, the actual class of the existing instance need not be of type `BusinessAddressBean`.

However, if an existing instance is not found and a new instance has to be created, then the value of the `class` attribute specifies that the actual instance of the bean that is created must be of the `BusinessAddressBean` class. This `useBean` declaration is equivalent to the following code:

```
AddressBean address = (AddressBean)
                      session.getAttribute("address");

if (address == null)
{
    address = new chapter14.BusinessAddressBean();
    session.setAttribute("address", address);
}
```

If we do not explicitly specify the `type` attribute, it is considered to be the same as `class`. The following two tags are equivalent:

```
<jsp:useBean id="address"
             class="chapter14.BusinessAddressBean"
             scope="session" />

<jsp:useBean id="address"
             type="BusinessAddressBean"
             class="chapter14.BusinessAddressBean"
             scope="session" />
```

Using serialized beans

In the following `useBean` declaration, the `beanName` attribute specifies the use of a serialized bean, `businessData.visitorAddresses.John`:

```
<jsp:useBean id="address"
             type="AddressBean"
             beanName="businessData.visitorAddresses.John"
             scope="session" />
```

In this case, the action first tries to locate an existing instance of the bean in the session scope. If the bean is not found, then the action creates an instance and initializes it with the serialized data present in the `businessData/visitorAddresses/John.ser` file. This method of locating and creating a bean is equivalent to the following code:

```
AddressBean address = (AddressBean)
    session.getAttribute("address");

if (address == null)
{
    ClassLoader classLoader = this.getClass().getClassLoader();

    address = (AddressBean)
        java.beans.Beans.instantiate(
            classLoader,
            "businessData.visitorAddresses.John");

    session.setAttribute("address", address);
}
```

In the following `useBean` declaration, the `beanName` attribute specifies the name of a class instead of a serialized bean:

```
<jsp:useBean id="address"
    type="AddressBean"
    beanName="AddressBean"
    scope="session" />
```

Here, the action first tries to locate an existing instance of the bean in the specified scope. If the bean is not found, the action creates an instance of the class specified by the `beanName` attribute. This action is equivalent to the following:

```
java.beans.Beans.instantiate(classLoader, "AddressBean");
```

Note that we do not specify the `class` attribute with the `beanName` attribute, because the class of the bean is determined either by the bean's serialized data itself or by the value of the `beanName` attribute. But the `type` attribute is required in order to determine the type of the declared variable. Thus, the value of the `type` attribute must be same as the class of the bean, a superclass of the bean, or an interface implemented by the bean.

Using a request-time attribute expression with the beanName attribute

Since the value of the `beanName` attribute can also be specified as a request-time attribute expression, it can be useful in deciding the resource to be used as a bean at request time. Consider the following example:

```
<%@ page import="chapter14.AddressBean, java.io.*" %>
<%
    String theBeanName = null;

    String name = request.getParameter("name");
```

```

        if (name!=null && !name.equals(""))
        {
            theBeanName = "businessData.visitorAddresses. " + name;
        }
        else
        {
            //Name not specified.

            if ("Business".equals(request.getParameter("newType")))
            {
                theBeanName = "BusinessAddressBean";
            }
            else
            {
                theBeanName = "ResidentialAddressBean";
            }
        }
    }
%>

<jsp:useBean id="address"
            type="AddressBean"
            beanName="<%= theBeanName %>" />

```

If the value of the request parameter name is John (or Mary), then the action will try to locate a serialized file named `businessData/visitorAddresses/John.ser` (or `businessData/visitorAddresses/Mary.ser`) in the classpath.

If the name request parameter is not specified, the `useBean` action will try to create an instance of the `BusinessAddressBean` or `ResidentialAddressBean` class, depending on the `newType` request parameter.

Notice that the `type` attribute in the action specifies the type as `AddressBean`. This ensures that the code will work regardless of the actual type of the instance at runtime. For example, the serialized bean of John may be of the type `BusinessAddressBean`, while that of Mary may be of the type `ResidentialAddressBean`. This is true even for the new instances created based on the `newType` parameter. Also, the `import` directive imports only one class, `chapter14.AddressBean`. Since the two derived classes are used only within `String` literals, we do not have to import them explicitly.

In either of the two cases, if the specified resource, serialized bean, or class is not found, then a `java.lang.InstantiationException` is thrown.

Using the type attribute

The following action uses the `type` attribute without the `class` or `beanName` attribute. This is useful if we want to locate an existing bean object but do not want to create a new instance even if an existing instance is not available:

```

<jsp:useBean id="address" type="AddressBean" scope="session" />

```

If the located object is not of type `AddressBean`, and if it is not a subtype of `AddressBean`, a `ClassCastException` is thrown. On the other hand, if the

object could not be located in the specified scope at all, no new object is created and a `java.lang.InstantiationException` is thrown.

Initializing bean properties

A limitation of using JavaBeans is that they are instantiated by the JSP engine using a no-argument constructor. Because of this limitation, we cannot initialize the beans by passing parameters to constructors. To overcome this, the JSP specification allows us to provide a body for the `<jsp:useBean>` tag, as shown by this example:

```
<jsp:useBean id="address" scope="session" class="AddressBean" >
  <%
    address.setStreet("123 Main St. ");
  %>
</jsp:useBean>
```

The previous code conveys two things to the JSP engine:

- 1 If the bean named `address` is already present in the session scope, then the JSP engine should skip the body of the `<jsp:useBean>` tag and use the bean object as it is. In this case, the nested scriptlet code is not executed.
- 2 If the bean named `address` is not already present in the session scope, then the JSP engine should create a new instance of the bean class `AddressBean`, add the instance to the session scope with the name `address`, and execute the body of the `<jsp:useBean>` tag before continuing. In this case, the nested scriptlet code is executed, allowing us to set the `street` property to an initial value of `"123 Main St."` each time the bean is instantiated.

It is the second point that gives us the ability to initialize newly created beans. In this example, we have used a scriptlet to set the `street` property of the bean. However, a cleaner way to initialize the properties is to use the `<jsp:setProperty>` action. We will discuss this action in section 14.2.2.

In practice, in addition to initializing the bean, the body of the `<jsp:useBean>` tag can be used to write any valid JSP code (HTML, scriptlets, expressions, and so forth). Just remember that it is executed only when the `<jsp:useBean>` tag requires that a bean be created.

Scope of the declared variable

When a `<jsp:useBean>` declaration is enclosed inside a Java programming language block using scriptlets and a pair of curly braces, the scope of the variable declared by the action also gets restricted to the enclosing block. Consider the following code, which will not execute correctly:

```
<%@ page language="java" import="AddressBean" %>
<html><body>

  <%
```

```

        if (true)
        {
    %>
        <jsp:useBean id="address" class="AddressBean" />
    <%
        }
    %>
<B>Some HTML here</B>

    <%
        if (true)
        {
            out.print("Zip: "+ address.getZip()); //error here
        }
    %>

</body></html>

```

In this example, the `<jsp:useBean>` declaration is enclosed inside a block using a pair of curly braces. This marks the scope of the `address` variable declared by the action. When this variable is used within the `out.print()` method, which is in a different block, the compiler flags an error indicating that the `address` variable is not defined.

However, even though the `address` variable is out of scope because of the pair of curly braces, the `address` bean is still reachable in the page scope. We can access it by using the implicit variable `pageContext`:

```

    <%
        if (true)
        {
            AddressBean address = (AddressBean)
                pageContext.getAttribute("address");

            out.print("Zip: "+ address.getZip()); //ok
        }
    %>

```

If the `<jsp:useBean>` action specifies a different scope, such as the request, session, or application scope, then we can use the corresponding implicit variable and the `getAttribute()` method to access the declared bean in that scope.

14.2.2 Mutating properties using `<jsp:setProperty>`

The `<jsp:setProperty>` action assigns new values to the bean's properties. It has four attributes, as described in table 14.3.

Table 14.3 The `<jsp:setProperty>` attributes

Attribute Name	Description
name	The name by which the bean is identified in the JSP page
property	The name of the property of the bean, which is to be given a new value

continued on next page

Table 14.3 The <jsp:setProperty> attributes (continued)

Attribute Name	Description
value	The new value to be assigned to the property
param	The name of the parameter available in the <code>HttpServletRequest</code> , which is to be assigned as a new value to the property of the bean

The name attribute

The name attribute identifies a particular instance of an existing bean. Therefore, the name attribute is mandatory. The bean must have already been declared by a previous `<jsp:useBean>` action, and the value of the name attribute must be the same as the value of the `id` attribute specified by the `<jsp:useBean>` action.

The property attribute

The `property` attribute specifies the property of the bean to be set. The JSP engine calls the `setXXX()` method on the bean based on the specified property. Thus, this attribute is also mandatory.

The value attribute

The `value` attribute specifies the new value to be set for the bean's property. This attribute can also accept a request-time attribute expression.

The param attribute

The `param` attribute specifies the name of the request parameter. If the request contains the specified parameter, then the value of that parameter is used to set the bean's property.

The `value` and `param` attributes are never used together and are both optional. If neither of the two is specified, it is equivalent to having the same value for both `param` and `property`, and the JSP engine searches for a request parameter with the name that is same as the `property` attribute.

Let's look at the following examples to understand how these attributes are used. For each of the examples, assume that we have already declared the use of the bean as follows:

```
<%@ page import="chapter14.AddressBean" %>
<jsp:useBean id="address" class="chapter14.AddressBean" />
```

Note that in all the examples below, we have shown the equivalent scriptlet code that we can write instead of the standard actions, `getProperty` and `setProperty`. The JSP engine, however, does not always translate these actions into such code. It uses reflection to check if the bean has the specified properties and then calls the appropriate methods.

Using the value attribute

These two actions instruct the JSP engine to use the bean named `address` and set its `city` and `state` properties to the values "Albany" and "NY", respectively:

```
<jsp:setProperty name="address" property="city" value="Albany" />
<jsp:setProperty name="address" property="state" value="NY" />
```

They are equivalent to the scriptlet code:

```
<%
    address.setCity("Albany");
    address.setState("NY");
%>
```

The following example uses a request-time expression for the value attribute:

```
<% String theCity = getCityFromSomewhere(); %>
<jsp:setProperty name="address"
    property="city"
    value="<%= theCity %>" />
```

Using the param attribute

In this case, instead of specifying the values using the value attributes, we have specified the request parameter names using the param attributes:

```
<jsp:setProperty name="address" property="city" param="myCity" />
<jsp:setProperty name="address" property="state" param="myState" />
```

This instructs the engine to get the values of the `myCity` and `myState` request parameters and set them to the "city" and "state" properties, respectively. Thus, the above tags are equivalent to the following scriptlet code:

```
<%
    address.setCity(request.getParameter("myCity"));
    address.setState(request.getParameter("myState"));
%>
```

Using the default param mechanism

The technique of setting the properties shown in the previous example is used when the names of the request parameters do not match the names of the bean properties. If the names of the request parameters match the names of the bean properties, we do not need to specify either the param or the value attribute, as shown here:

```
<jsp:setProperty name="address" property="city" />
<jsp:setProperty name="address" property="state" />
```

In this case, the bean properties are set using the corresponding values from the request parameters. The above tags are equivalent to

```
<jsp:setProperty name="address" property="city" param="city" />
<jsp:setProperty name="address" property="state" param="state" />
```

These in turn are equivalent to the following scriptlet code:

```
<%
    address.setCity(request.getParameter("city"));
    address.setState(request.getParameter("state"));
%>
```

You might ask, “What if there is no such parameter in the request?” If the parameter is not present in the request, or if it has a value of “ ” (empty string), the `<jsp:setProperty>` action has no effect and the property retains its original value.

Setting all the properties in one action

The following is a shortcut to set all the properties of a bean in a single action:

```
<jsp:setProperty name="address" property="*" />
```

Instead of setting each property of the `address` bean one by one, we can set all the properties to the respective values present in the request parameters using a value of “*” for the `property` attribute. The above tag is equivalent to the following scriptlet code:

```
<%
    address.setStreet(request.getParameter("street"));
    address.setCity(request.getParameter("city"));
    address.setState(request.getParameter("state"));
    address.setZip(request.getParameter("zip"));
%>
```

Obviously, the names of the request parameters must match the names of the properties. As we mentioned earlier, if there is no matching parameter in the request for a particular property, the value of that property remains unchanged. Also, in all the above cases where we use the `param` attribute, if the request parameter has multiple values, then only the first value is used.

14.2.3 Accessing properties using `<jsp:getProperty>`

The `<jsp:getProperty>` action is used to retrieve and print the values of the bean properties to the output stream. The syntax of this action is quite simple:

```
<jsp:getProperty name="beanInstanceName"
    property="propertyName" />
```

It has only two attributes, `name` and `property`, which are both mandatory. As in the `<setProperty>` action, the `name` attribute specifies the name of the bean instance as declared by a previous `<useBean>` action and the `property` attribute specifies the property whose value is to be printed.

The following actions instruct the JSP engine to print out the values of the `state` and the `zip` properties of the `address` bean:

```
<jsp:getProperty name="address" property="state" />
<jsp:getProperty name="address" property="zip" />
```


They are equivalent to the scriptlet code:

```
<%
    out.print(address.getState());
    out.print(address.getZip());
%>
```

The code in listing 14.4 locates an instance of `AddressBean` in the session scope and prints out its properties in a tabular format.

Listing 14.4 addressDisplay.jsp

```
<%@ page import="chapter14.AddressBean" %>
<jsp:useBean id="address" class="chapter14.AddressBean" scope="session"/>

<html><body>
<table>
  <tr>
    <td>Street</td>
    <td><jsp:getProperty name="address" property="street"/></td>
  </tr>
  <tr>
    <td>City</td>
    <td><jsp:getProperty name="address" property="city"/></td>
  </tr>
  <tr>
    <td>State</td>
    <td><jsp:getProperty name="address" property="state"/></td>
  </tr>
  <tr>
    <td>Zip</td>
    <td><jsp:getProperty name="address" property="zip"/></td>
  </tr>
</table>
</body></html>
```

Quizlet

Q: Consider the following code from a file named `addressInput.jsp`:

```
<%@ page import="chapter14.AddressBean" %>

<jsp:useBean id="address" class="chapter14.AddressBean"
scope="request" />

<jsp:setProperty name="address" property="*" />

<jsp:forward page="addressDisplay.jsp" />
```

Can the `addressDisplay.jsp` file access the address bean declared in the `addressInput.jsp` file and print its values using `<jsp:getProperty>`?

- A:** Yes, the page `addressDisplay.jsp` file can print the values of the bean properties using `<jsp:getProperty>` provided it also contains a `<jsp:useBean>` declaration that is identical to the one shown in `addressInput.jsp` and that the declaration appears before the `<jsp:getProperty>` declaration.

14.3 JAVABEANS IN SERVLETS

We know that JSP pages are converted into servlets at translation time, which means that the beans that we use in our JSP pages are actually used from a servlet. This implies that we can use JavaBeans from servlets, too. This section discusses the ways in which we can share beans between JSP pages and servlets. The exam requires you to know the servlet code that is equivalent to using beans in the different scopes: request, session, and application.

Suppose a JSP page uses three beans, each with a different scope—request, session, and application—declared as

```
<jsp:useBean id="address1" class="chapter14.AddressBean"
scope="request" />
<jsp:useBean id="address2" class="chapter14.AddressBean"
scope="session" />
<jsp:useBean id="address3" class="chapter14.AddressBean"
scope="application" />
```

Listing 14.5 shows how to achieve the same functionality in the servlet code.

Listing 14.5 Using JavaBeans in servlets

```
import javax.servlet.*;
import javax.servlet.http.*;
import chapter14.AddressBean;

public class BeanTestServlet extends HttpServlet
{
    public void service(HttpServletRequest request,
                        HttpServletResponse response)
                        throws java.io.IOException,
                        ServletException
    {
        AddressBean address1 = null;
        AddressBean address2 = null;
        AddressBean address3 = null;

        //Get address1 using the parameter request
        synchronized(request)
        {
            address1 = (AddressBean)
                request.getAttribute("address1");

            if (address1==null)
```

```

        {
            address1 = new AddressBean();
            request.setAttribute("address1", address1);
        }
    }

    //Get address2 using HttpSession
    HttpSession session = request.getSession();

    synchronized(session)
    {
        address2 = (AddressBean)
            session.getAttribute("address2");

        if (address2==null)
        {
            address2 = new AddressBean();
            session.setAttribute("address2", address2);
        }
    }

    // Get address3 using ServletContext
    ServletContext servletContext = this.getServletContext();

    synchronized(servletContext)
    {
        address3 = (AddressBean)
            servletContext.getAttribute("address3");

        if (address3==null)
        {
            address3 = new AddressBean();
            servletContext.setAttribute("address3", address3);
        }
    }
}

} //service
} //class

```

This simple example demonstrated the use of the three container objects—`HttpServletRequest`, `HttpSession`, and `ServletContext`—that allow us to share JavaBeans between servlets and JSP pages in the three scopes—request, session, and application. An important point to remember here is that we have to synchronize access to the three container objects because other servlets and JSP pages may be accessing the same objects simultaneously in more than one thread. If we want to use serialized beans from a servlet, we can use the following method:

```

java.beans.Beans.instantiate(
    this.getClass().getClassLoader(),
    "businessData.John");

```

We will not discuss this method because you are not required to know its details for the exam. Please refer to the API documentation for more information.

Quizlet

Q: Consider the following servlet code. How will you achieve the same effect in a JSP page?

```
AddressBean address;
ServletContext servletContext =
    this.getServletContext ();

synchronized(servletContext)
{
    address = (AddressBean)
        servletContext.getAttribute("address");

    if (address==null)
    {
        address = new BusinessAddressBean();
        address.setCity("Greenwich");
        address.setState("Connecticut");
        servletContext.setAttribute("address", address);
    }
}
```

A: You should notice four points about this code. First, the code uses `ServletContext` to get and set the named object `address`. This means that it uses the application scope. Second, if the object is not found, the code creates a new instance of the class with the new keyword and does not use the `java.beans.Beans` mechanism. This means we should use the `class` attribute instead of the `beanName` attribute. Third, the declared variable `address` is of type `AddressBean`, while the new instance created is of type `BusinessAddressBean`. This means that we must use the `type` attribute with `AddressBean` as its value but that the `class` attribute must have `BusinessAddressBean` as its value. Fourth, it sets two properties, `city` and `state`, whenever the object `address` is not found in the application scope and a new instance is created. This means we must initialize the bean using `<jsp:setProperty>` tags within the opening and closing `<jsp:useBean>` tags. Thus, it is equivalent to the following JSP code:

```
<jsp:useBean id="address"
    type="AddressBean"
    class="chapter14.BusinessAddressBean"
    scope="application" >
    <jsp:setProperty name="address" property="city"
        value="Greenwich" />
    <jsp:setProperty name="address" property="state"
        value="Connecticut" />
</jsp:useBean>
```

14.4 ACCESSING JAVABEANS FROM SCRIPTING ELEMENTS

As we have seen, one of the main advantages of using JavaBeans in JSP pages is that they help to keep the code clean when they are used with the standard actions. However, JavaBeans can also be used in scripting elements. Say, for instance, that a bean has some processing capabilities in addition to its normal function of holding a set of properties. The bean may retrieve values from the database as it is initialized; in such cases, it may have methods that are not, or cannot be, implemented as setters and getters. For example, suppose we have a `UserBean` that stores user profiles, and we need to set the login and password properties submitted by the user and then load the user information from the database. The following code snippet shows how to do this:

```
<%@ page import="chapter14.UserBean" %>

<jsp:useBean id="user" class="chapter14.UserBean" scope="session">

    <jsp:setProperty name="user" property="login" />
    <jsp:setProperty name="user" property="password" />

    <%
        //The bean is used in a scriptlet here.
        //Load the user information from the database.
        user.initialize();
    %>

</jsp:useBean>
```

Here, we first create an instance of the `UserBean` using the `<jsp:useBean>` action. Then we set its login and password properties using the `<jsp:setProperty>` action. But after that we want to initialize it using its `initialize()` method. Since there is no standard JSP action that can be used to achieve this, we have to use a scriptlet³ as shown above. Within the scriptlet, we can use the `user` variable to refer to the bean instance since the `<jsp:useBean>` action declares it automatically.

Another reason for using beans in scriptlets and expressions is that the standard action `<jsp:getProperty>` writes out the property value directly into the output stream. It cannot be used for writing conditional logic or for passing it as a value to an attribute. For example, suppose `UserBean` has a property, named `loginStatus`, which is set to `true` or `false` depending on whether the login attempt was successful. We cannot use the `<jsp:getProperty>` action in an `if` condition to test it. The following is not valid:

```
<%
    if (<jsp:getProperty
        name="user"
        //error here
```

³ We can also define and use custom tags to work with beans instead of scriptlets. Custom tags are explained in chapters 15, 16, and 17.

```

        property="loginStatus" />
    }
}
%>

```

Similarly, if `UserBean` has a property named `preferredHomePage` that stores a URL to the user's preferred home page, then `<jsp:getProperty>` cannot be used to pass request-time values to the `<jsp:forward>` action. The following is not valid:

```

<jsp:forward page="<jsp:getProperty
              name="user"
              property="preferredHomePage" />" //error here
/>

```

In such cases, we have to use the bean in a scriptlet and an expression in this way:

```

<% if (user.getLoginStatus()) { %>
    <jsp:forward page="<%=user.getPreferredHomePage()%>" />
} else {
    <jsp:forward page="loginError.jsp" >
<% } %>

```

Quizlet

Q: What is wrong with the following code?

```

<jsp:useBean id="address"
             class="chapter14.AddressBean"
             beanName="businessData.visitorAddresses.John" />

```

A: We cannot use the attributes `beanName` and `class` in the same `<jsp:useBean>` declaration.

Q: What is wrong with the following code?

```

<jsp:setProperty name="address"
                 param="state"
                 value="FL" />

```

A: We have to use the mandatory attribute `property` to specify the property of the bean. `param` specifies the request parameter whose value is to be used. We cannot use the `param` and `value` attributes in the same `<jsp:setProperty>` action.

Q: How can we get all of the properties of a bean in a single JSP action?

A: We can set all of the properties of a bean in a single action:

```

<jsp:setProperty name="beanName" property="*" />

```

But there is no way to get all of the properties of a bean in a single action.

14.5 MORE ABOUT PROPERTIES IN JAVABEANS

In the `AddressBean` example we have used throughout this chapter, all of the properties have been of type `java.lang.String`. However, a bean can have any type of property, such as

- Primitive data types (`int`, `char`, `boolean`, etc.)
- Wrapper object types (`java.lang.Integer`, `java.lang.Character`, `java.lang.Boolean`, etc.)
- Other object types
- Array types (`int []`, `Integer []`, etc.)

In this section, we will take a closer look at the way JSP pages manage these nonstring data types in JavaBeans. To illustrate this in the short examples that follow, we will be using a `UserBean` class that stores three different types of data:

```
public class chapter14.UserBean{  
  
    private int visits;           //An example of primitive type  
    private Boolean valid;       //An example of wrapper type  
    private char[] permissions; //An example of index type  
    //appropriate setters and getters go here  
}
```

The `visits` variable counts the number of visits by this user. The `valid` property indicates whether this bean instance has been initialized and is valid or not. Notice that the `permissions` property is an array of `char`. Such properties are called *indexed properties*. We will learn more about them in section 14.5.2.

14.5.1 Using nonstring data type properties

Request parameters are always of the type `String`. Hence, if we are expecting a value that is of a type other than `String`, then we have to do the necessary conversions in scriptlets before using the value. For example, the following scriptlet converts an incoming request parameter from a `String` to an `int` and to an `Integer`:

```
<%  
String numAsString = null;  
int numAsInt = 0;  
Integer numAsInteger = null;  
try{  
    numAsString = request.getParameter("num");  
    numAsInt = Integer.parseInt(numAsString);  
    numAsInteger = Integer.valueOf(numAsString);  
}  
catch(NumberFormatException nfe){  
}  
%>
```

However, if a JavaBean has nonstring properties, then the standard JSP actions `<jsp:setProperty>` and `<jsp:getProperty>` perform the necessary conversions automatically, as explained below.

Automatic type conversion in `<jsp:setProperty>`

When we use literal values to set nonstring properties in a bean using the `<jsp:setProperty>` action, the container performs the appropriate conversion from `String` to the property type. For example, in the following actions, the engine converts the literals `30` and `true` to `int` and `Boolean`, respectively:

```
<jsp:setProperty name="user" property="visits" value="30" />
<jsp:setProperty name="user" property="valid" value="true" />
```

These actions are equivalent to the following scriptlet:

```
<%
    user.setVisits(Integer.valueOf("30").intValue());
    user.setValid(Boolean.valueOf("true"));
%>
```

The type conversion occurs automatically even in the case of request parameters. For example, if we do not specify any value in the `<jsp:setProperty>` action as shown below but pass the values using a query string in the URL, the JSP engine automatically performs the conversions:

```
<jsp:setProperty name="user" property="visits" />
<jsp:setProperty name="user" property="valid" />
```

This will work without errors when we call the JSP page using the following URL:

```
http://localhost:8080/chapter14/test.jsp?visits=30&valid=true
```

However, if we use request-time attribute expressions in the `<jsp:setProperty>` action, then no such automatic conversion happens. Consider the following:

```
<%
    String anIntAsString = "30";
    String aBoolAsString = "true";
%>

<jsp:setProperty name="user"
    property="visits"
    value="<%= anIntAsString %>" />

<jsp:setProperty name="user"
    property="valid"
    value="<%= aBoolAsString %>" />
```

This example will not compile because we have passed a request-time attribute expression value that evaluates to a `String` in both actions instead of an `int` and a `Boolean`, respectively. To make it work, we have to explicitly convert the values from a

String to int for the `visits` property and from a String to a Boolean for the `valid` property, as follows:

```
<jsp:setProperty
  name="user"
  property="visits"
  value="<%= Integer.valueOf(anIntAsString).intValue() %>" />

<jsp:setProperty
  name="user"
  property="valid"
  value="<%= Boolean.valueOf(aBoolAsString) %>" />
```

Automatic type conversion in <jsp:getProperty>

When we use nonstring data type properties of a bean in the `<jsp:getProperty>` action, the action takes care of the conversion from the given type to String. For example, the following actions use an int and a Boolean to print the values:

```
<jsp:getProperty name="user" property="visits" />
<jsp:getProperty name="user" property="valid" />
```

They are equivalent to the following scriptlet:

```
<%
  out.print(user.getVisits()); //getVisits() returns int
  out.print(user.getValid()); //getValid() returns Boolean
%>
```

14.5.2 Using indexed properties

Indexed properties can be used to associate multiple values to a single property. We can set an indexed property in one of two ways:

- We can set it automatically from the request parameter.
- We can set it using a request-time expression and explicitly pass an array of the desired type.

Setting an indexed property from request parameters

The following action sets the indexed property `permissions` to the values received in the request parameter:

```
<jsp:setProperty name="user" property="permissions" />
```

Suppose we access the JSP page with this URL:

```
http://localhost:8080/chapter14/test.jsp?
  permissions=XYZ&permissions=PQR&permissions=L
```

The engine will create an array of `chars` with the length of the array equal to the number of request parameter values for the parameter named `permissions`. Since the URL specifies three values for `permissions` as `permissions=XYZ&permissions=PQR&permissions=L`, it will create an array of three `chars`. Then, for

each individual element of the char array, it will convert the parameter value from String to char using `String.charAt(0)`. Thus, the `<jsp:setProperty>` action will use the values X, B, and L. It is equivalent to

```
char charArr[] = new char[3];

charArr[0] = (request.getParameterValues("permissions"))[0].charAt(0);
// "XYZ".charAt(0)

charArr[1] = (request.getParameterValues("permissions"))[1].charAt(0);
// "PQR".charAt(0)

charArr[2] = (request.getParameterValues("permissions"))[2].charAt(0);
// "L".charAt(0)
```

Setting an indexed property using request-time attribute expressions

The following action sets the indexed property permissions using a request-time attribute expression:

```
<%!
    char myPermissions[] = {'A', 'B', 'C' };
%>

<jsp:setProperty
    name="user"
    property="permissions"
    value="<%= myPermissions %>" />
```

This action is simple and will use the values A, B, and C.

The approaches we have discussed for setting indexed properties apply to all of the data types. Whether they are `String`, primitive data types, wrapper objects, or other objects, the parameter values are converted to their respective types for each individual element of the array.

Getting indexed properties

When we get indexed properties from a bean using the `<jsp:getProperty>` action, it is equivalent to calling `out.print(property-type [])`, where `property-type` is the data type of the indexed property. Consider the following example:

```
<jsp:getProperty name="user" property="permissions" />
```

This is equivalent to the following scriptlet:

```
<%
    out.print(user.getPermissions()); // getPermissions returns char[]
%>
```

However, since an array in Java is considered to be an `Object`, the `<jsp:getProperty>` action is not very useful for indexed properties; it simply prints the internal reference of the `Object` in memory. The output of the action might look something like this:

```
char []@0xcafebabe
```

Thus, to print a particular property, or all the properties, we have to use scripting elements:

```
<%
    char[] permissions = user.getPermissions();

    if (permissions != null)
    {
        for (int p = 0; p<permissions.length; p++ )
        {
            Permission is <%= permissions[p] %> <br>
        }
    }
%>
```

Here we are using scriptlets to get the permissions, and we are using an expression to print them.

14.6 SUMMARY

JSP technology is designed to take advantage of the power of JavaBeans components. In this chapter, we discussed JavaBeans from the JSP developer's point of view. We noted the value they provide by reducing the code in JSP pages, thereby increasing the readability of the pages. In JSP pages, any class can be used as a JavaBean as long as it has a public constructor with no arguments and private properties that are accessed by public getter and setter methods. The JSP specification provides the following standard actions to use JavaBeans in JSP pages: `<jsp:useBean>`, `<jsp:setProperty>`, and `<jsp:getProperty>`.

Servlets can also access JavaBeans objects. We reviewed the servlet code to access JavaBeans in the different scopes, and we compared it with the equivalent JSP code. At times, we need to access a JavaBean from scripting elements, so we discussed using beans with scriptlets and expressions and under what circumstances we would do that.

Not all properties of JavaBean are of the type `java.lang.String`. We reviewed the way that the `<jsp:setProperty>` and `<jsp:getProperty>` standard actions handle using nonstring data types, and then we looked at indexed properties and how they are managed in JSP pages.

With the end of this chapter, you should be ready to answer exam questions about the way JavaBeans can be declared, initialized, and used in JSP pages as well as in servlets and how the standard actions—`<jsp:useBean>`, `<jsp:setProperty>`, and `<jsp:getProperty>`—help us to reduce the use of scriptlets and write cleaner JSP pages.

Now that we have a good understanding of the standard tags, in the next chapter we will explore the use of custom tags in JSP pages.

14.7 REVIEW QUESTIONS

- Which of the following is a valid use of the `<jsp:useBean>` action? (Select one)
 - `<jsp:useBean id="address" class="chapter14.AddressBean" />`
 - `<jsp:useBean name="address" class="chapter14.AddressBean"/>`
 - `<jsp:useBean bean="address" class="chapter14.AddressBean" />`
 - `<jsp:useBean beanName="address" class="chapter14.AddressBean" />`
- Which of the following is a valid way of getting a bean's property? (Select one)
 - `<jsp:useBean action="get" id="address" property="city" />`
 - `<jsp:getProperty id="address" property="city" />`
 - `<jsp:getProperty name="address" property="city" />`
 - `<jsp:getProperty bean="address" property="*" />`
- Which of the following are valid uses of the `<jsp:useBean>` action? (Select two)
 - `<jsp:useBean id="address" class="chapter14.AddressBean" name="address" />`
 - `<jsp:useBean id="address" class="chapter14.AddressBean" type="AddressBean" />`
 - `<jsp:useBean id="address" beanName="AddressBean" class="chapter14.AddressBean" />`
 - `<jsp:useBean id="address" beanName="AddressBean" type="AddressBean" />`
- Which of the following gets or sets the bean in the `ServletContext` container object? (Select one)
 - `<jsp:useBean id="address" class="chapter14.AddressBean" />`
 - `<jsp:useBean id="address" class="chapter14.AddressBean" scope="application" />`
 - `<jsp:useBean id="address" class="chapter14.AddressBean" scope="servlet" />`
 - `<jsp:useBean id="address" class="chapter14.AddressBean" scope="session" />`
 - None of the above
- Consider the following code:

```
<html><body>
<jsp:useBean id="address" class="chapter14.AddressBean"
scope="session" />
state = <jsp:getProperty name="address" property="state" />
</body></html>
```

Which of the following are equivalent to the third line above? (Select three)

```

a <% state = address.getState(); %>
b <% out.write("state = "); out.print(address.getState()); %>
c <% out.write("state = "); out.print(address.getstate()); %>
d <% out.print("state = " + address.getState()); %>
e state = <%= address.getState() %>
f state = <%! address.getState(); %>

```

6. Which of the options locate the bean equivalent to the following action? (Select three)

```

<jsp:useBean id="address" class="chapter14.AddressBean"
scope="request" />

```

```

a request.getAttribute("address");
b request.getParameter("address");
c getServletContext().getRequestAttribute("address");
d pageContext.getAttribute("address", PageContext.REQUEST_SCOPE);
e pageContext.getRequest().getAttribute("address");
f pageContext.getRequestAttribute("address");
g pageContext.getRequestParameter("address");

```

7. Consider the following code for `address.jsp`:

```

<html><body>
<jsp:useBean id="address" class="chapter14.AddressBean" />
<jsp:setProperty name="address" property="city" value="LosAngeles" />
<jsp:setProperty name="address" property="city" />
<jsp:getProperty name="address" property="city" />
</body></html>

```

What is the output if the above page is accessed via the URL

```

http://localhost:8080/chap14/address.jsp?city=Chicago&city=Miami

```

Assume that the `city` property is not an indexed property. (Select one)

```

a LosAngeles
b Chicago
c Miami
d ChicagoMiami
e LosAngelesChicagoMaimi
f It will not print anything because the value will be null or "".

```

8. Consider the following code:

```

<html><body>

<{%>
<jsp:useBean id="address" class="chapter14.AddressBean"
scope="session" />

```

```
<%}%>
//1
</body></html>
```

Which of the following can be placed at line //1 above to print the value of the street property? (Select two)

- a `<jsp:getProperty name="address" property="street" />`
- b `<% out.print(address.getStreet()); %>`
- c `<%= address.getStreet() %>`
- d `<%= ((AddressBean) session.getAttribute("address")).getStreet() %>`
- e None of the above; the bean is nonexistent at this point.

9. Consider the following code:

```
<html><body>
<%{%>
<jsp:useBean id="address" class="chapter14.AddressBean"
scope="session" />
<%}%>
<jsp:useBean id="address" class="chapter14.AddressBean"
scope="session" />
<jsp:getProperty name="address" property="street" />
</body></html>
```

Which of the following is true about the above code? (Select one)

- a It will give translation-time errors.
- b It will give compile-time errors.
- c It may throw runtime exceptions.
- d It will print the value of the street property.

10. Consider the following servlet code:

```
//...
public void service (HttpServletRequest request,
                    HttpServletResponse response)
    throws IOException, ServletException
{
    //1
}
```

Which of the following can be used at //1 to retrieve a JavaBean named address present in the application scope? (Select one)

- a `getServletContext().getAttribute("address");`
- b `application.getAttribute("address");`

```
c request.getAttribute("address",APPLICATION_SCOPE);
d pageContext.getAttribute("address",APPLICATION_SCOPE);
```

11. Consider the following code, contained in a file called `this.jsp`:

```
<html><body>
<jsp:useBean id="chapter14.address" class="AddressBean" />
<jsp:setProperty name="address" property="*" />
<jsp:include page="that.jsp" />
</body></html>
```

Which of the following is true about the `AddressBean` instance declared in this code? (Select one)

- a The bean instance will not be available in `that.jsp`.
- b The bean instance may or may not be available in `that.jsp`, depending on the threading model implemented by `that.jsp`.
- c The bean instance will be available in `that.jsp`, and the `that.jsp` page can print the values of the beans properties using `<jsp:getProperty />`.
- d The bean instance will be available in `that.jsp` and the `that.jsp` page can print the values of the bean's properties using `<jsp:getProperty />` only if `that.jsp` also contains a `<jsp:useBean/>` declaration identical to the one in `this.jsp` and before using `<jsp:getProperty/>`.

12. Consider the following code contained in a file named `this.jsp` (the same as above, except the fourth line):

```
<html><body>
<jsp:useBean id="address" class="chapter14.AddressBean" />
<jsp:setProperty name="address" property="*" />
<%@ include file="that.jsp" %>
</body></html>
```

Which of the following is true about the `AddressBean` instance declared in the above code? (Select one)

- a The bean instance will not be available in `that.jsp`.
- b The bean instance may or may not be available in `that.jsp`, depending on the threading model implemented by `that.jsp`.
- c The bean instance will be available in `that.jsp`, and the `that.jsp` page can print the values of the bean's properties using `<jsp:getProperty />`.
- d The bean instance will be available in `that.jsp`, and the `that.jsp` page can print the values of the bean's properties using `<jsp:getProperty />` only if `that.jsp` also contains a `<jsp:useBean/>` declaration identical to the one in `this.jsp` and before using `<jsp:getProperty/>`.



CHAPTER 15

Using custom tags

- | | |
|--|--|
| 15.1 Getting started 286 | 15.4 Using the JSP Standard Tag Library (JSTL) 298 |
| 15.2 Informing the JSP engine about a custom tag library 288 | 15.5 Summary 305 |
| 15.3 Using custom tags in JSP pages 293 | 15.6 Review questions 305 |

EXAM OBJECTIVES

- 6.6** Configure the deployment descriptor to declare one or more tag libraries, deactivate the evaluation language, and deactivate the scripting language.
- 9.1** For a custom tag library or a library of Tag Files, create the ‘taglib’ directive for a JSP page.
(Section 15.2)
- 9.2** Given a design goal, create the custom tag structure in a JSP page to support that goal.
(Sections 15.2, 15.3, and 15.4)
- 9.3** Given a design goal, use an appropriate JSP Standard Tag Library (JSTL v1.1) tag from the “core” tag library.
(Section 15.4)

INTRODUCTION

As we saw in chapter 12, “Reusable web components,” and chapter 14, “Using JavaBeans,” the JSP specification provides standard XML type tags, called JSP actions, that instruct the JSP engine to take some action in a predefined manner.

Although very useful, the standard tags provide just a basic set of features. As your web application grows, you will find that these standard JSP tags are somewhat restrictive and don’t provide support for the presentation logic that is required for formatting dynamic data. For instance, you may be forced to write too much of your presentation code in JSP scriptlets. Moreover, you may have to copy and paste those presentation scriptlets onto multiple pages of the application. At this point, you need a way to put that presentation logic in one place and reuse it wherever it is required. The JSP technology provides a feature that allows you to do just that. You can create new tags and define their behavior according to your needs. These user-defined tags are called *custom tags*.

In addition, we’ll present a predeveloped set of tags that make up the “core” library of the JSP Standard Tag Library (JSTL). Rather than build your own, you can use these tags to perform flow control and conditional processing in your JSP.

15.1 GETTING STARTED

Custom tags do not introduce any new syntax. They are similar to the standard JSP actions and follow the same XML syntax format. In that sense, it may be more accurate to refer to them as custom *actions* rather than *tags*. Custom tags allow us to move the presentation logic outside the JSP pages into independent Java classes, thereby centralizing the implementation of the presentation logic and increasing maintainability. By using custom tags in JSP pages instead of scriptlets, we avoid duplicating the presentation logic; removing the scriptlets also makes the pages less cluttered and easier to read.

15.1.1 New terms

New concepts bring new terms. So let’s begin our foray into the world of custom tags by becoming familiar with the terminology.

Tag handler

The JSP specification defines a *tag handler* as a runtime, container-managed object that evaluates custom actions during the execution of a JSP page.

In practical terms, a tag handler is a Java class that implements one of two types of interfaces. The first type contains the “Classic” tag interfaces—`Tag`, `IterationTag`, or `BodyTag`—and will be discussed in the following chapter. The second type contains the “Simple” tag interface, `SimpleTag`, and will be presented in chapter 17. For now, just remember that while executing a JSP file, if the JSP engine encounters a custom tag, it calls the methods on the tag’s handler class to do the actual work.

Tag library

The JSP specification defines a *tag library* as a collection of actions that encapsulate some functionality to be used from within a JSP page.

Typically, we would not create just one tag to fulfill a particular requirement. Rather, we would design and develop a set of custom tags that work together and help solve a recurring requirement. Such a set of custom tags is called a tag library.

Tag library descriptor

When we use custom tags in a JSP page, the JSP engine needs to know the tag handler classes for these tags, in which tag library they are located, and how are they used. This meta-information is stored in a file called the *tag library descriptor* (TLD).

Types of URIs

In a JSP page, we reference the tag libraries through URIs. Table 15.1 describes the three types of URIs that are used in a JSP page.

Table 15.1 Types of URIs for referring to a tag library

Type	Description	Examples
Absolute URI	A URI that has a protocol, a hostname, and optionally a port number.	<code>http://localhost:8080/taglibs</code> <code>http://www.manning.com/taglibs</code>
Root Relative URI	A URI that starts with a / and that does not have a protocol, a hostname, or a port number. It is interpreted as relative to the document root of the web application.	<code>/helloLib</code> <code>/taglib1/helloLib</code>
Non-root Relative URI	A URI that does not start with a / and that does not have a protocol, a hostname, or a port number. It is interpreted as relative to either the current JSP page or the <code>WEB-INF</code> , depending on where it is used.	<code>HelloLib</code> <code>taglib2/helloLib</code>

15.1.2 Understanding tag libraries

You don't need to create your own tag libraries to handle many common functions. A variety of custom tag libraries are available on the Internet that you can use in our JSP pages. For example, the libraries provided by the Jakarta Apache Project at `http://jakarta.apache.org/taglibs` contain several frequently used features, such as text manipulation or date manipulation. Sun Microsystems has also developed a JSP Standard Tag Library that we will discuss later in this chapter.

To use an existing tag library, you need to know two things:

- How to inform the JSP engine of the location of the TLD file of a tag library
- How to use the custom tags provided by the tag library in JSP pages

The exam objectives covered in this chapter focus on these two points.

However, if none of the existing libraries suits your needs and if you plan to implement a custom tag library on your own, you need to know two more things:

- How to implement the tag handlers for your tag library
- How to describe your tag library in a TLD file

We will discuss these two topics in the next chapter, where we will develop a sample tag library.

15.2 **INFORMING THE JSP ENGINE ABOUT A CUSTOM TAG LIBRARY**

In the JSP syntax, we can import a new tag library into a JSP page using a `taglib` directive:

```
<%@ taglib prefix="test" uri="sampleLib.tld" %>
```

If we are using the XML syntax, the library information in the JSP document is included in the `<jsp:root>` element:

```
<jsp:root
  xmlns:jsp="http://java.sun.com/JSP/Page"
  xmlns:test="sampleLib.tld"
  version="2.0" >
  ...JSP PAGE...
</jsp:root>
```

The above declarations inform the engine that the page uses custom tags with the prefix `test` and that these tags are described in the file `sampleLib.tld`. In this example, the value of the URI attribute provides a nonroot, or page-relative path, to the TLD file. The JSP engine searches for the file in the same directory as the JSP page.

Though keeping all the JSP pages and the TLD files in the same directory is the simplest way to use a `taglib` directive, it has two major drawbacks: security and flexibility.

Let's look at security first. Suppose the URL of the JSP page is `http://www.someserver.com/sample.jsp`. A visitor can view the contents of your library without much effort by typing the URL `http://www.someserver.com/sampleLib.tld`.

Of course, we can configure the web server to restrict access to all the TLD files, or better yet, we can put the TLD files under the `/WEB-INF` directory and use the path `/WEB-INF/sampleLib.tld` to access them. However, we would still have the problem of flexibility. If we wanted to switch to a newer version of the library, say `sampleLib_2.tld`, then we would have to manually modify all the JSP pages that are affected by this change. Further, third-party custom tag libraries are packaged and shipped as JAR files. How would we indicate the location of a TLD file in such cases?

To avoid such problems, JSP provides a cleaner solution for indicating the use of tag libraries. The JSP container maintains a map between the URIs that we use in `taglib` directives and the actual physical location of the TLD files in the file system. With this approach, instead of using a page-relative path, we use an absolute URI path:

```
<%@ taglib prefix="test"
      uri="http://www.someserver.com/sampleLib" %>
```

When the JSP engine reads the above URI, it refers to its internal map to find the corresponding TLD file location.

Thus, by creating a level of indirection, this approach solves both the security and the flexibility problems. The actual TLD file can reside in the `WEB-INF` directory or even in a JAR file, hidden away from the visitors. If a newer version is released, all we have to do is update the mapping between the URI and the actual path.

NOTE The use of an absolute URL does not mean that the JSP engine will actually download the TLD file or the tag library classes from the specified URL. Thus, in the above example, the JSP engine will not try to locate the library at `http://www.someserver.com/sampleLib`. Consider the URI as just a name that is mapped to the actual location of the TLD file somewhere on the local machine.

In the following sections, we will discuss the possible locations for TLD files, how the mappings between the URIs and TLD locations are created, and how the JSP engine interprets the different values of URIs specified in the `taglib` directives.

15.2.1 Location of a TLD file

A TLD file can reside in one of two types of places. First, it can be placed in any directory of a web application; for example:

```
<docroot>/sampleLib.tld
<docroot>/myLibs/sampleLib.tld
<docroot>/WEB-INF/sampleLib.tld
<docroot>/WEB-INF/myLibs/sampleLib.tld
```

We usually keep the TLD file in a directory, instead of a JAR file, during the development of a tag library. This speeds up the development and testing cycles, during which we design new tags, add new handler classes, and modify the TLD file frequently. However, once development is finished, we package the handler classes and the TLD file of the library as a JAR file. This file is then deployed under the `<doc-root>/WEB-INF/lib` directory along with other jarred classes, such as servlets and third-party tools.

The JSP specification mandates that, when deployed in a JAR file, a TLD file be placed either directly under or inside a subdirectory of the `META-INF` directory. In addition, the name of the TLD file must be `taglib.tld`. Thus, a JAR file containing a packaged tag library is typically structured like this:

```
myPackage/myTagHandler1.class
myPackage/myTagHandler2.class
myPackage/myTagHandler3.class
META-INF/taglib.tld
```

The JSP container will recognize either of these two locations, a directory or a JAR, as a path to a TLD file. This path is called the *TLD resource path*.

Let's return for a moment to our discussion on the mapping between a URI and the location of a TLD file. We can now see that this mapping is actually between a URI and a TLD resource path, where the TLD resource path is either the path to the TLD file or to the JAR file containing the TLD file. This mapping is referred to as the *taglib map*.

15.2.2 Associating URIs with TLD file locations

The JSP container populates the taglib map in three ways:

- First, the container reads the user-defined mapping entries present in the deployment descriptor. This is called *explicit mapping*. We will learn how to add new entries to the deployment descriptor in the next section.
- Then, the container reads all the `taglib.tld` files present in the packaged JARs. For each jarred TLD file that contains information about its own URI, the JSP container automatically creates a mapping between the specified URI and the current location of the JAR file. This is called *implicit mapping*. We will learn how to add the URI information to a TLD file in the next chapter, where we will create a custom tag library.
- Finally, the JSP container adds entries for the URIs that are known to the container by default. These URIs are called *well-known URIs*. The `<jsp:root>` element of a JSP document contains an example of a well-known URI:

```
http://java.sun.com/JSP/Page
```

The container itself provides the implementation for all the tags in this library. This is actually another form of implicit mapping.

15.2.3 Understanding explicit mapping

We use the `<taglib>` element of the deployment descriptor file, `web.xml`, to associate user-defined URIs with TLD resource paths. This is the syntax:

```
<!ELEMENT taglib (taglib-uri, taglib-location)>
```

Each `<taglib>` element associates one URI with one location. It contains two subelements:

- `<taglib-uri>`. This is the user-defined URI. Its value can be an absolute URI, a root-relative URI, or a nonroot relative URI.
- `<taglib-location>`. This is the TLD resource path. Its value can be either a root-relative URI or a nonroot relative URI, and it must point to a valid TLD resource path.

Listing 15.1 illustrates the use of the `<taglib>` element in the deployment descriptor.

Listing 15.1 web.xml

```
<?xml version="1.0" encoding="ISO-8859-1"?>

<!DOCTYPE web-app
  PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
  "http://java.sun.com/dtd/web-app_2_3.dtd">

<web-app>

  <!-- other elements ... -->

  <!-- Taglib 1 -->

  <taglib>
    <taglib-uri>
      http://www.manning.com/studyKit
    </taglib-uri>
    <taglib-location>
      /myLibs/studyKit.tld
    </taglib-location>
  </taglib>

  <!-- Taglib 2 -->

  <taglib>
    <taglib-uri>
      http://www.manning.com/sampleLib
    </taglib-uri>
    <taglib-location>
      yourLibs/sample.jar
    </taglib-location>
  </taglib>

</web-app>
```

In listing 15.1, there are two `<taglib>` elements. The first `<taglib>` element associates the URI `http://www.manning.com/studyKit` with the TLD resource path `/myLibs/studyKit.tld`, while the second `<taglib>` element associates the URI `http://www.manning.com/sampleLib` with the TLD resource path `yourLibs/sample.jar`.

15.2.4 Resolving URIs to TLD file locations

Once the taglib mapping between the URIs and the TLD resource paths has been created, the JSP pages can refer to these URIs using the `taglib` directives:

```
<%@ taglib prefix="study" uri="http://www.manning.com/studyKit" %>
<%@ taglib prefix="sample" uri="http://www.manning.com/sampleLib" %>
```

When the JSP engine parses a JSP file and encounters a `taglib` directive, it checks its taglib map to see if a mapping exists for the `uri` attribute of the `taglib` directive:

- If the value of the `uri` attribute matches any of the `<taglib-uri>` entries, the engine uses the value of the corresponding `<taglib-location>` to locate the actual TLD file.
 - If the `<taglib-location>` value is a root-relative URI (that is, it starts with a `/`), the JSP engine assumes the location to be relative to the web application's document root directory. Thus, the location for the URI `http://www.manning.com/studyKit` in listing 15.1 will resolve to the TLD file `<doc-root>/myLibs/studyKit.tld`.
 - If the `<taglib-location>` value is a nonroot relative URI (that is, it starts without a `/`), the JSP engine prepends `/WEB-INF/` to the URI and assumes the location to be relative to the web application's document root directory. Thus, the location for the URI `http://www.manning.com/sampleLib` in listing 15.1 will resolve to the TLD path `<doc-root>/WEB-INF/yourLibs/sample.jar`.
- If the value of the `uri` attribute of the `taglib` directive does not match any of the `<taglib-uri>` entries, then the following three possibilities arise:
 - If the specified `uri` attribute is an absolute URI, then it is an error and is reported at translation time.
 - If the specified `uri` attribute is a root-relative URI, it is assumed to be relative to the web application's document root directory.
 - If the specified `uri` attribute is a nonroot relative URI, it is assumed to be relative to the current JSP page. Thus, if the JSP file `<doc-root>/jsp/test.jsp` contains the directive `<%@ taglib prefix="test" uri="sample.tld" %>`, the engine will expect to find the `sample.tld` file at `<doc-root>/jsp/sample.tld`.

Quizlet

Q: Consider the following `taglib` directive appearing in a JSP file. What will happen if the URI used in this directive is not mapped to a TLD file in the deployment descriptor?

```
<%@ taglib uri="www.manning.com/hello.tld" prefix="a" %>
```

A: The URI `www.manning.com/hello.tld` does not contain a protocol; therefore, it is not an absolute URI. It does not start with a `/`, so it is not a root-relative URI either. It is a page-relative URI. After failing to find an entry in the map, the engine will search for the file `hello.tld` at the location relative to the current page. Suppose the JSP page is at location

```
C:\tomcat\webapps\chapter15\test.jsp
```

The engine will look for the file at

```
C:\tomcat\webapps\chapter15\www.manning.com\hello.tld
```

If it is unable to find `hello.tld` at this location either, it will flag an error.

Q: What is wrong with the following `taglib` declaration?

```
<taglib>
  <taglib-uri>http://myLibs.com</taglib-uri>
  <taglib-location>http://yourLibs.com</taglib-location>
</taglib>
```

A: The value of `<taglib-location>` cannot be an absolute URI.

15.2.5 Understanding the prefix

As we discussed in chapter 11, “The JSP technology model—the basics,” each standard JSP action tag has a tag name that is made up of two parts, a prefix and an action, separated by a colon. For example, in `<jsp:include>` and `<jsp:forward>`, the prefix is `jsp`, while the actions are `include` and `forward`, respectively. Custom tags use the same syntax:

```
<myPrefix:myCustomAction>
```

Since we can use multiple libraries in one JSP page, the prefix differentiates between tags that belong to different libraries. For example:

```
<%@ taglib prefix="compA" uri="mathLibFromCompanyA" %>
<%@ taglib prefix="compB" uri="mathLibFromCompanyB" %>

<!-- Uses a tag from Company A -->
<compA:random/>

<!-- Uses a tag from Company B-->
<compB:random/>
```

In the above code snippet, the prefixes `compA` and `compB` enable the JSP engine to identify the libraries to which the tags belong.

In order to avoid conflicts between the user-defined tags and the standard tags provided by the JSP implementations, there are certain restrictions on the value of the `prefix` attribute. For example, we cannot use `jsp` as a prefix for custom tag libraries because it is already used as a prefix for standard actions, such as `<jsp:include>`, `<jsp:forward>`, and `<jsp:useBean>`.

In addition to `jsp`, the specification has reserved the following prefix values, which means we cannot use them in a `taglib` directive: `jspx`, `java`, `javax`, `servlet`, `sun`, and `sunw`. Thus, the following directive is invalid:

```
<%@ taglib prefix="sun" uri="myLib" %>
```

15.3 USING CUSTOM TAGS IN JSP PAGES

In section 15.2, we learned how to import custom tags using a `taglib` directive. Now, we will see how to use different types of custom tags in a JSP page. These types include:

- Empty tags
- Tags with attributes
- Tags with JSP code
- Tags with nested tags

To illustrate the usage of these tag types, we have used the tags of a sample tag library in the following sections. We will learn how to build these tags in the next chapter.

15.3.1 Empty tags

Empty tags do not have any body content. They are written in two ways. They can consist of a pair of opening and closing tags without anything in between:

```
<prefix:tagName></prefix:tagName>
```

They can also be formatted as a single self-tag:

```
<prefix:tagName />
```

The self-tag has a forward slash / at the end of the tag.

In listing 15.2, we have an empty tag named `required` that embeds the character * in the output HTML. This tag is useful while accepting <FORM> input from users.

Listing 15.2 Usage of an empty tag

```
<%@ taglib uri="sampleLib.tld" prefix="test" %>

<html>
Please enter your address and click submit.<br>
The fields marked with a <test:required /> are mandatory.

<form action="validateAddress.jsp">
<table>

<tr>
<td><test:required /> Street 1</td>
<td><input TYPE='text' NAME='street1'></td>
</tr>

<tr>
<td> Street 2</td>
<td><input TYPE='text' NAME='street2'></td>
</tr>

<tr>
<td> Street 3</td>
<td><input TYPE='text' NAME='street3'></td>
</tr>

<tr>
<td><test:required/> City </td>
<td><input TYPE='text' NAME='city'></td>
</tr>
```

```

<tr>
  <td><test:required/> State </td>
  <td><input TYPE='text' NAME='state'></td>
</tr>

<tr>
  <td><test:required /> Zip </td>
  <td><input TYPE='text' NAME='zip'></td>
</tr>

</table>

<input TYPE='submit' >

</form>
</html>

```

Figure 15.1 shows the output of listing 15.2 in a browser.

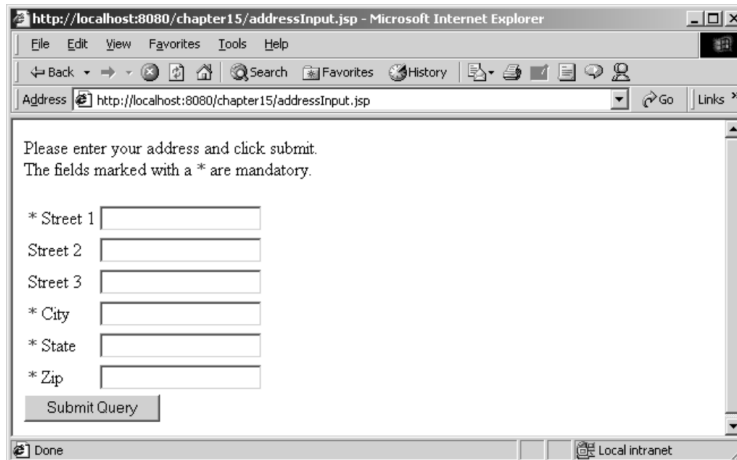


Figure 15.1
A JSP page using
the required tag

Although simple, this tag is quite useful. The page author does not have to scatter `*` everywhere in the JSP page. In addition, if you decide later on that you want to use a different color for the `*` or you want to use an image instead of `*`, you can just modify the tag handler class and the change will be reflected application-wide without modifying any JSP page.

15.3.2 Tags with attributes

Just like standard tags, custom tags can have attributes. In our sample library, we have a tag named `greet` that prints the greeting `Hello` in the output. It accepts an attribute named `user` to print the user's name:

```

<html><body>

  <%@ taglib prefix="test" uri="sampleLib.tld" %>

```

```

    <h3><test:greet user="john" /></h3>
</body></html>

```

When used this way, it will print the user's name in the greeting as Hello john! in the browser.

In the same way that standard tags may have attributes that are mandatory, attributes for custom tags also may be defined as mandatory. If we do not specify the mandatory attributes, the JSP engine flags an error at translation time. On the other hand, if we do not specify the nonmandatory attributes, the tag handler uses the default values. The default values depend on the implementation of the tag handler. The attribute values can be either constants or JSP expressions:

```

<prefix:tagName attrib1="fixedValue"
    attrib2="<%= someJSPExpression %>"
    attrib3= ...
/>

```

The expressions are evaluated at request time and passed to the corresponding tag handler. Instead of passing a string literal john, we can use a JSP expression to make it more flexible:

```

<html><body>
    <%@ taglib prefix="test" uri="sampleLib.tld" %>
    <h3>
    <test:greet
        name='<%= request.getParameter("username") %>'
    />
    </h3>
</body></html>

```

Thus, attributes to a tag are like parameters to a method call. The tag designer can customize the behavior of a tag by specifying attributes and values. However, we cannot pass in any arbitrary attribute-value pairs. As we will see in the next chapter, the tag library designer defines the following in a TLD file:

- A set of valid attributes names
- Whether or not an attribute is mandatory
- The data type of the values
- Whether the value of an attribute has to be specified at translation time using a string literal or if it can be specified as a request-time expression

15.3.3 Tags with JSP code

A tag may contain JSP code enclosed within the opening and closing tags. This code is called the *body content* of the tag. It can be any valid JSP code, which means it can be text, HTML, a scriptlet, an expression, and so forth. Our sample library has a tag named `if` that accepts a Boolean attribute. It either includes the body in the output or skips the body altogether based on the value of the attribute passed:

```

<html><body>
  <%@ taglib uri="sampletaglib.tld" prefix="test" %>

  <test:if condition="true">
    Anything that is to be printed when the condition is true goes here.
    Name is: <%= request.getParameter("name") %>
  </test:if>

</body></html>

```

In the above code, the `<test:if>` tag is passed a value of `true` for the attribute `condition`. It executes the body of the tag, and the contents are included in the output. If we set the value of the `condition` attribute to `false`, it will skip the body and the contents will not be included in the output.

15.3.4 Tags with nested custom tags

Nonempty tags can also include other custom tags in their body. These types of tags are called *nested tags*. In our sample library, we have three tags—`<switch>`, `<case>`, and `<default>`—that help us write switch-case statements in JSP pages:

```

<html><body>
  <%@ taglib uri="sampleLib.tld" prefix="test" %>

  <test:switch conditionValue='<%= request.getParameter("action") %>' >

    <test:case caseValue="sayHello">
      Hello!
    </test:case>

    <test:case caseValue="sayGoodBye" >
      Good Bye!!
    </test:case>

    <test:default>
      I am Dumb!!!
    </test:default>

  </test:switch>

</body></html>

```

In the above code, the `<case>` and `<default>` tags are nested inside the `<switch>` tag. Depending on the value of the `conditionValue` attribute, the body of an appropriate case tag executes.

Note that the opening and closing tags of a nested tag and its enclosing tag cannot overlap. The following is syntactically incorrect:

```

<test:tag1>
  <test:tag2>
</test:tag1>
  </test:tag2>

```

15.4 USING THE JSP STANDARD TAG LIBRARY (JSTL)

Now that you're familiar with the process of incorporating tag libraries into your JSPs, we can present a freely available tag library as an example. This section will cover Sun's latest version (1.1) of its JSTL. This library consists of a number of sublibraries, each of which provides tags for a specific group of functions. Specifically, these sublibraries include

- *core*—Tags for general purpose processing
- *xml*—Tags for parsing, selecting, and transforming XML data
- *fmt*—Tags for formatting data for international use
- *sql*—Tags for accessing relational databases
- *functions*—Tags for manipulating Strings and collections

All of these are useful, but the SCWCD exam only requires that you understand the tags in the core library. The tags in this library are both very useful and easy to integrate into JSPs. Once again, the new SCWCD exam includes this material to provide an alternative to adding scripts (declarations, expressions, and scriptlets) to your pages.

But before we investigate what tags are included in the core library, we need to show you how to add the JSTL library within your Tomcat installation. This won't be tested, but it will enable you to code and execute the examples in this section.

15.4.1 Acquiring and installing the JSTL

There are two JAR files that provide JSTL capabilities to JSPs. The first, `jstl.jar`, provides the API classes for the library. The second, `standard.jar`, provides the library's implementation classes. Tomcat 5.0 holds both, but they're hiding in the examples. To acquire them, copy the two libraries from

```
C:\jakarta-tomcat-5.0.25\webapps\jsp-examples\WEB-INF\lib
```

and add them to the `lib` directory within your application's `WEB-INF` folder.

Because you've added these files to the `lib` directory, you don't need to update your deployment descriptor. The container will find them automatically. But you do need to reference the library in your JSP by using the `taglib` directive. Here's an example:

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
```

This will allow you to reference JSTL tags of the core library by using the prefix `c`.

Now that you know how to incorporate the core library into your page and access its tags, we can describe what they do. Sun classifies them into four categories, and table 15.2 lists them and their tags.

We'll start with by describing the tags in the first category.

Table 15.2 JSTL tags categorized by function

JSTL tag category	JSTL tags	Tag descriptions
General purpose	<code><c:catch></code>	Catches exceptions within a variable
	<code><c:out></code>	Displays contents within the page
Variable support	<code><c:set></code>	Sets the value of an EL variable
	<code><c:remove></code>	Removes an EL variable
Flow control	<code><c:if></code>	Alters processing according to an attribute equaling a value
	<code><c:choose></code>	Alters processing according to an attribute equaling a set of values
	<code><c:forEach></code>	Repeats processing for each object in a collection
	<code><c:forTokens></code>	Performs processing for each substring in a given text field
URL handling	<code><c:url></code>	Rewrites URLs and encodes their parameters
	<code><c:import></code>	Accesses content outside the web application
	<code><c:redirect></code>	Tells the client browser to access a different URL

15.4.2 General purpose JSTL tags: `<c:catch>` and `<c:out>`

The first two core tags, `<c:catch>` and `<c:out>`, make it possible to perform Java processing without using JSP scripts. The first tag allows you to catch processing errors within the JSP instead of bringing up an error page. The second functions similarly to the JSP expression script, and makes its contents available for display.

Normally, any exceptions thrown within a JSP will be sent to the error page. However, you may want to perform different error-handling routines for different actions within the page. The `<c:catch>` tag won't perform these routines by itself, but it will store the thrown exception within a variable named by the `var` attribute. Here's an example:

```
<c:catch var="e">
    actions that might throw an exception
</c:catch>
```

If the code between the tags throws an exception, its value will be stored in a variable named `e`, which has page scope. Otherwise, `e` will be removed if it existed.

The second tag in this category, `<c:out>`, is even simpler to use. It functions just like the JSP script expression, represented by `<%=` and `/>`. It has one required attribute, `value`, and the tag functions by displaying the value of `value`. An example is shown here:

```
<c:out value="\${number}" />
```

This example will display the value of the `number` expression within the page.

This tag is simple to use and understand, but you should remember two points. First, if the `value` attribute contains a `<`, `>`, `'`, `"`, or `&` character, the tag will display

the corresponding character code, such as `>` for `>`. Second, you can specify a default value if the attribute's variable hasn't been initialized. This is done with a second attribute called `default`, which isn't required. This is shown in the following line:

```
<c:out value="\${color}" default="red" />
```

Just as the `<c:out>` tag can be used in place of JSP script expressions, we need tags that replace declarations. For this, the core library provides variable support tags.

15.4.3 Variable support JSTL tags: `<c:set>` and `<c:remove>`

Although the EL can manipulate variables in a number of ways, it can't set their value or remove them from scope. But with the core library's `<c:set>` and `<c:remove>` tags, you can perform these operations without resorting to JSP scripts.

The first tag in this category, `<c:set>`, can be used to set the values of both variables and objects, such as `JavaBeans` and `Map` instances. Also, the values can be specified from within the tag or inside its body content.

To use `<c:set>` with a variable, you need to specify the variable's name with the `var` attribute. Then, you can specify its value either with a `value` attribute or use the tag's body content. For example, the two tags

```
<c:set var="num" value="\${4*4}" />
```

and

```
<c:set var="num">
  \${8*2}
</c:set>
```

both set the value of the `num` variable to 16. The second method allows you to insert tag operations within the `<c:set>` tags. For example,

```
<c:set var="num">
  <c:out value="\${8+8}" />
</c:set>
```

will accomplish the same result as the two preceding examples.

Along with setting the values of variables, the `<c:set>` tags also allow you to work with `JavaBeans` and `java.util.Map` objects. The only difference is that the two methods use different attributes. When setting `JavaBeans` and `Maps`, you specify the object's name with the `target` attribute and its property (the `JavaBean`'s member variable or the `Map`'s key) with the `property` attribute. To set the property's value, you can use the `value` attribute or the body content, just as you can with setting variables.

For example, if you want to set the `zipcode` property of a `JavaBean` object called `customer1`, you can use

```
<c:set target="customer1" property="zipcode">
  55501
</c:set>
```

or

```
<c:set target="customer1" property="zipcode" value="55501">
```

The tag usage for Map objects is similar, except that the `property` attribute reflects the name of one of the Map's entries.

The `<c:remove>` tag is used to remove a variable from its scope. You specify the variable's name with the required `var` attribute, and its scope with the optional `scope` attribute. If `scope` isn't specified, then the container will look first at the page, then the request, then the session, and finally the application scope.

As a simple example, you can remove the `num` variable from the session scope with the following tag:

```
<c:remove var="num" scope="session" />
```

Unlike `<c:set>`, the `<c:remove>` tag can't be used with JavaBeans or Map objects.

15.4.4 Flow control JSTL: `<c:if>`, `<c:choose>`, `<c:forEach>`, and `<c:forTokens>`

Most Java classes incorporate some form of flow control to change processing according to a variable's value. Using `for` and `while` statements, you can control how many times a task should be repeated. With `if` and `switch-case` statements, you can control which task should be performed. Before JSP 2.0, the only way to control page processing was to use scriptlets. But JSTL provides four tags—`<c:if>`, `<c:choose>`, `<c:forEach>`, and `<c:forTokens>`—that provide flow control without JSP scripts.

Conditional processing with JSTL

The first two tags, `<c:if>` and `<c:choose>`, function very similarly to regular Java's `if` and `switch-case` statements. One important difference is that there is nothing like an `else` tag in the core library. Another is that each of these tags requires an attribute called `test`, which you need to set equal to a boolean expression. For example, the following code will only display the value of `x` if it's equal to 9:

```
<c:if test="\${x == '9'}">
  ${x}
</c:if>
```

It's important to notice the single quotes surrounding the variable's test value. These must be present for the container to recognize the comparison.

The `<c:choose>` tag doesn't take any attributes by itself, but contains a number of `<c:when>` tags that perform separate comparisons with their `test` attributes. For example, the following code displays different text based on the `color` variable:

```
<c:choose>
  <c:when test="\${color == 'white'}">
    Light!
  </c:when>
  <c:when test="\${color == 'black'}">
```



```

    Dark!
  </c:when>
  <c:otherwise>
    Colors!
  </c:otherwise>
</c:choose>

```

Just as Java's `switch` statement can contain a default entry when none of its other conditions are met, JSTL provides an optional `<c:otherwise>` tag as the default option.

JSTL iteration

The core library's `<forEach>` and `<forEachTokens>` tags allow you to repeat processing of the tag's body content. Using these tags, you can control the number of iterations in three ways:

- with a range of numbers: using `<forEach>` and its `begin`, `end`, and `step` attributes
- with the elements in a Java collection: using `<forEach>` and its `items` attribute
- with the tokens in a `String`: using `<forEach>` and its `items` attribute

The first method of iteration is very straightforward, and works like the traditional Java `for` loop. First, the `<forEach>` tag creates a variable specified by its `var` attribute. Then, it initializes the variable to its `begin` value and continues processing body content until the variable equals `end`. The following JSP counts from 0 to 30 and displays every third number:

```

<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<html><body>
  <c:forEach var="x" begin="0" end="30" step="3">
    ${x}
  </c:forEach>
</body></html>

```

The `<forEach>` tag can also be used to iterate over the elements of a Java collection object, such as a `Vector`, `List`, or `Map`. This tag processes its body content once for each element, and you can access these elements by using the `var` attribute. The following code cycles through the elements of `numArray` and sets their values to 100:

```

<c:forEach var="num" items="{numArray}">
  <c:set var="num" value="100" />
</c:forEach>

```

In the `<forEachTokens>` tag, the `items` attribute is a `String` made up of tokens separated by delimiters. If you think of this `String` as a collection of substrings, then you'll see how closely it resembles the previous `<c:forEach>` usage. For example, `Tokens.jsp` in code listing 15.3 creates a table and populates its cells with the tokens of `numlist`.

Listing 15.3 Tokens.jsp

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<html><body>
  <c:set var="numList" value="one,two,three,four,five,six" />
  Output of the forTokens tag:<p>
  <table border="1">
    <c:forTokens var="num" items="\${numList}" delims=",">
      <tr><td>\${num}</td></tr>
    </c:forTokens>
  </table>
</body></html>
```

The result is shown in figure 15.2.

In this code, the `delims` attribute tells the container to tokenize `items` by commas. Then, each token is placed in its own table row and cell. Again, it's important to remember the basics of HTML for the SCWCD exam.

Accessing URL information with JSTL

The last category of tags in the core library deals with URL accessing. The three main tags are

- `<c:url>`, which rewrites URLs and encodes their parameters
- `<c:import>`, which accesses content outside the web application, and
- `<c:redirect>`, which tells the client browser to access a different URL

If the client's browser doesn't accept cookies, then you need to rewrite the URL to maintain session state. The core library provides `<c:url>` for this purpose. You specify the base URL with the `value` attribute, and the transformed URL is displayed by the `JspWriter` or saved to a variable named by the optional `var` attribute.

A simple example is

```
<c:url value="/page.html" var="pagename"/>
```

You can also specify the scope of `var` with an optional `scope` attribute, whose values include `page`, `request`, `session`, and `application`.

Since the `value` parameter starts with a forward slash, the tag inserts the name of the context before the URL. For example, the value of `var` in the preceding line of code would be

```
/contextname/page.html
```

if the browser accepts cookies. If the container doesn't find a cookie with the current session, it will rewrite the URL with its session ID number. In this case, the result would be




Figure 15.2 A JSP showing the result of the `<c:forTokens>` tag

```
/contextname/page.html;jsessionid=jsessionid
```

You can also specify the context with the optional `context` attribute. If value doesn't start with a forward slash, no context name will be added.

You can also add parameters to the URL by using the `<c:param>` tag in the body of `<c:url>`. The following code shows how this is done:

```
<c:url value="/page.html" var="pagename">
  <c:param name="param1" value="{2*2}"/>
  <c:param name="param2" value="{3*3}"/>
</c:url>
```

The parameters within the `<c:param>` tag are specified by their name and value attributes. If the browser accepts cookies, the value of the `var` attribute would be

```
/contextname/page.html?param1=4&param2=9
```

The tag processing will encode the URL and its parameters as needed.

You can access URLs with the `include` directive, but if you want to incorporate content located outside the servlet container, you need the `<c:import>` tag. This tag adds content to the JSP referenced by the `url` attribute. The URL can be relative or absolute, but if it begins with a forward slash, it is treated as absolute within the application. You can set the URL's context with the `context` attribute.

As with `<c:url>`, you can save the imported content within a variable specified by the `var` attribute. You can also set the variable's scope with the `scope` attribute, or control its encoding with `charEncoding`. Finally, just as with the `<c:url>` tag, you can add parameters to the URL with `<c:param>` tags in the `<c:import>` body.

The following example creates a variable called `newstuff` and sets it equal to the content of the `content.html` URL. In accessing the URL, the `<c:import>` tag appends two parameters, `par1` and `par2`.

```
<c:import url="/content.html" var="newstuff" scope="session">
  <c:param name="par1" value="val1"/>
  <c:param name="par2" value="val2"/>
</c:import>
```

The last (and simplest) tag in this category is `<c:redirect>`. This tag functions identically to the `HttpServletResponse`'s `sendRedirect()` method. It sends a redirect response to the client and tells it to access the URL specified by the `url` attribute. As with the `<c:url>` and `<c:import>` tags, you can specify the URL's context with `context` and add parameters with `<c:param>` tags.

The code snippet that follows provides an example of how the `<c:redirect>` tag directs processing to a new URL:

```
<c:redirect url="/content.html">
  <c:param name="par1" value="val1"/>
  <c:param name="par2" value="val2"/>
</c:redirect>
```

The JSTL library provides many more tags than those in the core library, but these should be sufficient to perform most of the processing tasks normally accomplished by JSP scripts. With these tags, you can perform flow control, manipulate and access URLs, and set and display values of JSP variables. You should keep these tags and the JSTL in mind, both for the SCWCD exam and for general JSP development.

15.5 SUMMARY

Custom tags are action elements on JSP pages that are mapped to tag handler classes in a tag library. Tag libraries allow us to use independent Java classes to manage the presentation logic of the JSP pages, thereby reducing the use of scriptlets and leveraging existing code to accelerate development time. In this chapter, we learned the basic terms and usage of tag libraries, including how to explicitly associate a URI with a TLD file using the `<taglib>` element in the deployment descriptor.

We discussed the use of the `taglib` directive to import custom tags into a JSP page. Then we examined several different types of custom tags that we use in JSP pages: empty tags, tags with attributes, tags with a body, and tags that contain nested tags.

You should now be able to answer exam questions about the declaration of a tag library in the deployment descriptor, the various ways of importing a custom tag library for use in JSP pages, and the use of different types of custom tags in JSP pages.

In the next chapter, we will learn more about the TLD file and how to implement our own custom tags.

15.6 REVIEW QUESTIONS

- Which of the following elements are required for a valid `<taglib>` element in `web.xml`? (Select two)
 - `uri`
 - `taglib-uri`
 - `tagliburi`
 - `tag-uri`
 - `location`
 - `taglib-location`
 - `tag-location`
 - `tagliblocation`
- Which of the following `web.xml` snippets correctly defines the use of a tag library? (Select one)
 - ```
<taglib>
 <uri>http://www.abc.com/sample.tld</uri>
 <location>/WEB-INF/sample.tld</location>
</taglib>
```

- b** `<tag-lib>`  
`<taglib-uri>http://www.abc.com/sample.tld</taglib-uri>`  
`<taglib-location>/WEB-INF/sample.tld</taglib-location>`  
`</tag-lib>`
- c** `<taglib>`  
`<taglib-uri>http://www.abc.com/sample.tld</taglib-uri>`  
`<taglib-location>/WEB-INF/sample.tld</taglib-location>`  
`</taglib>`
- d** `<tag-lib>`  
`<taglib>http://www.abc.com/sample.tld</taglib-uri>`  
`<taglib>/WEB-INF/sample.tld</taglib-location>`  
`</tag-lib>`

3. Which of the following is a valid `taglib` directive? (Select one)

- a** `<% taglib uri="/stats" prefix="stats" %>`
- b** `<%@ taglib uri="/stats" prefix="stats" %>`
- c** `<%! taglib uri="/stats" prefix="stats" %>`
- d** `<%@ taglib name="/stats" prefix="stats" %>`
- e** `<%@ taglib name="/stats" value="stats" %>`

4. Which of the following is a valid `taglib` directive? (Select one)

- a** `<%@ taglib prefix="java" uri="sunlib"%>`
- b** `<%@ taglib prefix="jspx" uri="sunlib"%>`
- c** `<%@ taglib prefix="jsp" uri="sunlib"%>`
- d** `<%@ taglib prefix="servlet" uri="sunlib"%>`
- e** `<%@ taglib prefix="sunw" uri="sunlib"%>`
- f** `<%@ taglib prefix="suned" uri="sunlib"%>`

5. Consider the following `<taglib>` element, which appears in a deployment descriptor of a web application:

```
<taglib>
 <taglib-uri>/accounting</taglib-uri>
 <taglib-location>/WEB-INF/tlds/SmartAccount.tld</taglib-location>
</taglib>
```

Which of the following correctly specifies the use of the above tag library in a JSP page? (Select one)

- a** `<%@ taglib uri="/accounting" prefix="acc"%>`
- b** `<%@ taglib uri="/acc" prefix="/accounting"%>`
- c** `<%@ taglib name="/accounting" prefix="acc"%>`
- d** `<%@ taglib library="/accounting" prefix="acc"%>`
- e** `<%@ taglib name="/acc" prefix="/accounting"%>`

6. You are given a tag library that has a tag named `printReport`. This tag may accept an attribute, `department`, which cannot take a dynamic value. Which of the following are correct uses of this tag? (Select two)

**a** `<mylib:printReport/>`  
**b** `<mylib:printReport department="finance"/>`  
**c** `<mylib:printReport attribute="department" value="finance"/>`  
**d** `<mylib:printReport attribute="department"  
attribute-value="finance"/>`  
**e** `<mylib:printReport>  
<jsp:attribute name="department" value="finance" />  
</mylib:printReport>`

7. You are given a tag library that has a tag named `getMenu`, which requires an attribute, `subject`. This attribute can take a dynamic value. Which of the following are correct uses of this tag? (Select two)

**a** `<mylib:getMenu />`  
**b** `<mylib:getMenu subject="finance"/>`  
**c** `<% String subject="HR";%>  
<mylib:getMenu subject="<%=subject%"/>`  
**d** `<mylib:getMenu> <jsp:param subject="finance"/> </mylib:getMenu>`  
**e** `<mylib:getMenu>  
<jsp:param name="subject" value="finance"/>  
</mylib:getMenu>`

8. Which of the following is a correct way to nest one custom tag inside another? (Select one)

**a** `<greet:hello>  
<greet:world>  
</greet:hello>  
</greet:world>`  
**b** `<greet:hello>  
<greet:world>  
</greet:world>  
</greet:hello>`  
**c** `<greet:hello  
<greet:world/>  
>`  
**d** `<greet:hello>  
</greet:hello>  
<greet:world>  
</greet:world>`

9. Which of the following elements can you use to import a tag library in a JSP document? (Select one)
- a `<jsp:root>`
  - b `<jsp:taglib>`
  - c `<jsp:directive.taglib>`
  - d `<jsp:taglib.directive>`
  - e We cannot use custom tag libraries in XML format.
10. Using `c` to represent the JSTL library, which of the following produces the same result as `<%= var %>`? (Select one)
- a `<c:set value=var>`
  - b `<c:var out=${var}>`
  - c `<c:out value=${var}>`
  - d `<c:out var="var">`
  - e `<c:expr value=var>`
11. Which attribute of `<c:if>` specifies the conditional expression? (Select one)
- a `cond`
  - b `value`
  - c `check`
  - d `expr`
  - e `test`
12. Which of the following JSTL `forEach` tags is valid?
- a `<c:forEach varName="count" begin="1" end="10" step="1">`
  - b `<c:forEach var="count" begin="1" end="10" step="1">`
  - c `<c:forEach test="count" beg="1" end="10" step="1">`
  - d `<c:forEach varName="count" val="1" end="10" inc="1">`
  - e `<c:forEach var="count" start="1" end="10" step="1">`
13. Which tags can be found in a JSTL `choose`? (Select two)
- a `case`
  - b `choose`
  - c `check`
  - d `when`
  - e `otherwise`



## CHAPTER 16

---

# *Developing “Classic” custom tag libraries*

- 16.1 Understanding the tag library descriptor 310
- 16.2 The Tag Extension API 318
- 16.3 Implementing the Tag interface 320
- 16.4 Implementing the IterationTag interface 329
- 16.5 Implementing the BodyTag interface 333
- 16.6 Extending TagSupport and BodyTagSupport 338
- 16.7 What’s more? 347
- 16.8 Summary 348
- 16.9 Review questions 349

### **EXAM OBJECTIVES**

- 10.1** Describe the semantics of the “Classic” custom tag event model when each event method (doStartTag, doAfterBody, and doEndTag) is executed, and explain what the return value for each event method means; and write a tag handler class. (Sections 16.1 through 16.5)
- 10.2** Using the PageContext API, write tag handler code to access the JSP implicit variables and access web application attributes. (Section 16.3)
- 10.3** Given a scenario, write tag handler code to access the parent tag and an arbitrary tag ancestor. (Section 16.3)



## INTRODUCTION

In chapter 15, “Using custom tags,” we introduced JSP custom tags and tag libraries. We discussed the way that we use custom tags in JSP pages and the process of importing an existing tag library into a JSP page using the `taglib` directive. In this chapter, we will learn how to develop our own custom libraries according to the “Classic” model of tag library development.

## 16.1 UNDERSTANDING THE TAG LIBRARY DESCRIPTOR

The tag library descriptor (TLD) file contains the information that the JSP engine needs to know about the tag library in order to interpret the custom tags on a JSP page. Let’s take a close look at the elements of the TLD and how they are used to describe a tag library.

A *tag library descriptor* is an XML document that follows the DTD designated by the JSP specification so that it can be created, read, and understood by all kinds of users, including human users and the JSP engine, as well as other development tools. In essence, it informs the user of a tag library about the usage and behavior of the tags that the library provides.

On the exam, you may be asked to identify the correct format and usage of the different elements in a TLD file. In addition, given a properly formatted TLD file, you may be asked to identify the correct usage of the corresponding tags and their attributes in JSP pages.

Listing 16.1 is an example of a TLD. We have bolded four elements—`<taglib>`, `<tag>`, `<body-content>`, and `<attribute>`—which you need to be familiar with to do well on the exam. We will discuss all of these elements in the following sections.

### Listing 16.1 sampleLib.tld: A sample tag library descriptor

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE taglib PUBLIC
 "-//Sun Microsystems, Inc.//DTD JSP Tag Library 1.2//EN"
 "http://java.sun.com/dtd/web-jsptaglibrary_1_2.dtd" >
<taglib>
 <tlib-version>1.0</tlib-version>
 <jsp-version>2.0</jsp-version>
 <short-name>test</short-name>
 <uri>http://www.manning.com/sampleLib</uri>
 <tag>
 <name>greet</name>
 <tag-class>sampleLib.GreetTag</tag-class>
 <body-content>empty</body-content>
 <description>Prints Hello and the user name</description>
```

```

 <attribute>
 <name>user</name>
 <required>>false</required>
 <rtexprvalue>>true</rtexprvalue>
 </attribute>
 </tag>
</taglib>

```

This code demonstrates three important things:

- First, a tag library descriptor file, like all XML files, starts with the line `<?xml version="1.0" encoding="ISO-8859-1">`, which specifies the version of XML and the character set that the file is using.
- Second, it has a DOCTYPE declaration that identifies the DTD for this document. In the case of a tag library descriptor that conforms to the JSP 1.2 specification, the DOCTYPE declaration must be

```

<!DOCTYPE taglib PUBLIC
"-//Sun Microsystems, Inc.//DTD JSP Tag Library 1.2//EN"
"http://java.sun.com/dtd/web-jsptaglibrary_1_2.dtd" >

```

- Finally, it shows that all the contents of the TLD come under the `<taglib>` element. In other words, the `<taglib>` element is the root element of a TLD.

**NOTE** Unlike the deployment descriptor (`web.xml`), the extension of a TLD file is typically `.tld` and not `.xml`.

For all of the examples in this chapter, assume that the name of the TLD file is `sampleLib.tld`, that it resides in the `WEB-INF` directory of the web application named `chapter16`, and that the JSP code snippets that use the custom tags follow the `taglib` directive declared as

```

<%@ taglib prefix="test" uri="/WEB-INF/sampleLib.tld" %>

```

Therefore, all of the tags in these examples will be qualified with the prefix `test`.

### 16.1.1 The `<taglib>` element

The `<taglib>` element is the top-level, or root, element of the tag library descriptor. It contains other second-level elements that provide information about the library as a whole, such as the version of the library and the version of the JSP specification the library conforms to. The following is the DTD for the `<taglib>` element:

```

<!ELEMENT taglib (tlib-version, jsp-version, short-name,
uri?, display-name?, small-icon?, large-icon?,
description?, validator?, listener*, tag+) >

```

As we can see from this DTD, the three subelements—`<tlib-version>`, `<jsp-version>`, and `<short-name>`—are mandatory and appear exactly once. There

must be at least one `<tag>` subelement and zero or more `<listener>` elements. All other subelements are optional and can occur no more than once. Table 16.1 briefly describes their use.

**Table 16.1** The subelements of `<taglib>`

Element	Description	Occurrence
<code>tlib-version</code>	Specifies the version of the tag library.	Exactly once
<code>jsp-version</code>	Specifies the version of JSP that this tag library depends on. For example, a value of 2.0 informs the JSP container that the implementation classes are using features of the JSP Specification 2.0 and that the container must be 2.0 compatible to be able to use this library.	Exactly once
<code>short-name</code>	Specifies a preferred prefix value for the tags in the library. Usually used by page authoring tools.	Exactly once
<code>uri</code>	A URI for identifying this tag library. This is the implicit way of adding <code>&lt;taglib-uri&gt;</code> and <code>&lt;taglib-location&gt;</code> pair entries into the taglib map. The value of this element is used as the URI for this library and the actual physical location of this TLD file is used as the location of the library.	At most once
<code>display-name</code>	A short name that can be displayed by page authoring tools.	At most once
<code>small-icon</code>	A small icon that can be used by tools.	At most once
<code>large-icon</code>	A large icon that can be used by tools.	At most once
<code>description</code>	Any text describing the use of this taglib.	At most once
<code>validator</code>	Information about this library's <code>TagLibraryValidator</code> .	At most once
<code>listener</code>	Specifies event listener classes. We saw in Chapter 6 that we can specify listeners such as <code>HttpSessionListener</code> or <code>ServletContextListener</code> in the <code>web.xml</code> file. Similarly, if a tag library needs such listeners, we can specify the listener classes using this element. The container obtains these listener classes in exactly the way it obtains them from <code>web.xml</code> .	Any number of times
<code>tag</code>	Consists of subelements providing descriptions of a single tag. For multiple tags, we use multiple <code>&lt;tag&gt;</code> elements.	At least once

Notice that the file `sampleLib.tld` shown in listing 16.1 contains a `<uri>` element:

```
<uri>http://www.manning.com/sampleLib</uri>
```

This is the implicit way of adding a `<taglib-uri>` and `<taglib-location>` pair entry into the taglib map. If you recall the discussion of the taglib map in chapter 15, the JSP engine reads all of the `taglib.tld` files present in the packaged JAR files. For each jarred TLD file that contains information about its own URI, the JSP container automatically creates a mapping between the specified URI and the current location of the JAR file. In this case, if the `sampleLib.tld` file is renamed as `taglib.tld` and kept in the `META-INF` directory inside a JAR file, the above

<uri> element will cause the JSP container to create an implicit mapping of the URI `http://www.manning.com/sampleLib` with the actual location of the JAR file. We can then import the library into JSP pages using the `taglib` directive, as shown here:

```
<%@ taglib prefix="test" uri="http://www.manning.com/sampleLib" %>
```

However, if the deployment descriptor file contains an explicit mapping for the same URI, then the explicit mapping takes precedence over such implicit mappings.

## 16.1.2 The <tag> element

The <taglib> element may contain one or more <tag> elements. Each <tag> element provides information about a single tag, such as the tag's name, that will be used in the JSP pages, the tag's handler class, the tag's attributes, and so forth. The <tag> element is defined as follows:

```
<!ELEMENT tag (name, tag-class, tei-class?, body-content?,
 display-name?, small-icon?, large-icon?,
 description?, variable*, attribute*, example?) >
```

As we can see from this definition, the two subelements, <name> and <tag-class>, are mandatory and appear exactly once. There can be zero or more <variable> and <attribute> elements. All other subelements are optional and can occur at most once. Table 16.2 gives a brief description of each of the subelements.

**Table 16.2** The subelements of <tag>

Element	Description	Occurrence
name	The unique tag name.	Exactly once
tag-class	The tag handler class that implements <code>javax.servlet.jsp.tagext.Tag</code> .	Exactly once
tei-class	An optional subclass of <code>javax.servlet.jsp.tagext.TagExtraInfo</code> .	At most once
body-content	The content type for the body of the tag. Can be <i>empty</i> , <i>JSP</i> , or <i>tagdependent</i> . The default is <i>JSP</i> .	At most once
display-name	A short name that is intended to be displayed by development tools.	At most once
small-icon	A small icon that can be used by development tools.	At most once
large-icon	A large icon that can be used by development tools.	At most once
description	Specifies any tag-specific information.	At most once
variable	Specifies the scripting variable information.	Any number of times
attribute	Describes an attribute that this tag can accept.	Any number of times
example	Optional informal description of an example of using this tag.	At most once

The `<name>` element specifies the name of the tag that is to be used in the JSP pages, and the `<tag-class>` element specifies the fully qualified class name that implements the functionality of this tag. The class that we specify here must implement the `javax.servlet.jsp.tagext.Tag` interface.

We can define multiple tags that have different names and the same tag class. For example:

```
<tag>
 <name>greet</name>
 <tag-class>sampleLib.GreetTag</tag-class >
</tag>

<tag>
 <name>welcome</name>
 <tag-class>sampleLib.GreetTag</tag-class>
</tag>
```

In a JSP page, both of the tags, `<test:greet>` and `<test:welcome>`, will invoke the same handler class, `sampleLib.GreetTag` (assuming that the JSP page uses `test` as the prefix for this tag library).

However, we cannot define more than one tag with the same name, because the engine would not be able to resolve the tag handler class while de-referencing the tag name. Thus, the following is illegal:

```
<tag>
 <name>greet</name>
 <tag-class>sampleLib.GreetTag</tag-class >
</tag>

<tag>
 <name>greet</name>
 <tag-class>sampleLib.WelcomeTag</tag-class>
</tag>
```

### 16.1.3 The `<attribute>` element

The `<attribute>` element is a third-level element in a TLD and is a child of the `<tag>` element. If a custom tag accepts attributes, then the information about each attribute is specified using an `<attribute>` element. Each `<attribute>` element can have five subelements that provide the following information about the attribute:

- The attribute's name that will be used in the JSP pages
- The attribute's data type (`int`, `Boolean`, etc.)
- Whether or not the attribute is mandatory
- Whether or not the attribute can accept values at request time
- A brief description of the attribute

Here is the definition for the `<attribute>` element:

```
<!ELEMENT attribute (name, required?, rtexprvalue?,
 type?, description?) >
```

As we can see from the definition, only the `<name>` subelement is mandatory and must appear exactly once. All other subelements are optional and can occur no more than once. Table 16.3 describes each of the subelements.

**Table 16.3 The subelements of `<attribute>`**

Element	Description
<code>name</code>	The name of the attribute.
<code>required</code>	A value that specifies whether the attribute is required or optional. The default is <code>false</code> , which means optional. If this is set to <code>true</code> , then the JSP page must pass a value for this attribute. Possible values are <code>true</code> , <code>false</code> , <code>yes</code> , and <code>no</code> .
<code>rtexprvalue</code>	A value that specifies whether or not the attribute can accept request-time expression values. The default is <code>false</code> , which means it cannot accept request-time expression values. Possible values are <code>true</code> , <code>false</code> , <code>yes</code> , and <code>no</code> .
<code>type</code>	The data type of the attribute. This may be used only when <code>&lt;rtexprvalue&gt;</code> is set to <code>true</code> . It specifies the return type of the expression, using a request-time attribute expression: <code>&lt;%= %&gt;</code> . The default value is <code>java.lang.String</code> .
<code>description</code>	Some text describing the attribute for documentation purposes.

Consider a tag element that is defined as follows:

```
<tag>
 <name>greet</name>
 <tag-class>sampleLib.Greet</tag-class>
 <attribute>
 <name>user</name>
 <required>false</required>
 <rtexprvalue>true</rtexprvalue>
 </attribute>
</tag>
```

The above tag indicates that it accepts an attribute named `user`. Since the value of the `<required>` tag is `false`, a JSP page author may choose not to use the attribute-value pair. Further, since `<rtexprvalue>` is `true`, a JSP page author may use a request-time expression value. Therefore, the following lines from a JSP page are valid usages of this tag:

```
<test:greet /> ← Does not use user attribute Uses request-time
 expression
<test:greet user='<%= request.getParameter("user") %>' /> ←
```

In the previous tag definition, if we were to use a value of `true` for the `<required>` element, then we would have to specify the attribute-value pair for the given attribute in the JSP page. Thus, the first tag, `<test:greet />`, would generate a translation-time error because it does not specify a value for the `user` attribute. Moreover, if we were to use a value of `false` for the `<rtexprvalue>` element, then we must provide a value for the attribute in the JSP page that is not a request-time expression. Therefore, `<test:greet user="john" />` would be fine, but the second line in

the code, `<test:greet user="<%=...%>" />`, which uses a request-time attribute value, would generate a translation-time error.

### 16.1.4 The `<body-content>` element

The `<body-content>` element is a third-level element in a TLD and is a direct child of the `<tag>` element. This element does not have any subelements and can have one of three values:

- `empty`. Specifies that the body of the tag must be empty
- `JSP`. Specifies that the body of the tag can accept any normal JSP code
- `tagdependent`. Specifies that the content is not to be interpreted by the JSP engine and is tag dependent

Let's look at the details for each of these.

#### ***empty***

Some tags require a body, while others do not. In chapter 15 (section 15.3), we discussed the tags `<test:required>`, `<test:greet>`, and `<test:if>`. The `<test:required>` and `<test:greet>` tags did not have any body content. By their mere presence, they generated some output in the final HTML. On the other hand, the purpose of the `<test:if>` tag was to contain a set of statements that could either be skipped or executed as required.

A value of `empty` for the `<body-content>` element indicates that a tag does not support any body content. If the page author provides any content, the JSP engine flags an error at translation time. The following example declares the `<greet>` tag and specifies that its body content should be empty:

```
<tag>
 <name>greet</name>
 <tag-class>sampleLib.GreetTag</tag-class>
 <body-content>empty</body-content>

 <attribute>
 <name>user</name>
 <required>>false</required>
 <rtexprvalue>>true</rtexprvalue>
 </attribute>
</tag>
```

Therefore, the following are valid usages of the `<test:greet>` tag shown above:

```
<test:greet />
<test:greet user="john" />
<test:greet></test:greet>
<test:greet user="john"></test:greet>
```

The following usages of the tag are invalid since they contain some body content between the start and end tags:

```
<test:greet>john</test:greet>
<test:greet><%= "john" %></test:greet>
<test:greet> </test:greet> ← Space is not the same as empty
<test:greet>
</test:greet> ← New Line is not the same as empty
```

## JSP

A value of `JSP` for the `<body-content>` element indicates that the tag can have any valid JSP code in its body. This means that it can take plain text, HTML, scripting elements, standard actions, or even other custom tags nested inside this tag. The body could even be empty. The following example declares the `<if>` tag and specifies that its body can contain any kind of JSP content:

```
<tag>
 <name>if</name>
 <tag-class>sampleLib.IfTag</tag-class>
 <body-content>JSP</body-content>
 <attribute>
 <name>condition</name>
 <required>>true</required>
 <rtexprvalue>>true</rtexprvalue>
 </attribute>
</tag>
```

At request time, the nested scriptlets and expressions are executed and the actions and the custom tags are invoked as usual. Thus, the following are valid usages for the `<if>` tag declared above:

```
<test:if condition="true" />
<test:if condition="true"> </test:if>
<test:if condition="true"> </test:if>
<test:if condition="true">
 <test:greet user="john" />
 <% int x = 2+3; %>
 2+3 = <%= x %>
</test:if>
```

## tagdependent

A value of `tagdependent` for the `<body-content>` element indicates that the tag is expecting the body to contain text that may not be valid JSP code. The JSP engine does not attempt to execute the body and passes it as is to the tag handler at request time. It is up to the tag handler class to process the body content as needed. This value



for the body content element is required if we want to introduce code snippets from other languages. For example, we can develop a tag that executes SQL statements and inserts the result set into the output:

```
<test:dbQuery>
 SELECT count(*) FROM USERS
</test:dbQuery>
```

The tag handler class of the `dbQuery` tag would handle everything regarding the database, such as opening a connection, firing an SQL query, and so forth. It will only need the actual SQL query string that is specified as the body of the above tag. For such a tag, the `<body-content>` element must be specified as `tagdependent`.

**NOTE** Since a TLD is an XML document, the following rules apply:

- The order of the different elements and subelements is important. For example, the `<body-content>` element must appear before the `<attribute>` element under the `<tag>` element.
- The tag names are case sensitive. Thus, `<Attribute>` is not a valid element; we must use `<attribute>`.
- The character - appearing in many of the elements is important. Thus, `<bodycontent>` and `<tagclass>` are both valid in JSP 1.1 but invalid in JSP 1.2. Questions on version-specific issues may not appear on the exam, but it is good to know these points because while practicing the examples with a JSP 1.2-compliant container, you will have to use `<body-content>` and `<tag-class>` if you use the following DOCTYPE:

```
<!DOCTYPE taglib PUBLIC
 "-//Sun Microsystems, Inc.//DTD JSP Tag Library 1.2//EN"
 "http://java.sun.com/dtd/web-jsptaglibrary_1_2.dtd" >
```

## 16.2 THE TAG EXTENSION API

The Tag Extension API is a set of interfaces and classes that forms a contract between the JSP container and the tag handler classes. Just as we need to know the Servlet API to write servlets, we need to know the Tag Extension API to write custom tags. This API consists of just one package: `javax.servlet.jsp.tagext`. It has four interfaces and 13 classes. Of these, the most important ones are shown in table 16.4.

**Table 16.4** Important classes and interfaces of the `javax.servlet.jsp.tagext` package

Interface Name	Description
Tag	The Tag interface is the base interface for all tag handlers and is used for writing simple tags. It declares six tag life-cycle methods, including the two most important ones: <code>doStartTag()</code> and <code>doEndTag()</code> . We implement this interface if we want to write a simple tag that does not require iterations or processing of its body content.

*continued on next page*

**Table 16.4 Important classes and interfaces of the `javax.servlet.jsp.tagext` package (continued)**

<b>Interface Name</b>	<b>Description</b>
<code>IterationTag</code>	<code>IterationTag</code> extends the <code>Tag</code> interface and adds one more method for supporting iterations: <code>doAfterBody()</code> .
<code>BodyTag</code>	<code>BodyTag</code> extends <code>IterationTag</code> and adds two methods for supporting the buffering of body contents: <code>doInitBody()</code> and <code>setBodyContent()</code> .
<b>Class Name</b>	<b>Description</b>
<code>TagSupport</code>	<code>TagSupport</code> implements the <code>IterationTag</code> interface and provides a default implementation for all its methods, acting as an <code>IterationTag</code> adapter class. We can use it as a base class for implementing simple and iterative custom tags.
<code>BodyTagSupport</code>	<code>BodyTagSupport</code> implements the <code>BodyTag</code> interface and provides a default implementation for all its methods, acting as a <code>BodyTag</code> adapter class. We can use it as a base class for implementing custom tags that process the contents of the body.
<code>BodyContent</code>	<code>BodyContent</code> extends the <code>JspWriter</code> class and acts as a buffer for the temporary storage of the evaluated body of a tag. This object is used only with the <code>BodyTag</code> interface or the <code>BodyTagSupport</code> class.

All custom tag handlers must implement one of these three interfaces either directly or indirectly (by extending from the adapter classes).

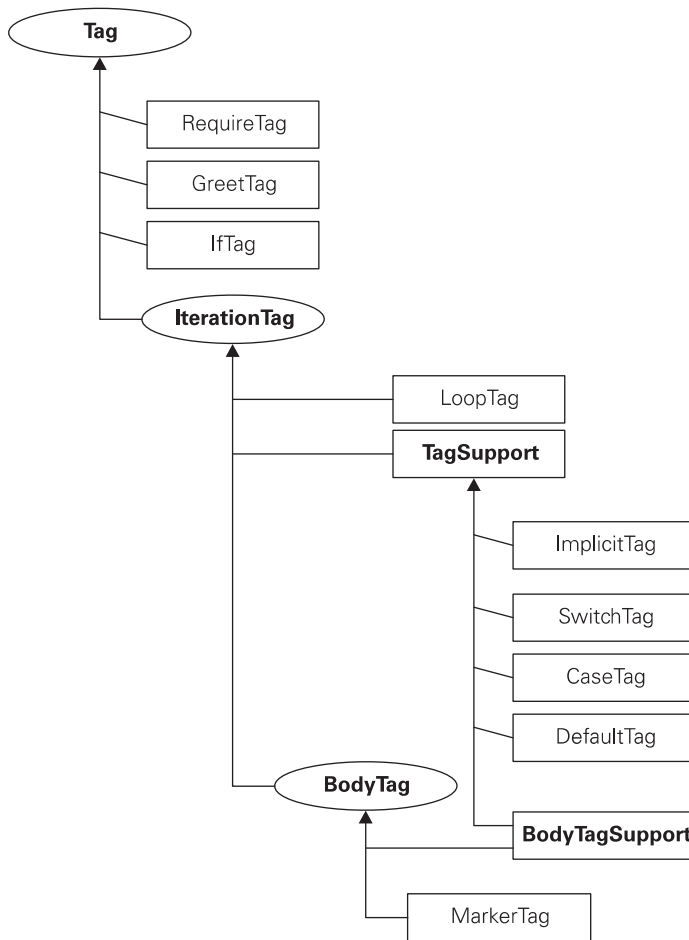
In addition to the interfaces and classes in table 16.4, the `Tag` handler classes use the exception classes shown in table 16.5, which are defined in the `javax.servlet.jsp` package.

**Table 16.5 Exception classes of `javax.servlet.jsp`**

<b>Class Name</b>	<b>Description</b>
<code>JspException</code>	<code>JspException</code> is derived from <code>java.lang.Exception</code> . It is a generic exception that is known to the JSP engine. All uncaught <code>JspExceptions</code> result in an invocation of the error-page machinery. The important methods <code>doStartTag()</code> , <code>doInitBody()</code> , <code>doAfterBody()</code> , and <code>doEndTag()</code> all throw <code>JspException</code> .
<code>JspTagException</code>	<code>JspTagException</code> extends <code>JspException</code> . Tag handler classes can use this exception to indicate unrecoverable errors.

Observe that the `JspTagException` class belongs to the `javax.servlet.jsp` package and not to the `javax.servlet.jsp.tagext` package, as you might expect.

In the next few sections, we are going to examine the interfaces and classes that are listed in table 16.4, using examples from the `sampleLib` package that is downloadable from the Manning web site. Figure 16.1 shows the inheritance relationship between the interfaces and classes of the API and our sample classes. The oval-shaped objects represent the interfaces, the square objects represent the classes, and the arrows represent generalization. The names of the standard classes and interfaces present in the `javax.servlet.jsp.tagext` package are in bold. The other nine classes are the examples that we are going to look at in the following sections.



**Figure 16.1**  
A diagram of the  
Tag Extension API  
and the examples  
in this chapter

## 16.3 IMPLEMENTING THE TAG INTERFACE

The Tag interface is the base interface for all custom tag handlers. It provides the basic life-cycle methods that the JSP engine calls on the tags. Table 16.6 gives a brief description of the methods and constants defined by the Tag interface.

**Table 16.6** Methods and constants of the  
`javax.servlet.jsp.tagext.Tag` interface

Method Name	Description
<code>int doStartTag()</code>	Called when the opening tag is encountered
<code>int doEndTag()</code>	Called when the closing tag is encountered
<code>Tag getParent()</code>	Returns the handler class object of the closest enclosing tag of this tag

*continued on next page*

**Table 16.6 Methods and constants of the `javax.servlet.jsp.tagext.Tag` interface (continued)**

Method Name	Description
<code>void release()</code>	Called on a tag handler to release resources
<code>void setPageContext (PageContext)</code>	Sets the current page context
<code>void setParent (Tag)</code>	Sets the parent (closest enclosing tag handler) of this tag handler
Constant	Description
<code>EVAL_BODY_INCLUDE</code>	Possibly returned by value for <code>doStartTag()</code> Instructs the JSP engine to evaluate the tag body and include it in the output
<code>SKIP_BODY</code>	Possibly returned by value for <code>doStartTag()</code> Instructs the JSP engine not to evaluate the tag body and not to include it in the output
<code>EVAL_PAGE</code>	Possibly returned by value for <code>doEndTag()</code> Instructs the JSP engine to evaluate the rest of the page and include it in the output
<code>SKIP_PAGE</code>	Possibly returned by value for <code>doEndTag()</code> Instructs the JSP engine not to evaluate the rest of the page and not to include it in the output

### 16.3.1 Understanding the methods of the Tag interface

Let's take a closer look at the methods in the life cycle of a custom tag. They are presented here in the sequence in which they are normally called.

#### ***The `setPageContext()` method***

The `setPageContext()` method is the first method that is called in the life cycle of a custom tag. The signature of `setPageContext()` is

```
public void setPageContext (PageContext);
```

The JSP container calls this method to pass the `pageContext` implicit object of the JSP page in which the tag appears. A typical implementation of this method is to save the `pageContext` reference in a private member for future use.

#### ***The `setParent()` and `getParent()` methods***

These methods are used when custom tags are nested one inside the other. In such cases, the outer tag is called the *parent* tag, while the inner tag is called the *child* tag. The signatures of these methods are

```
public void setParent (Tag);
public Tag getParent ();
```

The JSP container calls the `setParent()` method on the child tag and passes it a reference to the parent tag. The `getParent()` method is usually called by one of the

child tags and not directly by the JSP container. A typical implementation of these methods is to save the reference to the parent tag in a private member and return it when required. For example, if a JSP page has multiple tags at more than two levels of nesting, then the JSP engine passes each child tag a reference to its immediate parent. This allows the innermost tag to get a reference to the outermost tag by calling `getParent()` on its immediate parent and then again calling `getParent()` on the returned reference, working its way up the nested hierarchy.

### **The setter methods**

Attributes in custom tags are handled in exactly the same way properties are handled in JavaBeans. If a custom tag has any attributes, then for each attribute, the JSP engine calls the appropriate setter method to set its value at request time. Since the method signatures depend on the attribute names and types, these methods are not defined in the `Tag` interface but are invoked using the standard introspection mechanism that is used in JavaBeans. The setter methods are called after the calls to the `setPageContext()` and `setParent()` methods but before the call to `doStartTag()`.

### **The doStartTag() method**

After setting up the tag with appropriate references by calling the `setPageContext()`, `setParent()`, and setter methods, the container calls the `doStartTag()` method on the tag. The signature of `doStartTag()` is

```
public int doStartTag() throws JspException;
```

This method marks the beginning of the tag's actual processing, giving the tag handler a chance to do initial computations and to verify whether or not the attribute values passed in the setter methods are valid. If the initialization fails, the method may throw a `JspException` or a subclass of `JspException`, such as `JspTagException`, to indicate the problem.

After initialization, the `doStartTag()` method decides whether or not to continue evaluating its body content. As a result, it returns one of the two integer constants defined in the `Tag` interface: `EVAL_BODY_INCLUDE` or `SKIP_BODY`. A return value of `Tag.EVAL_BODY_INCLUDE` indicates that the body must be executed and that its output must be included in the response, while a return value of `Tag.SKIP_BODY` indicates that the body must be skipped and that it is not to be evaluated at all. This method cannot return any other value.

### **The doEndTag() method**

After the body of the tag is evaluated or skipped (depending on the return value of `doStartTag()`), the container calls the `doEndTag()` method. The signature of `doEndTag()` is

```
public int doEndTag() throws JspException;
```

This marks the end of the processing of the tag and gives the tag handler a chance to do the final cleanup for the tag. If anything fails during the cleanup process, the method may throw a `JspException` or a subclass of `JspException`, such as `JspTagException`, to indicate the problem.

Finally, the `doEndTag()` method decides whether or not to continue evaluating the rest of the JSP page. As a result, it returns one of the two integer constants defined in the `Tag` interface: `EVAL_PAGE` or `SKIP_PAGE`. A return value of `Tag.EVAL_PAGE` indicates that the rest of the JSP page must be evaluated and its output be included in the response, while a return value of `Tag.SKIP_PAGE` indicates that the rest of the JSP page must not be evaluated at all and that the JSP engine should return immediately from the current `_jspService()` method. If this page was forwarded or included from another JSP page or a servlet, only the current page evaluation is terminated, and if the page was included, the processing returns to the calling component. This method cannot return any other value.

### ***The release() method***

Finally, the container calls the `release()` method on the handler class when the tag handler object is no longer required. The signature of `release()` is

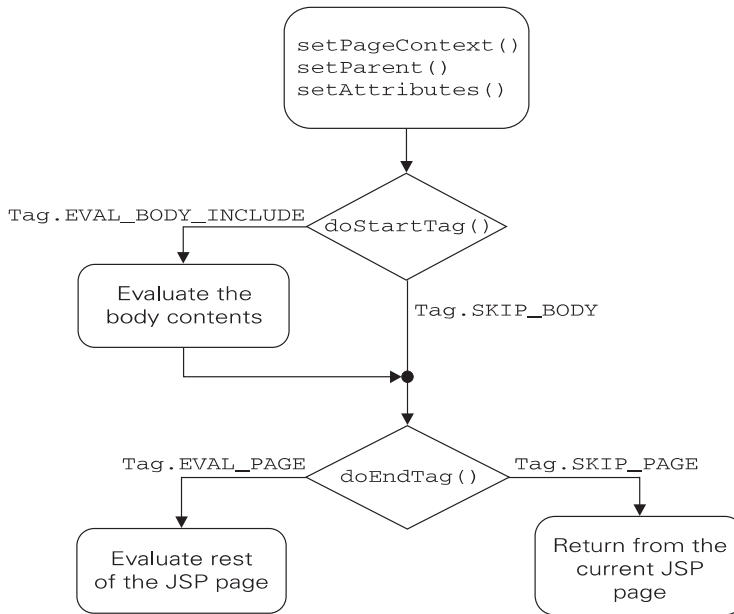
```
public void release();
```

A custom tag may occur multiple times on a JSP page. A single instance of the tag handler may be used to handle all of these occurrences. The JSP container calls the `release()` method on the handler class when the handler object is no longer required. It is important to note that this method is not called after every call to `doEndTag()`. It is called only once, when the container decides to put this instance out of service. For example, if the container implementation maintains a pool of tag handler instances, the container may reuse an instance of a tag by calling the sequence `setPageContext()`, `doStartTag()`, `doEndTag()` multiple times. The container calls the `release()` method only when the tag is to be removed permanently from the pool. This method can be used to release all resources acquired by the tag handler during its lifetime.

The flowchart in figure 16.2 shows the order of processing in a tag handler class that implements the `Tag` interface.

Let's now look at some examples of using the `Tag` interface. We will write tag handlers for the following cases:

- An empty tag that just prints HTML text
- An empty tag that accepts an attribute
- A nonempty tag (a tag with a body) that skips or includes its body content



**Figure 16.2**  
Flowchart for the  
Tag interface

### 16.3.2 An empty tag that prints HTML text

In chapter 15 (section 15.3), we used a tag named `required` on a JSP page to demonstrate the use of an empty tag. It prints the `*` character wherever it is placed on the page:

```
<test:required />
```

Listing 16.2 shows the implementation of the tag handler for this tag.

#### Listing 16.2 RequiredTag.java

```

package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class RequiredTag implements Tag
{
 private PageContext pageContext;
 private Tag parentTag;

 public void setPageContext(PageContext pageContext)
 {
 this.pageContext = pageContext;
 }

 public void setParent(Tag parentTag)
 {
 this.parentTag = parentTag;
 }
}

```

```

public Tag getParent()
{
 return this.parentTag;
}

public int doStartTag() throws JspException
{
 try
 {
 JspWriter out = pageContext.getOut();
 out.print("<font color='#ff0000'*");
 }
 catch(Exception e)
 {
 throw new JspException("Error in RequiredTag.doStartTag()");
 }

 return SKIP_BODY;
}

public int doEndTag() throws JspException
{
 return EVAL_PAGE;
}

//clean up the resources (if any)
public void release()
{
}
}

```

---

This code shows a simple tag handler class, `RequiredTag`, that implements the `Tag` interface and defines all six methods. First, we save the references to `pageContext` and the parent tag in the `setPageContext()` and `setParent()` methods. Note that in this tag, we do not have any attributes, and so we have not defined any setter methods. Then, in the `doStartTag()` method, we use the saved `pageContext` reference to get the output writer of the JSP page and print the HTML code:

```

JspWriter out = pageContext.getOut();
out.print("<font color='#ff0000'*");

```

We return `SKIP_BODY` in `doStartTag()` because we expect the page author to use this as an empty tag. Finally, we return `EVAL_PAGE` in `doEndTag()`, since we want the rest of the page to be executed normally. The following `<tag>` element describes this tag in a TLD file:

```

<tag>
 <name>required</name>
 <tag-class>sampleLib.RequiredTag</tag-class>
 <body-content>empty</body-content>
 <description>Prints * wherever it occurs</description>
</tag>

```



Notice that we have specified `<body-content>` as empty because we do not want to have any body content for this tag.

**NOTE** We could have gotten the output writer and written out the HTML code in `doEndTag()` instead of `doStartTag()`. In tags that are empty, or in tags where `doStartTag()` returns `SKIP_BODY`, such as the above one, it does not matter in which method you choose to print out the HTML code. However, if the tag has a body and if the `doStartTag()` method returns `EVAL_BODY_INCLUDE`, then anything that is printed in `doStartTag()` appears before the body content and anything that is printed in `doEndTag()` appears after the body content in the final output.

### 16.3.3 An empty tag that accepts an attribute

When a tag accepts attributes, there are three important things that we must do for each attribute:

- We must declare an instance variable in the tag class to hold the value of the attribute.
- If we do not want to make the attribute mandatory, we must either provide a default value or take care of the corresponding null instance variable in the code.
- We must implement the appropriate setter methods for each of the attributes.

Let's look at an implementation of the `<greet>` tag that accepts one attribute, `user`, and prints the word `Hello` followed by the user value, in the output HTML:

```
<test:greet />
<test:greet user='john' />
```

The tag prints only the word `Hello` if the `user` attribute is not specified. The following is a code snippet from the file `GreetTag.java`, which is on the Manning web site. We have omitted the parts of the code that are common to all the examples, such as package declarations, import statements, and `setPageContext()`:

```
public class GreetTag implements Tag
{
 //other methods as before

 //A String that holds the user attribute
 private String user;

 //The setter method that is called by the container
 public void setUser(String user) { this.user = user; }

 public int doStartTag() throws JspException
 {
 JspWriter out = pageContext.getOut();

 try
 {
```

```

 if (user==null)
 out.print("Hello!");
 else
 out.print("Hello "+user+"!");
 }
 catch(Exception e)
 {
 throw new JspException("Error in Greet.doStartTag()");
 }
 return SKIP_BODY;
}
}

```

This code is similar to the code for `RequiredTag` except that it has two extra members: a variable and a setter method for the `user` attribute. If and when the JSP engine encounters the `user` attribute in the tag, it calls the `setUser()` method, passing it the attribute's value. The `setUser()` method stores this value in the private instance variable, which is then used by the `doStartTag()` method. In this example, if the page author does not specify the `user` attribute in the `<test:greet>` tag, the `user` variable remains `null` and we print the word `Hello` in the output without a username.

The following `<tag>` element describes this tag in a TLD file:

```

<tag>
 <name>greet</name>
 <tag-class>sampleLib.GreetTag</tag-class>
 <body-content>empty</body-content>
 <description>Prints Hello user! wherever it occurs</description>
 <attribute>
 <name>user</name>
 <required>false</required>
 <rtexprvalue>true</rtexprvalue>
 </attribute>
</tag>

```

You may have been surprised at the use of the `user` variable as an instance variable. Using instance variables to keep request-specific information is very dangerous in servlets since they are not thread safe. However, in the case of custom tags, the onus of thread safety is on the container. It ensures that either a new tag handler instance is created for each occurrence of the tag in a JSP page, or if the container maintains a pool of instances, then an appropriate instance is reused from a pool and the attributes are reset. Thus, in the following example, the second occurrence of the `<test:greet>` tag will not use the value of the `user` attribute passed into the first occurrence:

```

<html><body>
 <test:greet user="john" />
 <test:greet />
</body></html>

```

### 16.3.4 A nonempty tag that includes its body content

The `doStartTag()` method can return either `EVAL_BODY_INCLUDE` or `SKIP_BODY` to include or skip the body content. Let's look at a tag that uses this feature to provide functionality similar to that of an `if` statement in a programming language. The following is a code snippet from the `IfTag.java` file you can find on the Manning web site:

```
public class IfTag implements Tag
{
 //other methods as before

 private boolean condition = false;

 public void setCondition(boolean condition)
 {
 this.condition = condition;
 }

 public int doStartTag() throws JspException
 {
 if (condition)
 return EVAL_BODY_INCLUDE;
 else
 return SKIP_BODY;
 }
}
```

In this example, we use a boolean attribute, `condition`, to determine whether the body needs to be included or skipped. In the `doStartTag()` method, depending on the condition value, we return either `Tag.EVAL_BODY_INCLUDE` or `Tag.SKIP_BODY`.

The following `<tag>` element describes the `<if>` tag in a TLD file:

```
<tag>
 <name>if</name>
 <tag-class>sampleLib.IfTag</tag-class>
 <body-content>JSP</body-content>

 <attribute>
 <name>condition</name>
 <required>true</required>
 <rtexprvalue>true</rtexprvalue>
 </attribute>
</tag>
```

Observe that we have specified the `<body-content>` as `JSP`. This allows us to write any valid JSP code in the body of the tag. If the `doStartTag()` method of the tag handler returns `EVAL_BODY_INCLUDE`, the body will be executed like normal JSP code; otherwise, it will be skipped altogether. We have specified the value of the `<required>` element as `true`, since we need it to decide whether or not to include the body content. The following code shows the usage of the `<if>` tag in the JSP page:

```

<%@ taglib prefix="test" uri="/WEB-INF/sampleLib.tld" %>
<% boolean debug = "true".equals(request.getParameter("debug")); %>

<html><body>
Hello

 <test:if condition="<%= debug %>" >
 DEBUG INFO:...
 </test:if>

</body></html>

```

When we access the above JSP page as

```
http://localhost:8080/chapter16/ifTest.jsp?debug=true
```

the output will be:

```

Hello
DEBUG INFO:...

```

If we pass `debug=false` in the query string of the URL, the output will not contain the second line.

## 16.4 IMPLEMENTING THE ITERATIONTAG INTERFACE

In the previous examples, we used the `Tag` interface to either include or skip the body content of the tag. However, if the body content was included, it was included only once. The `IterationTag` interface extends the `Tag` interface and allows us to include the body content multiple times, in a way that is similar to the loop functionality of a programming language. The `IterationTag` interface declares one method and one constant, as shown in table 16.7.

**Table 16.7** Methods of the `javax.servlet.jsp.tagext.IterationTag` interface

Method Name	Description
<code>int doAfterBody()</code>	This method is called after each evaluation of the tag body. It can return either of two values: <code>IterationTag.EVAL_BODY_AGAIN</code> or <code>Tag.SKIP_BODY</code> . The return value determines whether or not the body needs to be reevaluated.
Constant	Description
<code>EVAL_BODY_AGAIN</code>	Possible return value for <code>doAfterBody()</code> . This constant instructs the JSP engine to evaluate the tag body and include it in the output.

### 16.4.1 Understanding the `IterationTag` methods

Since `IterationTag` extends `Tag`, it inherits all the functionality of the `Tag` interface. The container sets up the iterative tag with appropriate references by calling the `setPageContext()` and `setParent()` methods, passes the attribute values using the setter methods, and calls `doStartTag()`. Depending on the return value of `doStartTag()`, the container either includes or skips the body content.

If `doStartTag()` returns `SKIP_BODY`, then the body is skipped and the container calls `doEndTag()`. In this case, the `doAfterBody()` method is never called on the iterative tag. However, if `doStartTag()` returns `EVAL_BODY_INCLUDE`, the body of the tag is evaluated, the result is included in the output, and the container calls `doAfterBody()` for the very first time.

### **The `doAfterBody()` method**

The `doAfterBody()` is the only method defined by the `IterationTag` interface. It gives the tag handler a chance to reevaluate its body. The signature of `doAfterBody()` is

```
public int doAfterBody() throws JspException;
```

If an error occurs during the invocation of the `doAfterBody()` method, it may throw a `JspException` or a subclass of it, such as `JspTagException`, to indicate the problem. If everything goes fine, it decides whether or not to reevaluate its body. To evaluate the body again, it will return the integer constant `EVAL_BODY_AGAIN`, which is defined in the `IterationTag` interface. This will cause an evaluation of the tag's body the second time, and after the evaluation, the JSP container will call `doAfterBody()` for the second time. This process continues until `doAfterBody()` returns `SKIP_BODY`, which is defined in the `Tag` interface. We cannot return any other value from this method.

Finally, the `doEndTag()` method is called, either because `doStartTag()` returns `SKIP_BODY`, or because `doAfterBody()` returns `SKIP_BODY`. The purpose and functionality of the `doEndTag()` method in the `IterationTag` interface are the same as in the `Tag` interface.

The flowchart in figure 16.3 shows the order of processing in a tag handler class that implements the `IterationTag` interface.

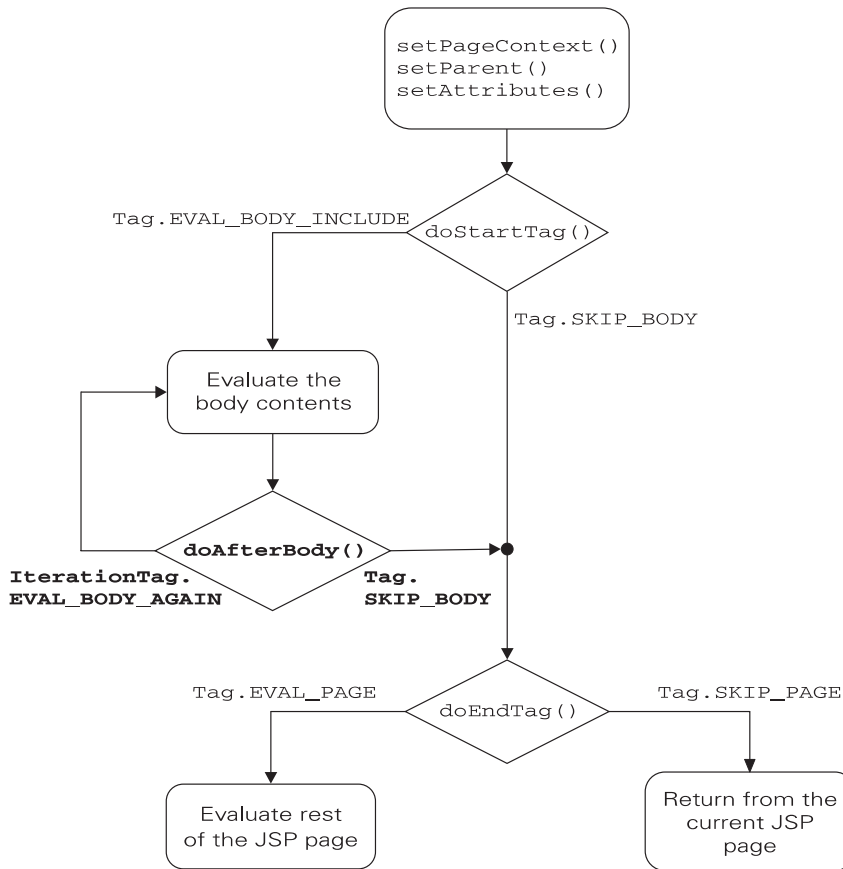
## **16.4.2 A simple iterative tag**

Let's now look at a tag that provides the functionality similar to that of a looping construct in a programming language. The following code snippet from a JSP page demonstrates the use of the loop tag:

```
<%@ taglib prefix="test" uri="/WEB-INF/sampleLib.tld" %>
<html><body>
 <test:loop count="5" >
 Hello World!

 </test:loop>
</body></html>
```

The above tag has an attribute named `count` that accepts integral values to specify the number of times the body of the tag should be executed. The above code prints `Hello World!` five times in the output.



**Figure 16.3** Flowchart for the `IterationTag` interface

Listing 16.3 shows the code for `LoopTag.java`. This tag handler class implements the `IterationTag` interface and thus provides implementation for the extra method `doAfterBody()` as well as the six methods of the `Tag` interface.

#### Listing 16.3 `LoopTag.java`

```

package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class LoopTag implements IterationTag
{
 private PageContext pageContext;
 private Tag parentTag;

 public void setPageContext(PageContext pageContext)
 {

```

```

 this.pageContext = pageContext;
 }

 public void setParent(Tag parentTag)
 {
 this.parentTag = parentTag;
 }

 public Tag getParent()
 {
 return this.parentTag;
 }

 //Attribute to maintain looping count
 private int count = 0;

 public void setCount(int count)
 {
 this.count = count;
 }

 public int doStartTag() throws JspException
 {
 if (count>0)
 return EVAL_BODY_INCLUDE;
 else
 return SKIP_BODY;
 }

 public int doAfterBody() throws JspException
 {
 if (--count > 0)
 return EVAL_BODY_AGAIN;
 else
 return SKIP_BODY;
 }

 public int doEndTag() throws JspException
 {
 return EVAL_PAGE;
 }

 public void release()
 {
 }
}

```

---

In listing 16.3, we have used the `count` variable to keep track of the number of iterations. For each invocation of the `doAfterBody()` method, we decrement the count by 1 and return `EVAL_BODY_AGAIN` until the count reaches 0. If the count reaches 0, we return `SKIP_BODY` to terminate the looping effect, which tells the container to skip further iterations and call `doEndTag()`.

The following `<tag>` element describes the loop tag in a TLD file:

```
<tag>
 <name>loop</name>
 <tag-class>sampleLib.LoopTag</tag-class>
 <body-content>JSP</body-content>
 <attribute>
 <name>count</name>
 <required>true</required>
 <rtexprvalue>true</rtexprvalue>
 </attribute>
</tag>
```

As you can see, describing an `IterationTag` is no different than defining a normal tag in a TLD. We have not informed the container explicitly whether or not this tag is iterative. The container introspects the interface implemented by the tag's class and calls `doAfterBody()` only if it finds that the class is an instance of `IterationTag`.

## 16.5 IMPLEMENTING THE BODYTAG INTERFACE

The `BodyTag` interface extends `IterationTag` and adds a new functionality that lets the tag handler evaluate its body content in a temporary buffer. This feature allows the tag to process the generated contents at will. For example, after evaluation, the tag handler can view the body content, discard it completely, modify it, or add more data to it before sending it to the output stream. Since it is derived from `IterationTag`, `BodyTag` can also handle the evaluation and processing of the content as many times as required.

The `BodyTag` interface declares two methods and a constant, as shown in table 16.8.

**Table 16.8** Methods of the `javax.servlet.jsp.tagext.BodyTag` Interface

Method Name	Description
<code>void</code>	
<code>setBodyContent (BodyContent)</code>	Called by the JSP container to pass a reference to a <code>BodyContent</code> object.
<code>void doInitBody()</code>	Called by the JSP container after calling <code>setBodyContent()</code> , to allow the tag handler class to perform initialization steps on <code>BodyContent</code> .
Constant	Description
<code>EVAL_BODY_BUFFERED</code>	A constant defined as a return value for <code>doStartTag()</code> and <code>doAfterBody()</code> for tags that want to buffer the evaluation of their content.

**NOTE** In JSP 1.1, there was another return value for `BodyTag.doAfterBody()`: `EVAL_BODY_TAG`. This value is now deprecated. If the exam asks about it, you should treat it the same as `IterationTag.EVAL_BODY_AGAIN` or `BodyTag.EVAL_BODY_BUFFERED`, as the case may be.



## 16.5.1 Understanding the methods of BodyTag

The `BodyTag` interface adds two new methods to handle the processing of the tag's body: `setBodyContent()` and `doInitBody()`.

### **The `setBodyContent()` method**

The JSP container calls the `setBodyContent()` method to pass an instance of `BodyContent` to the tag. The signature of `setBodyContent()` is

```
public void setBodyContent(BodyContent);
```

A typical implementation of this method is to save the `BodyContent` reference in a private member for future use.

### **The `doInitBody()` method**

The JSP container calls the `doInitBody()` method after calling `setBodyContent()`. The signature of `doInitBody()` is

```
public void doInitBody() throws JspException;
```

This method allows the tag handler class to initialize the `BodyContent` object, if required, before the actual evaluation process starts. Thus, if the initialization of `BodyContent` fails, the `doInitBody()` method may throw a `JspException` or a subclass of `JspException`, such as `JspTagException`, to indicate the problem.

Since `BodyTag` extends `IterationTag`, which in turn extends the `Tag` interface, `BodyTag` inherits all the functionality of `IterationTag` as well as `Tag`. The container sets up the implementation handler class with appropriate references by calling the `setPageContext()` and `setParent()` methods, passes the attribute values using the setter methods, and calls `doStartTag()`.

The `doStartTag()` method of a class that implements the `BodyTag` interface returns any one of three values: `EVAL_BODY_INCLUDE` or `SKIP_BODY` inherited from the `Tag` interface, or `EVAL_BODY_BUFFERED`, which is defined in the `BodyTag` interface. The actions taken by the JSP container for the return values `EVAL_BODY_INCLUDE` and `SKIP_BODY` are the same as for the `IterationTag` interface.

However, if `doStartTag()` returns `EVAL_BODY_BUFFERED`, the JSP container takes a different course of action. It first creates an instance of the `BodyContent` class. The `BodyContent` class is a subclass of `JspWriter` and overrides all the `print` and `write` methods of `JspWriter` to buffer any data written into it rather than sending it to the output stream of the response. The JSP container passes the newly created `BodyContent` instance to the tag handler using its `setBodyContent()` method, calls `doInitBody()` on the tag, and finally evaluates the body of the tag, filling the `BodyContent` buffer with the result of the body tag evaluation.

The container calls `doAfterBody()` after evaluating the body, writing the data directly into the output or buffering it, as the case may be. If the output was buffered, we can add, modify, or delete the contents of this buffer. Finally, this method returns

EVAL\_BODY\_AGAIN or EVAL\_BODY\_BUFFERED to continue evaluating the body in a loop, or returns SKIP\_BODY to terminate the loop.

Finally, the container calls doEndTag(), and, as with the other interfaces, the tag handler class that is implementing the BodyTag interface can return either SKIP\_PAGE or EVAL\_PAGE.

The flowchart in figure 16.4 shows the order of processing in a tag handler class that implements the BodyTag interface.

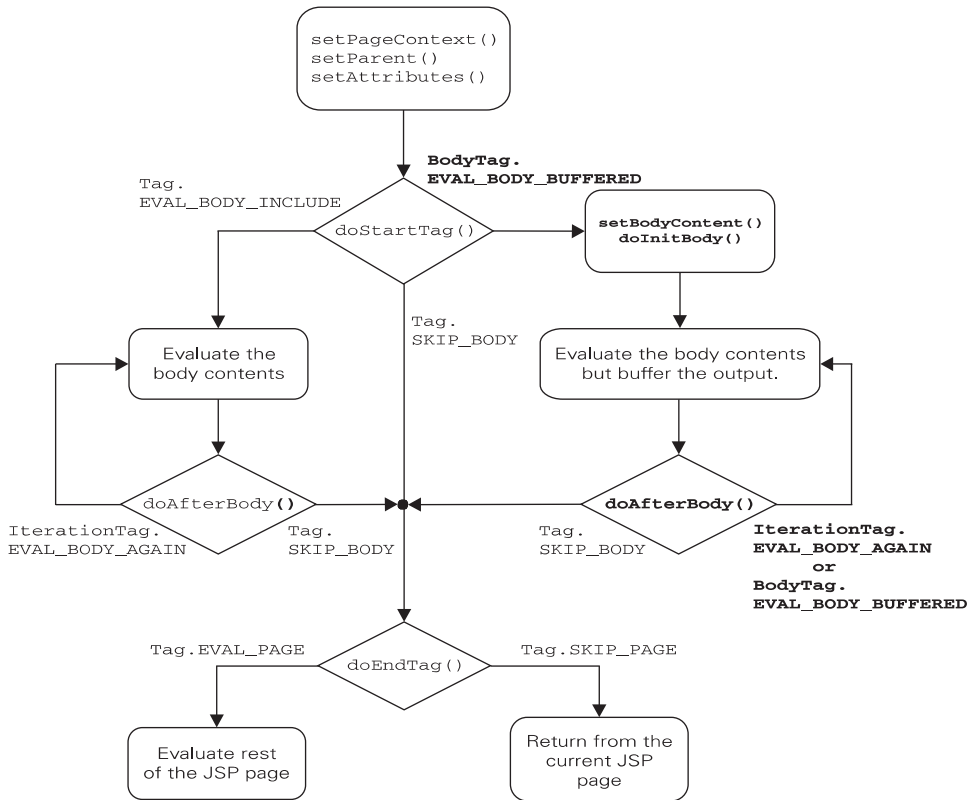


Figure 16.4 Flowchart for the BodyTag interface

## 16.5.2 A tag that processes its body

Let's now look at a tag named mark that displays certain characters in its tag body in boldface when they are specified in the search attribute of the tag. For example:

```

<test:mark search="s">
 she sells sea shells on the sea shore!
</test:mark>

```

will print

```

 she sells sea shells on the sea shore!

```

whereas if we pass `sh` to the `search` attribute, it will print

```
she sells sea shells on the sea shore!
```

This kind of feature is useful if your site maintains many informative documents and allows a search on them using a keyword. The output of the search engine can be nested inside the `mark` tag with the search string shown in bold.

Listing 16.4 contains the code for the tag handler class of this tag.

#### Listing 16.4 MarkerTag.java

```
package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class MarkerTag implements BodyTag {

 //INITIALIZATION
 private PageContext pageContext;
 private Tag parentTag;

 public void setPageContext(PageContext pageContext)
 {
 this.pageContext = pageContext;
 }

 public void setParent(Tag parentTag)
 {
 this.parentTag = parentTag;
 }

 public Tag getParent()
 {
 return this.parentTag;
 }

 //attributes
 private String search = null;

 public void setSearch(String search)
 {
 this.search = search;
 }

 //BODY CONTENT RELATED MEMBERS

 private BodyContent bodyContent;

 public void setBodyContent(BodyContent bodyContent)
 {
 this.bodyContent = bodyContent;
 }

 public void doInitBody() throws JspException
 {
 }
}
```

```

//START, ITERATE, AND END METHODS

public int doStartTag() throws JspException
{
 return EVAL_BODY_BUFFERED;
}

public int doAfterBody() throws JspException
{
 try{
 JspWriter out = bodyContent.getEnclosingWriter();

 String text = bodyContent.getString();

 int len = search.length();

 int oldIndex=0, newIndex=0;

 while((newIndex = text.indexOf(search,oldIndex))>=0){

 if (newIndex<oldIndex)
 {
 break;
 }
 out.print(text.substring(oldIndex,newIndex));
 out.print(""+search+"");
 oldIndex = newIndex + len;
 }

 out.print(text.substring(oldIndex));
 }
 catch(Exception e){
 e.printStackTrace();
 }

 return SKIP_BODY;
}

public int doEndTag() throws JspException {
 return EVAL_PAGE;
}

public void release()
{
}
}

```

In listing 16.4, we have used an object of type `BodyContent`.

Note that we have used `EVAL_BODY_BUFFERED` in the return value of `doStartTag()`. If we had returned `EVAL_BODY_INCLUDE` instead, it would have thrown a `NullPointerException` upon using the `bodyContent` object in the `doAfterBody()`:

```
JspWriter out = bodyContent.getEnclosingWriter();
```

This is because `setBodyContent()` is not called and the `bodyContent` object is not set if `doStartTag()` returns `EVAL_BODY_INCLUDE`.

## 16.6 EXTENDING TAGSUPPORT AND BODYTAGSUPPORT

Until now, we have been writing tag classes that directly implement the `Tag`, `IterationTag`, and `BodyTag` interfaces. We chose to do this in order to learn about the flow of events that occur during the execution of a tag in a JSP page. In practice, however, we do not need to write all of the methods ourselves. The API provides two adapter classes, `TagSupport` and `BodyTagSupport`, that implement the `IterationTag` interface and the `BodyTag` interface, respectively, and provide a default implementation of all the methods. Thus, we only need to override those methods that have to be customized.

### 16.6.1 The TagSupport class

The `TagSupport` class implements the `IterationTag` interface, and provides default implementations for each of the methods of the `Tag` and `IterationTag` interfaces. It adds some new convenience methods that allow us to maintain a list of named objects in a hashtable and a method to find an outer tag of a given class from an inner tag. Table 16.9 lists some of the important methods of the `TagSupport` class and a protected attribute.

**Table 16.9** Methods of the `TagSupport` class

Method	Description
<b>Important Overridden Methods and Their Default Return Values</b>	
<code>int doStartTag()</code>	Inherited from <code>Tag</code> . The default return value is <code>SKIP_BODY</code> .
<code>int doAfterBody()</code>	Inherited from <code>IterationTag</code> . The default return value is <code>SKIP_BODY</code> .
<code>int doEndTag()</code>	Inherited from <code>Tag</code> . The default return value is <code>EVAL_PAGE</code> .
<b>Methods Useful for Convenient Handling of Nested Tags</b>	
<code>void setParent(Tag)</code>	Accepts and maintains a reference to the parent tag.
<code>Tag getParent()</code>	Returns the reference to the parent tag.
<code>Tag findAncestorWithClass(Tag, Class)</code>	This is a static method. Given a reference to a tag and the desired class, it will find the closest ancestor of the given tag of the given class. Internally, it calls <code>getParent()</code> on each of the references returned.
<b>Convenience Methods Useful for Maintaining a Map of Name-Value Pairs, Especially for Accepting Tag Attributes</b>	
<code>void setValue(String, Object)</code>	Accepts a name-value pair. The name is a <code>String</code> and the value can be any <code>Object</code> .

*continued on next page*

**Table 16.9 Methods of the TagSupport class (continued)**

Method	Description
Object <code>getValue(String)</code>	Returns the object for the supplied String name.
Enumeration <code>getValues()</code>	Returns an Enumeration of all the values.
void <code>removeValue(String)</code>	Removes the name-value pair from the list for the given name.
<b>Protected Member That Can Be Used by Derived Classes</b>	
PageContext <code>pageContext</code>	A reference to the saved PageContext object.

## 16.6.2 The BodyTagSupport class

The BodyTagSupport class extends the TagSupport class, inheriting all the functionality shown in table 16.9. In addition, it implements the BodyTag interface and provides default implementations of the `setBodyContent()` and `doInitBody()` methods. It also provides two convenience methods for handling buffered output: `getBodyContent()` and `getPreviousOut()`. Table 16.10 lists some of the important methods of the BodyTagSupport class.

**Table 16.10 Methods of the BodyTagSupport class**

Method	Description
<b>Important Overridden Methods and Their Default Return Values</b>	
int <code>doStartTag()</code>	Inherited from Tag. The default return value is EVAL_BODY_BUFFERED.
int <code>doAfterBody()</code>	Inherited from IterationTag. The default return value is SKIP_BODY.
int <code>doEndTag()</code>	Inherited from Tag. The default return value is EVAL_PAGE.
<b>Methods Useful for Convenient Handling of the Buffered Output</b>	
void <code>setBodyContent(BodyContent)</code>	Accepts and maintains a reference to the BodyContent object. The BodyContent object is a wrapper around the actual JspWriter object. It acts as the current output stream, but does not write the output directly to the client. Instead, it buffers it for further processing.
BodyContent <code>getBodyContent()</code>	Returns the reference to the BodyContent object.
JspWriter <code>getPreviousOut()</code>	Returns the output writer object, which is wrapped inside the BodyContent object.

## 16.6.3 Accessing implicit objects

One of the greatest features of custom tags is their ability to access all the objects that are accessible to the JSP page in which they appear from within the tag handler classes. This is done with the help of the PageContext object, which is set by the container using the `setPageContext()` method before calling the `doStartTag()` method. Using the pageContext object, we can access any other object available to the page.

For example, we have used it to get the `JspWriter` object to write out HTML in some of the previous examples:

```
JspWriter out = pageContext.getOut();
```

Table 16.11 lists the methods to access the four implicit objects that act as scope containers for other objects.

**Table 16.11** The four implicit objects that act as scope containers

Scope	Implicit variable	Implicit Object Class	Getting the Object from within the Tag Class	
			Using Direct Getters	As Named Objects
Application	<code>application</code>	<code>ServletContext</code>	<code>pageContext.getServlet-Context()</code>	<code>pageContext.getAttribute(PageContext.APPLICATION);</code>
Session	<code>session</code>	<code>HttpSession</code>	<code>pageContext.getSession()</code>	<code>pageContext.getAttribute(PageContext.SESSION);</code>
Request	<code>request</code>	<code>ServletRequest</code>	<code>pageContext.getRequest()</code>	<code>pageContext.getAttribute(PageContext.REQUEST);</code>
Page	<code>pageContext</code>	<code>PageContext</code>	-	<code>pageContext.getAttribute(PageContext.PAGECONTEXT);</code>

Let's now look at a tag named `<implicit>` that accepts two attributes: `attributeName` and `scopeName`. It searches for an object with the given name in the given scope and will print it out in the output stream. The following JSP page code shows the intended usage of this tag:

```
<%@ taglib prefix="test" uri="/WEB-INF/sampleLib.tld" %>
<html><body>
 <%
 application.setAttribute("attribute1", "somestring");
 session.setAttribute("attribute2", new Boolean(true));
 request.setAttribute("attribute3", new Integer(5));
 %>

 <test:implicit attributeName="attribute1" scopeName="application"/>
 <test:implicit attributeName="attribute2" scopeName="session"/>
 <test:implicit attributeName="attribute3" scopeName="request"/>
</body></html>
```

In the above tag usage, we set `attribute1`, `attribute2`, and `attribute3` in the application, session, and request scopes, respectively. Our intention is now to

print the values of these attributes in the output HTML through our `<implicit>` tag. To do this, we pass the `attributeName` and `scopeName` to the three `<test:implicit>` tags one by one. The tags use these values to generate the following output:

```
<html><body>
 someString
 true
 5
</body></html>
```

Listing 16.5 contains the code for `ImplicitTag.java`, which implements the `<implicit>` tag by extending the `TagSupport` class.

#### Listing 16.5 `ImplicitTag.java`

```
package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class ImplicitTag extends TagSupport {

 public void setAttributeName(String name)
 {
 //Stores the passed object in the hashtable maintained by
 //TagSupport with the name "attributeName".
 setValue("attributeName", name);
 }

 public void setScopeName(String scope)
 {
 //Stores the passed object in the hashtable with
 //the name "scopeName".
 setValue("scopeName", scope);
 }

 //Our utility method to convert the scopeName String
 //to the integer constant defined in PageContext for each scope.
 //We need this method because we have to use
 //PageContext.getAttribute(String name, int scope) later.

 private int getScopeAsInt()
 {
 //Retrieve the scopeName value from the hashtable
 String scope = (String) getValue("scopeName");

 if ("request".equals(scope))
 return PageContext.REQUEST_SCOPE;

 if ("session".equals(scope))
 return PageContext.SESSION_SCOPE;

 if ("application".equals(scope))
 return PageContext.APPLICATION_SCOPE;
 }
}
```



```

 //Default is page scope
 return PageContext.PAGE_SCOPE;
 }

 public int doStartTag() throws JspException
 {
 try
 {
 JspWriter out = pageContext.getOut();

 String attributeName = (String) getValue("attributeName");
 int scopeConstant = getScopeAsInt();

 out.print(pageContext.getAttribute(attributeName, scopeConstant));

 return SKIP_BODY;
 }
 catch(Exception e)
 {
 throw new JspException("Error in Implicit.doAfterBody()");
 }
 }
}

```

---

There are three points worth noting in listing 16.5:

- We have implemented a setter method for each attribute. However, instead of defining a private instance variable for storing each attribute, we have used the hashtable maintained by the `TagSupport` class to store the attributes as name-value pairs.
- `getScopeAsInt()` is our utility method that returns the integer constant representing the scope name that is stored as a string in the hashtable. These constants are already defined by `PageContext` and are used by the `PageContext.getAttribute(String name, int scope)` method.
- Finally, in the `doStartTag()` method, we use the `PageContext.getAttribute(String name, int scope)` method to retrieve the value of the given attribute name from the given scope name. After printing the value of the attribute, we return `SKIP_BODY`, since we do not want to accept any body content for the `<implicit>` tag.

This tag can be easily described in a TLD as follows:

```

<tag>
 <name>implicit</name>
 <tag-class>sampleLib.ImplicitTag</tag-class>
 <body-content>empty</body-content>
 <attribute>
 <name>attributeName</name>
 <required>true</required>
 </attribute>
</tag>

```

```

 <attribute>
 <name>scopeName</name>
 <required>true</required>
 </attribute>
 </tag>

```

## 16.6.4 Writing cooperative tags

Since tags are usually designed and developed with a common pattern of problems in mind, we often build a group of tags that work together. Such tags are called *cooperative tags*. One of the simplest examples to demonstrate the usage of cooperative tags is to implement the switch-case functionality similar to the one provided by the Java programming language. Let us look at three tags—`<switch>`, `<case>`, and `<default>`—that can be used in a JSP page, as shown in listing 16.6.

**Listing 16.6** switchTest.jsp

```

<html><body>

 <%@ taglib prefix="test" uri="/WEB-INF/sampleLib.tld" %>

 <% String action = request.getParameter("action"); %>

 <test:switch conditionValue="<%= action %%" >

 <test:case caseValue="sayHello">
 Hello!
 </test:case>

 <test:case caseValue="sayGoodBye" >
 Good Bye!!
 </test:case>

 <test:default>
 I am Dumb!!!
 </test:default>

 </test:switch>

</body></html>

```

The `conditionValue` attribute of the `switch` tag acts like the `switch` condition of the Java `switch` statement, while the `caseValue` attribute of the `case` tag acts like the `case` value of the Java `switch` statement. Only those `case` tags whose `caseValue` attribute value matches the value of the `conditionValue` attribute of the `switch` tag should print their body contents. Thus, the above page should print `Hello!` in the browser when accessed through the URL

```
http://localhost:8080/chapter16/switchTest.jsp?action=sayHello
```

Let's now look at `SwitchTag.java`, `CaseTag.java`, and `DefaultTag.java`, which implement these tags (listing 16.7).

### Listing 16.7 SwitchTag.java

```
package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class SwitchTag extends TagSupport
{
 public void setPageContext(PageContext pageContext)
 {
 super.setPageContext(pageContext);

 //Sets the internal flag that tells whether or not a matching
 //case tag has been found to be false.
 setValue("caseFound", Boolean.FALSE);
 }

 //stores the value of the match attribute
 public void setConditionValue(String value)
 {
 setValue("conditionValue", value);
 }

 public int doStartTag() throws JspException
 {
 return EVAL_BODY_INCLUDE;
 }
}
```

The code for the `SwitchTag` class is quite simple and has just three methods: `setPageContext()`, a setter for the `conditionValue` attribute, and `doStartTag()`. The `caseFound` flag indicates whether or not a matching case tag has been found. The `pageContext()` method initializes it to `false`. We will show the use of this flag in the case and default tags shortly. The setter method stores the value of the `conditionValue` attribute using the `setValue()` method, as explained in the previous example. We don't want to do anything in the `doStartTag()` method, but we do want the body of the switch tag to be evaluated. However, the default implementation of `doStartTag()` provided by `TagSupport` returns `SKIP_BODY`, so we need to override it and return `EVAL_BODY_INCLUDE` instead.

Listing 16.8 is the code for the case tag handler.

### Listing 16.8 CaseTag.java

```
package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class CaseTag extends TagSupport
{
```

```

public void setCaseValue(String caseValue)
{
 setValue("caseValue",caseValue);
}

public int doStartTag() throws JspException
{
 //gets the reference of the enclosing switch tag handler.
 SwitchTag parent =
 (SwitchTag) findAncestorWithClass(this, SwitchTag.class);

 Object caseValue = this.getValue("caseValue");
 Object conditionValue = parent.getValue("conditionValue");

 //If the value of the caseValue attribute of this case tag
 //matches with the value of the conditionValue attribute of
 //the parent switch tag, it sets the caseFound flag to true and
 //includes the body; otherwise, it skips the body.

 if (conditionValue.equals(caseValue))
 {
 //Sets the caseFound flag to true
 parent.setValue("caseFound",Boolean.TRUE);

 //Includes the body contents in the output HTML
 return EVAL_BODY_INCLUDE;
 }
 else
 {
 return SKIP_BODY;
 }
}
}

```

There are two points worth noting in listing 16.8:

- The `setCaseValue()` method stores the value attribute using the `setValue()` method of `TagSupport`.
- In `doStartTag()`, we first get a reference of the parent switch tag using the `findAncestorWithClass()` method of `TagSupport`. Next, we retrieve the value of the `conditionValue` attribute from the parent tag handler reference and the value of the `caseValue` attribute from the current tag. If the two values match, we set the `caseFound` flag to `true` and return `EVAL_BODY_INCLUDE` so that the body of this case tag is included in the output. Otherwise, if the values don't match, we return `SKIP_BODY`.

Now, let's look at the code for the default tag handler in listing 16.9.

## Listing 16.9 DefaultTag.java

```
package sampleLib;

import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class DefaultTag extends TagSupport
{
 public int doStartTag() throws JspException
 {
 SwitchTag parent = (SwitchTag)
 findAncestorWithClass(this, SwitchTag.class);

 Boolean caseFound = (Boolean) parent.getValue("caseFound");

 //If the conditionValue attribute value of the switch tag
 //did not match with any of the caseValue attribute values,
 //then it includes the body of this tag; otherwise; it skips the body.

 if (caseFound.equals(Boolean.FALSE))
 {
 return EVAL_BODY_INCLUDE;
 }
 else
 {
 return SKIP_BODY;
 }
 }
}
```

The implementation of the default tag handler checks whether the `caseFound` attribute of the enclosing `SwitchTag` instance is set to `true`. If it is set to `false`, it means that none of the `caseValues` matched with the `conditionValue`, in which case the `DefaultTag` will include its own body. Otherwise, it will skip the body.

The following elements describe these tags in the TLD:

```
<tag>
 <name>switch</name>
 <tag-class>sampleLib.SwitchTag</tag-class>
 <body-content>JSP</body-content>
 <attribute>
 <name>conditionValue</name>
 <required>true</required>
 <rtexprvalue>true</rtexprvalue>
 </attribute>
</tag>

<tag>
 <name>case</name>
 <tag-class>sampleLib.CaseTag</tag-class>
```

```

<body-content>JSP</body-content>
<attribute>
 <name>caseValue</name>
 <required>true</required>
</attribute>
</tag>

<tag>
 <name>default</name>
 <tag-class>sampleLib.DefaultTag</tag-class>
 <body-content>JSP</body-content>
</tag>

```

## 16.7 WHAT'S MORE?

We have discussed the various types of tags, interfaces, and classes provided by the Tag Extension API, which is information that you need to know in order to do well on the exam. However, the story of custom tags does not end here. To get maximum benefit from the use of custom tags, you should learn how to use other classes and the interfaces of the Tag Extension API as well. For example, you can learn how to use the `TagLibraryValidator` and the `PageData` classes to validate the semantics of the tags used in a JSP page at translation time, or how to make scripting variables in a JSP page available—in a manner similar to `<jsp:useBean>`—via the `TagExtraInfo` and `TagVariableInfo` classes. You can also learn how to handle resources more efficiently by implementing the `TryCatchFinally` interface, and many other techniques for working with custom tags.

### ***The difference between custom tags and JavaBeans***

Many times developers ask when they should use custom tags and when they should use JavaBeans. They both are reusable components and help us to reduce the length of our scriptlets and make our JSP pages cleaner and more manageable. Then, how do we know when to use which type of component? For example, should the functionality of database access go into JavaBeans, or in custom tags?

After reading chapters 14, 15, and 16, you should know that the two types of components serve two different purposes. Here are some differences between the two that might help you determine which one to use:

- JavaBeans are the data handlers of JSP pages and aid in encapsulating data-management logic. They are used for storage. Tags, on the other hand, aid computational logic related to a particular request.
- Tags are thread safe; beans are not. Beans, like other separate utility classes, have to be made thread safe by the developers.
- Tags are aware of the environment (the page context) in which they execute. Beans are not.

- Tags remain in the translation unit. We can think of tags as events occurring in the execution of a JSP page. Beans are object stores that reside outside the translation unit.
- Tags can access implicit objects. Beans cannot.
- Tags only have page scope. They are created and destroyed within a single request and in a single page. They can access other objects in all the scopes, though. Beans, on the other hand, are themselves objects that reside in different scopes. Therefore, tags can access and manipulate beans, while beans do not access and manipulate tags.
- The Tag Extension API is designed closely with the concept of a JSP page in mind. They may not be used in other applications. Beans, on the other hand, are supposed to be reusable components and can be used by other containers.
- Tags are not persistent objects. Beans have properties, and properties have values. A set of values is called the *state* of the bean. This state can be persisted via serialization and reused later.

So to answer the question about whether to use beans or tags for database access, let's clarify that we can use beans to access a database, to encapsulate data, and to implement business logic rules to manipulate data in the beans. The code for managing the beans across scopes should be handled by custom tags. The code that uses the bean's properties and includes presentation logic should be placed in custom tags. Thus, tags are a preferred way of writing JSP pages that use JavaBeans.

## 16.8 SUMMARY

In this chapter, we learned how to create our own custom tag libraries according to the “Custom” model of development. The tag library descriptor (TLD) file contains the information that the JSP engine needs to know about the tag library in order to successfully interpret the custom tags on JSP pages. We discussed the TLD file and its three important elements: `<tag>`, `<attribute>`, and `<body-content>`.

The Tag Extension API consists of one package: `javax.servlet.jsp.tagext`, with 4 interfaces and 13 classes. We examined in detail the methods and constants of the three interfaces: `Tag`, `IterativeTag`, and `BodyTag`. We then saw how we can implement those interfaces in classes. In addition, the JSP API provides two adapter classes, `TagSupport` and `BodyTagSupport`, that implement the `IterationTag` interface and `BodyTag` interface, respectively, and that provide default implementation of all the methods.

At this point, you should be able to answer exam questions about the structure and format of the TLD file elements as well as questions based on the Tag Extension API and implementation of custom tag libraries.

## 16.9 REVIEW QUESTIONS

1. Which of the following is not a valid subelement of the `<attribute>` element in a TLD? (Select one)

- a `<name>`
- b `<class>`
- c `<required>`
- d `<type>`

2. What is the name of the tag library descriptor element that declares that an attribute can have a request-time expression as its value?

[ \_\_\_\_\_ ]

3. Consider the following code in a JSP page.

```
<% String message = "Hello "; %>

<test:world>
 How are you?
 <% message = message + "World! "; %>
</test:world>

<%= message %>
```

If `doStartTag()` returns `EVAL_BODY_BUFFERED` and `doAfterBody()` clears the buffer by calling `bodyContent.clearBody()`, what will be the output of the above code? (Select one)

- a Hello
- b Hello How are you?
- c Hello How are you? World!
- d Hello World!
- e How are you World!

4. Which of the following interfaces are required at a minimum to create a simple custom tag with a body? (Select one)

- a `Tag`
- b `Tag` and `IterationTag`
- c `Tag`, `IterationTag`, and `BodyTag`
- d `TagSupport`
- e `BodyTagSupport`

5. At a minimum, which of the following interfaces are required to create an iterative custom tag? (Select one)

- a `Tag`
- b `Tag` and `IterationTag`



- c** Tag, IterationTag, and BodyTag
  - d** TagSupport
  - e** BodyTagSupport
6. Which of the following methods is never called for handler classes that implement only the Tag interface? (Select one)
- a** setParent()
  - b** doStartTag()
  - c** doAfterbody()
  - d** doEndTag()
7. Which of the following is a valid return value for doAfterBody()? (Select one)
- a** EVAL\_BODY\_INCLUDE
  - b** SKIP\_BODY
  - c** EVAL\_PAGE
  - d** SKIP\_PAGE
8. Which element would you use in a TLD to indicate the type of body a custom tag expects?
- [ \_\_\_\_\_ ]
9. If the doStartTag() method returns EVAL\_BODY\_INCLUDE one time and the doAfterBody() method returns EVAL\_BODY\_AGAIN five times, how many times will the setBodyContent() method be called? (Select one)
- a** Zero
  - b** One
  - c** Two
  - d** Five
  - e** Six
10. If the doStartTag() method returns EVAL\_BODY\_BUFFERED one time and the doAfterBody() method returns EVAL\_BODY\_BUFFERED five times, how many times will the setBodyContent() method be called? Assume that the body of the tag is not empty. (Select one)
- a** Zero
  - b** One
  - c** Two
  - d** Five
  - e** Six

11. How is the `SKIP_PAGE` constant used? (Select one)
- a `doStartTag()` can return it to skip the evaluation until the end of the current page.
  - b `doAfterBody()` can return it to skip the evaluation until the end of the current page.
  - c `doEndTag()` can return it to skip the evaluation until the end of the current page.
  - d It is passed as a parameter to `doEndTag()` as an indication to skip the evaluation until the end of the current page.
12. Which of the following can you use to achieve the same functionality as provided by `findAncestorWithClass()`? (Select one)
- a `getParent()`
  - b `getParentWithClass()`
  - c `getAncestor()`
  - d `getAncestorWithClass()`
  - e `findAncestor()`
13. Consider the following code in a tag handler class that extends `TagSupport`:
- ```
public int doStartTag()
{
    //1
}
```
- Which of the following can you use at `//1` to get an attribute from the application scope? (Select one)
- a `getServletContext().getAttribute("name");`
 - b `getApplication().getAttribute("name");`
 - c `pageContext.getAttribute("name", PageContext.APPLICATION_SCOPE);`
 - d `bodyContent.getApplicationAttribute("name");`
14. Which types of objects can be returned by `PageContext.getOut()`? (Select two)
- a An object of type `ServletOutputStream`
 - b An object of type `HttpServletOutputStream`
 - c An object of type `JspWriter`
 - d An object of type `HttpJspWriter`
 - e An object of type `BodyContent`
15. We can use the directive `<%@ page buffer="8kb" %>` to specify the size of the buffer when returning `EVAL_BODY_BUFFERED` from `doStartTag()`.
- a True
 - b False



CHAPTER 17

Developing “Simple” custom tag libraries

- 17.1 Understanding SimpleTags 353
- 17.2 Incorporating SimpleTags in JSPs 357
- 17.3 Creating Java-free libraries with tag files 364
- 17.4 Summary 371
- 17.5 Review questions 372

EXAM OBJECTIVES

- 10.4** Describe the semantics of the “Simple” custom tag event model when the event method (doTag) is executed; write a tag handler class; and explain the constraints on the JSP content within the tag.
(Sections 17.1 and 17.2)
- 10.5** Describe the semantics of the Tag File model; describe the web application structure for tag files; write a tag file; and explain the constraints on the JSP content in the body of the tag.
(Section 17.3)

So far, you've seen the interfaces for classic tag development (`Tag`, `IterationTag`, and `BodyTag`) and their associated event methods (`doStartTag()`, `doAfterBody()`, and `doEndTag()`). These constructs give you a great deal of flexibility in building custom tag libraries, but coding can be time-consuming and complex. The JSP 2.0 standard reduces the burden by providing an alternate means of creating your tag libraries: the simple tag model.

With this new methodology, you only need to keep track of one interface, `SimpleTag`, and a single event method, `doTag()`. This way, you can concentrate on Java code instead of directing the web container's operation. JSP 2.0 also gives you different options for processing body content and tag attributes.

In addition, the new specification provides a new means of library development with *tag files*. Tag file processing is similar to regular JSP tag processing, but doesn't use tag library descriptors or tag handlers. Instead, tag files are coded with regular JSP syntax, and can include script elements. These building blocks simplify the process of library creation and make it more modular.

In practical web development, you can choose whatever tag library development method you prefer. But for the SCWCD exam, you need to become familiar with each. So, having discussed the classic way of building custom tag libraries, let's explore the simple method.

17.1 UNDERSTANDING SIMPLETAGS

Building custom tag libraries with `SimpleTags` is similar to the process we described in chapter 16. You still need to create a Java tag handler, reference the class in a tag library descriptor (TLD), and include the TLD in your JSP.

The differences between classic and simple development concern the processing needed in a Java-based tag handler. `SimpleTag` classes use fewer methods, interface differently with body content, and have different implicit objects available. This section covers the principles behind `SimpleTags` and how they reduce the difficulty of building tag libraries.

17.1.1 A brief example

Before we get into the theory of `SimpleTags`, you can appreciate how easy they are to use by looking at example code. Listing 17.1 shows a Java tag handler that sends a message to the JSP output.

Listing 17.1 `SimpleTagExample.java`

```
package myTags;

import java.io.*;
import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class SimpleTagExample extends SimpleTagSupport
{
    public void doTag() throws JspException, IOException
```

```

    {
        getJspContext().getOut().print(
            "I can't believe it's so simple!"
        );
    }
}

```

That's it! There are no concerns with `SKIP_BODY`, `EVAL_BODY`, `EVAL_PAGE`, or any of the return values associated with classic tags. There's no need to consider different interfaces depending on body content or iterations. Instead, there's just one method, `doTag()`, and a single line of code to send output to the JSP.

The web container can access this class through a tag library descriptor in the same way as a `Tag` or `BodyTag` class. For example, the snippet below matches the `SimpleTagExample` class with a tag name called "message":

```

<taglib>
..
<tag>
  <name>message</name>
  <tag-class>myTags.SimpleTagExample</tag-class>
  <body-content>empty</body-content>
  <description>Sends a message to the JSP</description>
</tag>
..
</taglib>

```

Then, the "message" tag can be inserted into a JSP page just as with the classic model.

The superclass of `SimpleTagExample`, `SimpleTagSupport`, makes all of this possible. To see why this is the case, you need to understand both it and its interface, `SimpleTag`. In particular, we'll present the methods contained in `SimpleTag` and its life cycle.

17.1.2 Exploring SimpleTag and SimpleTagSupport

In the classic method of building custom tag libraries, Java classes implement the `BodyTag` interface if body content needs processing, the `IterationTag` interface if multiple operations are required, or the `Tag` interface if neither is necessary. With the simple model, the `SimpleTag` interface can be used in all three cases. The relationship between these interfaces is shown in figure 17.1.




Figure 17.1
JSP tag
interfaces

The JSP 2.0 specification also provides a new adapter class for tag classes. Just as classic tag handlers extend `BodyTagSupport` or `TagSupport`, classes in the simple model can extend `SimpleTagSupport`. This class contains all of the methods needed to implement `SimpleTag`, and provides additional methods for extracting information from the web container.

Using the `SimpleTag` interface—methods and life cycle

Like those of `Tag` and `BodyTag`, the methods in the `SimpleTag` interface serve two purposes. First, they allow you to transfer information between your Java class and the JSP. Second, they are invoked by the web container to initialize `SimpleTag` operation. Table 17.1 lists these methods with descriptions.

Table 17.1 Methods of the `SimpleTag` interface

| Name | Description |
|-------------------------------|---|
| <code>setJspContext ()</code> | Makes the <code>JspContext</code> available for tag processing |
| <code>setParent ()</code> | Called by the web container to make the parent tag available |
| <code>setJspBody ()</code> | Makes the body content available for tag processing |
| <code>doTag ()</code> | Called by the container to begin <code>SimpleTag</code> operation |
| <code>getParent ()</code> | Called by the Java class to obtain its parent <code>JspTag</code> |

These methods are listed in the order that they are invoked in the `SimpleTag` life cycle, which has three main steps.

Step 1 Initialize the information associated with the `SimpleTag`

After the web container creates an instantiation of the `SimpleTag` class, it calls the `setJspContext ()` method. This method returns an instance of the `JspContext` class, which is the superclass of `PageContext`—the object returned by the `setPageContext ()` method in the `Tag` or `BodyTag` interface. Like the `PageContext`, the `JspContext` allows your Java class to access scoped attributes and implicit variables.

Most of the `JspContext` methods are similar to those in the `PageContext` class, but there are a few different methods that can be very useful. First, as shown in the example, the `getOut ()` method returns a `JspWriter` that you can use to send information to the JSP output stream. Also, there are two methods, `getExpressionEvaluator ()` and `getVariableResolver ()`, that allow you to access the Expression Language handling capability of the container. An important thing to keep in mind is that, while the `PageContext` relies on J2EE servlet processing, the `JspContext` class is meant to be technology-neutral and able to interface with different packages or languages.

After the `JspContext` has been initialized for the `SimpleTag`, the web container calls `setParent ()`. This method is invoked only if the `SimpleTag` is

surrounded by another set of tags. Because `setParent()` returns a `JspTag` object, the returned parent tag can implement the `Tag`, `BodyTag`, `IterationTag`, or the `SimpleTag` interface.

Step 2 Make body content available for `SimpleTag` processing

If there is any JSP code inside the tags, the web container invokes `setBody()` to make it available for the Java class. The method's return type is `JspFragment`. This class will be fully discussed later in this chapter, but for now, it is important to understand that a `JspFragment` contains regular JSP code (HTML, XML, tags, text) *without scripts*. So you can't include JSP declarations, expressions, or scriptlets inside `SimpleTags`. But EL terms can be added to the `JspFragment`.

Step 3 Invoke `doTag()`

The `doTag()` method of the `SimpleTag` interface combines the functions of the `Tag`'s `doStartTag()`, `doAfterBody()`, and `doEndTag()` methods. It doesn't return any values, and when it finishes, the web container returns to its previous processing tasks. Instead of calling special methods, you can control all of the iteration and body processing with regular Java commands.

It is important to understand why the `SimpleTag` interface is able to streamline the development of custom tag libraries. The reason has to do with JSP scripts. A great deal of the extra processing performed by a classic tag occurs because of the need to keep track of JSP scripts in the page. For example, if the tag body relies on a JSP variable declaration, then the tag processing needs to be able to access that variable.

With `SimpleTags`, this isn't an option. When the web container processes `SimpleTags`, it doesn't take JSP scripts into account. This makes for simpler coding and faster operation, but you need to keep this constraint in mind—both for the exam and your own web development.

Using `SimpleTagSupport`

The `SimpleTagSupport` class allows you to implement the `SimpleTag` interface without having to code each of its methods by yourself. Instead, it provides for each of the methods mentioned above, and three others, which are listed in table 17.2.

Table 17.2 Additional methods provided by the `SimpleTagSupport` class

| Name | Description |
|--------------------------------------|--|
| <code>getJspContext()</code> | Returns the <code>JspContext</code> for processing in the tag |
| <code>getJspBody()</code> | Returns the <code>JspFragment</code> object for body processing in the tag |
| <code>findAncestorWithClass()</code> | Returns the ancestor tag with the specified class |

These methods are similar to those in the `TagSupport` and `BodyTagSupport` classes, with two exceptions. First, the first two methods return a `JspContext` and a `JspFragment` instead of a `PageContext` and a `BodyContent` object. Second,

SimpleTagSupport leaves out many of the methods in TagSupport and BodyTagSupport that deal with tag processing, such as `release()`.

Now that you've seen how the SimpleTag interface and the SimpleTagSupport class functions, you can appreciate why Sun included them in the new JSP specification. In the next section, we will use these data structures to build practical JSPs.

Quizlet

- Q:** Which of the following methods aren't immediately available for a subclass of SimpleTagSupport?
- a** `getJspBody()`;
 - b** `getJspContext().getAttribute("name")`;
 - c** `getParent()`;
 - d** `getBodyContent()`;
- A:** The answer is option d. The `getBodyContent()` method is provided by the BodyTagSupport class, and returns a BodyContent object. Instead, SimpleTagSupport invokes the `getJspBody()` method, which returns a JspFragment.

17.2 INCORPORATING SIMPLETAGS IN JSPS

The process of building a SimpleTag library and using its JSP tags is similar to that for classic tag libraries, but there are important differences between the two. In particular, SimpleTags process tag attributes and body content differently than Tag, IterationTag, or BodyTag classes. In this section, we'll make these characteristics apparent by building a tag library and JSP for calculating square roots.

Each SimpleTag class performs its main processing inside the `doTag()` method, but the structure of the class also depends on attribute tags and body content. To present these classes, we'll proceed from the simple to complex. This means starting with an empty SimpleTag.

17.2.1 Coding empty SimpleTags

Empty SimpleTag classes are used to send static information to the JSP. In this case, we'll start with a short class that sends a simple mathematical expression to a JspWriter. This may seem trivial, but we'll add more as we explore the SimpleTag interface in greater depth.

Listing 17.2 presents `MathTag.java`, located in the `myTags` package.

Listing 17.2 MathTag.java

```
package myTags;

import java.io.*;
import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;
```



```

public class MathTag extends SimpleTagSupport
{
    int x = 289;

    public void doTag() throws JspException, IOException
    {
        getJspContext().getOut().print(
            "The square root of " + x +
            " is " + Math.sqrt(x) + "."
        );
    }
}

```

After the web container creates an instance of `MathTag`, it will make the `JspContext` available. Since there are no nested tags or body content, it will then invoke the `doTag()` method directly.

A tag library descriptor is needed to tell the web container how to match the `MathTag` class with its JSP tag. Listing 17.3 shows `MathTag.tld`, which also informs the web container that the `MathTag` class has no attributes and doesn't process body content.

Listing 17.3 `MathTag.tld`

```

<!DOCTYPE taglib PUBLIC
    "-//Sun Microsystems, Inc.//DTD JSP Tag Library 1.2//EN"
    "http://java.sun.com/j2ee/dtd/web-jspstaglibrary_1_2.dtd">
<taglib>
    <uri>www.manning.com/scwcd/math</uri>
    <tlib-version>1.0</tlib-version>
    <jsp-version>2.0</jsp-version>
    <tag>
        <name>sqrt</name>
        <tag-class>myTags.MathTag</tag-class>
        <body-content>empty</body-content>
        <description>
            Sends a math expression to the JSP
        </description>
    </tag>
</taglib>

```

In this example, we'll use the URI specified in the TLD instead of updating the deployment descriptor. We lose the capability of centralized library referencing, but the code is simpler and we can create the JSP directly. The JSP itself, shown in listing 17.4, tells the web container the library's URI, and then uses the empty tag to display the math statement.

Listing 17.4 math.jsp

```
<%@ taglib prefix="math" uri="www.manning.com/scwcd/math" %>
<html><body>
  <math:sqrt />
</body></html>
```

The result, shown in figure 17.2, shows that the JSP works as desired.

Now that you've seen how to build a basic SimpleTag-based JSP, we can add more powerful features. Next, we'll add dynamic attributes to the SimpleTag.

The square root of 289 is 17.0.

Figure 17.2 Static output from an empty SimpleTag instance

17.2.2 Adding dynamic attributes to SimpleTags

In the previous chapter, we showed how tags implementing Tag, IterationTag, and BodyTag process attributes with JSP 1.x. By adding a setter method for the given attribute (setXYZ() for the XYZ attribute), you can incorporate its value into your Java class. Then, you need to update the TLD to tell the web container what attributes it should accept.

This process remains the same using the simple model of tag library creation, but what if you don't know the name of the tag's attributes? What if you don't know how many there are? With JSP 1.x, you face serious problems. But JSP 2.0 provides the DynamicValues interface, which allows you to process multiple, unspecified attributes with a single method, setDynamicAttribute().

To show how this works, we're going to add static and dynamic attributes to the MathTag example. This time, the JSP will display a table of math functions whose entries are determined by the tag's attributes. Listing 17.5 updates MathTag.tld to tell the web container what kind of attributes to expect in the JSP.

Listing 17.5 MathTag.tld (Updated)

```
<!DOCTYPE taglib PUBLIC
  "-//Sun Microsystems, Inc.//DTD JSP Tag Library 1.2//EN"
  "http://java.sun.com/j2ee/dtd/web-jsptaglibrary_1_2.dtd">
<taglib>
  <uri>www.manning.com/scwcd/math</uri>
  <tlib-version>1.0</tlib-version>
  <jsp-version>2.0</jsp-version>
  <tag>
    <name>functions</name>
    <tag-class>myTags.MathTag</tag-class>
    <body-content>empty</body-content>
    <attribute>
      <name>num</name>
```

```

        <required>true</required>
        <rtexprvalue>true</rtexprvalue>
    </attribute>
    <dynamic-attributes>
        true
    </dynamic-attributes>
    <description>
        Sends a math expression to the JSP
    </description>
</tag>
</taglib>

```

The body content remains empty, but there are now two elements for attributes. The first, `<attribute>`, tells the web container about a static attribute called `num`, which is required and can be dynamically calculated at runtime. The second, `<dynamic-attributes>`, tells the web container that the tag may contain other attributes besides `num`, and it should create a `Map` to hold their names and values.

Listing 17.6 updates the `MathTag.java` code to reflect the new attribute processing. This class creates a `String` called `output` that is updated by the `setDynamicAttribute()` method. The web container calls this method each time it encounters an attribute not mentioned in the TLD. Once it finishes reading the attributes, it invokes `doTag()`, which sends the `String` to the JSP for display.

Listing 17.6 `MathTag.java` (Updated)

```

package myTags;

import java.io.*;
import javax.servlet.jsp.*;
import javax.servlet.jsp.tagext.*;

public class MathTag extends SimpleTagSupport
    implements DynamicAttributes
{
    double num = 0;
    String output = "";

    public void setNum(double num)
    {
        this.num = num;
    }

    public void setDynamicAttribute(String uri, String localName,
        Object value ) throws JspException
    {
        double val = Double.parseDouble((String)value);
        if (localName == "min")
        {
            output = output + "<tr><td>The minimum of "+num+" and "+
                val + "</td><td>" + Math.min(num, val) + "</td></tr>";
        }
    }
}

```

```

    }
    else if (localName == "max")
    {
        output = output + "<tr><td>The maximum of "+num+" and "+
            val + "</td><td>" + Math.max(num, val) + "</td></tr>";
    }
    else if (localName == "pow")
    {
        output = output + "<tr><td>"+num+" raised to the "+val+
            " power"+"</td><td>"+Math.pow(num, val)+"</td></tr>";
    }
}

public void doTag() throws JspException, IOException
{
    getJspContext().getOut().print(output);
}
}

```

After the web container initializes the `JspContext` and parent tag (if necessary), it processes the tag's attributes. If the attribute is static, such as `num`, it calls the setter method with the name of the attribute—`setNum()` in our example. If the attribute isn't mentioned in the TLD, then it is dynamic, and the container invokes `setDynamicAttribute()`.

Since this method does most of the work in the example, it's important to learn how it functions. It provides three items of information: a `URI String` representing the attribute's namespace, a `String` containing its name, and the `Object` containing its value. After converting the value into a double, `MathTag` continues processing according to the attribute's name. If the name is `min`, `max`, or `pow`, then `output` is updated with a new table row.

Listing 17.7 shows how this tag is coded in the JSP. Note that the static attribute, `num`, needs to be included *first*. This way, its value will be available when the rest of the attributes are processed.

Listing 17.7 `MathTag.jsp` (Updated)

```

<%@ taglib prefix="math" uri="www.manning.com/scwcd/math" %>
<html><body>
    Math Functions:<p>
    <table border="1">
    <math:functions num="{3*2}" pow="2" min="4" max="8"/>
    </table>
</body></html>


```

The use of HTML tables and EL may seem unnecessary. But in the exam, Sun will make the JSP code as complicated as possible. So make sure you have a solid grasp of both topics.

We can specify the `num` attribute with EL because the TLD sets its `<rtexprvalue>` tag to `true`. But the dynamic attributes don't have this option. If you try to use EL to set the values of `pow`, `min`, and `max`, you'll get an error.

Figure 17.3 shows the JSP's output.

Although the `DynamicAttributes` interface is new, you can still extend its usage to the classic `Tag`, `IterationTag`, and `BodyTag` classes. But as we've shown, building `SimpleTags` requires less code and complexity. Let's finish our discussion on this topic by looking at how `SimpleTags` process body content.

The screenshot shows a web page titled "Math Functions:". Below the title is a table with three rows. The first row contains the text "6.0 raised to the 2.0 power" and the value "36.0". The second row contains the text "The minimum of 6.0 and 4.0" and the value "4.0". The third row contains the text "The maximum of 6.0 and 8.0" and the value "8.0".

Math Functions:	
6.0 raised to the 2.0 power	36.0
The minimum of 6.0 and 4.0	4.0
The maximum of 6.0 and 8.0	8.0

Figure 17.3 Dynamic output from a `SimpleTag` implementing `DynamicAttributes`

17.2.3 Processing body content inside `SimpleTags`

In the classic model, `BodyTag` classes acquire the text and code inside their JSP tags by invoking `getBodyContent()`. This returns a `BodyContent` object that can be converted into a `String` or a `Reader`. This means that you can parse through the body and alter it if needed.

These options aren't available with `SimpleTags`. If you want to access the body, the `getJspBody()` method will return a `JspFragment`. This object only has two methods. The first, `getJspContext()`, returns the `JspContext` associated with the fragment. The second, `invoke()`, executes the JSP code and directs its output to the `JspWriter`. Neither method allows you to access and manipulate the body's contents as you can in the classic model.

Further, a `SimpleTag`'s body must not contain scripts—no declarations, expressions, or scriptlets. So it is invalid for a `SimpleTag`'s tag library descriptor to specify its `<body-content>` as `JSP`. Therefore, if you want to process a `SimpleTag`'s body, you need to set its `<body-content>` to `tagdependent` or `scriptless`. This is an important constraint to remember.

Because `SimpleTag` development doesn't add any new capabilities for processing body content, we will present code snippets instead of a new example. The code below shows how the `doTag()` method acquires the `SimpleTag`'s body and directs it to the JSP for display:

```
public void doTag() throws JspException, IOException
{
    getJspContext().getOut().print(output);
    getJspBody().invoke(null);
}
```

Note that the `invoke()` method requires an argument specifying the `JspWriter` that will receive the `JspFragment`'s output. In this case, the null argument directs the output to the `JspWriter` returned by `getJspContext().getOut()`.

As shown here, the only change required in the TLD is the `<body-content>`. Since JSP is invalid and `empty` is erroneous, we'll set the value to `tagdependent`:


```
<tag>
  <name>functions</name>
  <tag-class>myTags.MathTag</tag-class>
  <body-content>tagdependent</body-content>
  ...
</tag>
```

For the JSP, we'll add a new row to the table by including it as body content. This is shown in the code that follows. Because the `SimpleTag` executes the `JspFragment` last, this will be the last row of the table.

```
<math:functions num="{3*2}" pow="2" min="4" max="8">
  <td>This is the body of the SimpleTag.</td>
</math:functions>
```

Figure 17.4 shows the output of the new JSP.

Although the `SimpleTag` class reduces the amount of Java needed to create custom tag libraries, it still requires building and compiling Java classes. To the creators of the JSP 2.0 specification, this is still too much work. So they came up with an even simpler way of building tag libraries. With *tag files*, you don't need TLDs or Java at all! In the next section, we'll see how this new method works.



Math Functions:	
6.0 raised to the 2.0 power	36.0
The minimum of 6.0 and 4.0	4.0
The maximum of 6.0 and 8.0	8.0
This is the body of the SimpleTag.	

Figure 17.4 Output updated with `SimpleTag` body content

Quizlet

- Q:** What is the main difference between a TLD for `SimpleTags` and a TLD for a classic `Tag`?
- A:** `SimpleTag` TLDs cannot set their `<body-value>` elements equal to JSP. This is because a `SimpleTag` cannot process script elements in body content.

17.3 CREATING JAVA-FREE LIBRARIES WITH TAG FILES

JSTL and EL reduce the amount of Java in a JSP and `SimpleTags` reduce the amount of Java in a tag handler. But tag files remove the need for Java programming altogether. As long as you understand the JSP syntax, you can now build custom tags for your pages.

We'll begin our discussion of tag files with a simple example. Next, we'll cover the directives that enable you to communicate information to the web container. Finally, we'll look at fragments and how they are processed with tag file actions.

17.3.1 Introducing tag files

At its simplest, a tag file is a file made up of JSP code that has a `.tag` or `.tagx` extension. It can include EL expressions, directives, and custom and standard tags. Unlike `SimpleTag` JSPs, tag files can also contain script elements. In fact, the only JSP elements that can't be used in tag files are page attributes.

To see how tag files work, let's start with a simple example. Listing 17.8 presents `example.tag`, which displays a sequence of six numbers.

Listing 17.8 `example.tag`

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<c:forTokens items="0 1 1 2 3 5" delims=" " var="fibNum">
  <c:out value="\${fibNum}"/>
</c:forTokens>
```

Listing 17.8 contains just regular JSP code using the JSTL `forTokens` action and an EL expression. The JSP presented in listing 17.9 accesses this tag and displays its output.

Listing 17.9 `example.jsp`

```
<%@ taglib prefix="ex" tagdir="/WEB-INF/tags" %>
<html><body>
  The first six numbers in the Fibonacci sequence are:
  <ex:example/>
</body></html>
```

The code may look trivial, but this new capability is important. The key is simplicity. You don't need a background in Java to create custom tags with tag files. You don't have to compile Java classes and keep track of their packages. You don't even need tag library descriptors. Thanks to the JSP 2.0 specification, any developer of presentation logic can now create a custom tag library.

The new specification also makes it simple to integrate tag files within a JSP. There are only two steps:

- 1 Add a `taglib` directive to the JSP with a `prefix` attribute and the `tagdir` attribute equal to `/WEB-INF/tags`.
- 2 Place a tag containing the prefix and the name of the tag file (without the extension) wherever you want the file's JSP code invoked.

This brings up an important question. If Java-based tags need TLDs to locate their classes, how do these tags locate their tag files? To answer this, we need to look at how the web container accesses and processes tag files.

17.3.2 Tag files and TLDs

In the example JSP above, the `tagdir` attribute is set to `/WEB-INF/tags`. This is necessary since the web container automatically looks there for tag files. Then, the container builds an implicit tag library and TLD for this directory and each subdirectory beneath it. The good news is that you don't have to create TLDs for tag files. The bad news is that your tag files *must* be in `/WEB-INF/tags/` or a subdirectory.

But if you deploy your tag files inside a JAR, the situation changes. In this case, you need to create a tag library descriptor for your files. This TLD is similar to regular TLDs, but instead of matching tags to tag handlers, it matches names of tag files to their paths.

To make this possible, tag file TLDs use `<tag-file>` elements in place of `<tag>` elements. The definition of a `<tag-file>` element is as follows:

```
<!ELEMENT tag-file (description?, display-name?,  
                    icon?, name, path, example?, tag-extension?) >
```

The only necessary subelements are `<name>`, which specifies the tag file name without its suffix, and `<path>`, which specifies the file's path from the archive's root. Therefore, `<path>` must begin with `/META-INF/tags`. Here is an example TLD for an archived tag file:

```
<taglib>  
...  
  <uri>www.manning.com/scwcd/example</uri>  
  <tag-file>  
    <name>example</name>  
    <path>/META-INF/tags/example.tag</path>  
  </tag-file>  
</taglib>
```

This TLD must be located in the `META-INF` directory and the tag file(s) must be placed in `META-INF/tags` or a subdirectory. An example directory structure is shown here:

```
META-INF/  
example.tld  
tags/  
example.tag
```


Since the tag file isn't located in or under `/WEB-INF/tags`, you can't use the `tagdir` attribute in the `taglib` directive. Instead, you need to specify the TLD's URI (`www.manning.com/scwcd/example`) using the `URI` attribute. For this example, the following JSP directive will tell the web container where to find the tag file's TLD:

```
<%@ taglib prefix="ex" uri="www.manning.com/scwcd/example" %>
```

Other important differences between tag file TLDs and tag TLDs concern the `<attribute>` and `<body-content>` elements. Tags furnish this information in their TLDs, but tag files can't. Instead, tag files use a special set of directives. They tell the web container how to process the tag file, and it is important to understand how they work.

17.3.3 Controlling tag processing with tag file directives

JSPs contain three different kinds of directives: `page`, `taglib`, and `include`. Tag files remove the `page` directive and add three more. The first, `variable`, creates and initializes a variable for use in tag processing. The second, `tag`, tells the web container how to process the tag file. The third, `attribute`, describes the attributes that can be used in the tag. We'll investigate each of these, and provide snippets of example code.

Creating JSP variables with the variable directive

In the previous chapter, we showed how to declare JSP variables in tag library descriptors by adding `<variable>` elements. Then, you can assign and display the variable with JSTL actions and EL expressions.

Tag files provide a similar capability with the `variable` directive. The attributes of this directive are the same as the `<variable>` subelements, using `scope` to define the variable's visibility, and `name-given` and `name-from-attribute` to provide the variable's name. The only difference is the `alias` attribute, which provides a local name for the variable when its real name is determined by an attribute value (using `name-from-attribute`).

As an example, if the tag file contains the directive

```
<%@ variable name-given="x" %>
```

then the JSP can set the variable's value with the JSTL action

```
<c:set var="x">
  Hooray!
</c:set>
```

and display this value inside the JSP with `#{x}`.

The important point about the `variable` directive is that you don't need to rely on script declarations to declare variables in a JSP. But this is a minor function. The `tag` directive accomplishes much more.

Using the tag directive in tag files

The first new directive, `tag`, works like the `page` directive in a JSP. It provides the web container with settings that apply to the entire file. Table 17.3 describes the attributes that can be specified within a tag file's `tag` directive.

Table 17.3 Tag file attributes within the `tag` directive

Name	Description
<code>body-content</code>	Similar to the TLD subelement—can be <code>empty</code> , <code>tagdependent</code> , or <code>scriptless</code> . Set to <code>scriptless</code> by default.
<code>description</code>	Optional <code>String</code> statement describing the tag file.
<code>display-name</code>	<code>String</code> used by XML tools. Set to the name of the tag file (without the extension by default).
<code>dynamic-attributes</code>	Tells the container to create a named <code>Map</code> to hold unspecified attributes and their values.
<code>example</code>	<code>String</code> providing an instance of the tag's usage.
<code>import</code>	Adds a class, interface, or package to the tag processing.
<code>isELIgnored</code>	Specifies whether EL constructs will be ignored.
<code>language</code>	Sets the programming language used in the tag file. "Java" by default.
<code>large-icon</code>	Path to the large image representing the tag.
<code>page-encoding</code>	Specifies the character encoding of the tag file.
<code>small-icon</code>	Path to the small image representing the tag.

One attribute that has no JSP counterpart is `dynamic-attributes`. This works like the TLD `<dynamic-attributes>` subelement, but instead of directing attributes to a Java method, the web container updates a local variable specified by the directive. For example, the tag file that follows uses the `tag` directive to send dynamic attribute data to `attrib`. This data is then displayed using the JSTL `forEach` action.

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<%@ tag dynamic-attributes="attrib" %>
<c:forEach items="{attrib}" var="att">
    ${att.value}<br>
</c:forEach>
```

The JSP code shown next accesses this tag file (`dynatt.tag`) and sets the names and values of the tag's attributes. When the JSP is invoked, it will display a list of these values.

```
<%@ taglib prefix="dyn" tagdir="/WEB-INF/tags" %>
<html><body>
    <dyn:dynatt first="first" second="second" third="third"/>
</body></html>
```

Now that you've learned how to specify dynamic attributes in tag files, it's important to understand how to add static attributes. This requires the `attribute` directive.

Adding static attributes with the `attribute` directive

Dynamic attributes provide flexibility, but if you already know your tag's attributes, you can inform the web container in advance with static attributes. Traditional tags have `<attribute>` subelements in TLDs for this purpose. But to set attributes in tag files, you need `attribute` directives.

The attributes associated with the `attribute` directive are similar to the subelements of the TLD's `<tag>` element. The `name` attribute provides identification, `required` informs the web container whether the attribute must be present, and `rtexprvalue` tells the container that the attribute's value can be determined at runtime.

A brief example will show how `attribute` directives are used in tag files. The following tag file snippet sends output to the JSP according to the value of `x`.

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<%@ attribute name="x" required="true" %>
<c:choose>
  <c:when test='${x == "yes"}'>
    Yippee!
  </c:when>
  <c:otherwise>
    Rats!
  </c:otherwise>
</c:choose>
```

Then, the following JSP snippet uses the tag file and sets its attribute:

```
<%@ taglib prefix="attr" tagdir="/WEB-INF/tags" %>
<html><body>
  <attr:statatt x="yes" />
</body></html>
```

The `attribute` directive also allows you to insert JSP code into static attributes by setting its `fragment` attribute equal to `true`. However, to process this fragment, you need to look outside tag file directives, and concern yourself with standard actions.

17.3.4 Processing fragments and body content with tag file actions

JSPs provide a set of standard actions to direct the web container's processing of the page. Tag files can use all of these, and provide two more. The first, `jsp:invoke`, makes use of the fragment declared in the `attribute` directive. The second, `jsp:doBody`, processes the tag's body content.

Manipulating fragments with the `jsp:invoke` action

`SimpleTag` classes retrieve body content by calling `getJspBody()`, which returns a `JspFragment`. Then, to direct the fragment's output to a `JspWriter`, the tag handler calls the fragment's `invoke()` method. This method's argument determines which `JspWriter` object will receive the fragment's output.

The `jsp:invoke` action performs essentially the same function as `invoke()`, but is used for attributes declared as fragments, not for body content. Also, this action can do more with the `JspFragment` than just directing it to a `JspWriter`. It can convert the fragment to a `String` or a `Reader` object. However, just as with `SimpleTags`, tag files *cannot* process script elements (declarations, expressions, scriptlets) inside body content.

Table 17.4 lists and describes the attributes needed to configure this action in tag files.

Table 17.4 Tag file attributes within the `tag` directive

Name	Description
<code>fragment</code>	Identifies the <code>JspFragment</code> for processing.
<code>var</code>	Name of the <code>String</code> used to contain the <code>JspFragment</code> . Cannot be used with <code>varReader</code> .
<code>varReader</code>	Name of the <code>Reader</code> used to contain the <code>JspFragment</code> . Cannot be used with <code>var</code> .
<code>scope</code>	Scope of the stored <code>JspFragment</code> . Must be <code>page</code> , <code>request</code> , <code>session</code> , or <code>application</code> .

Of these attributes, only `fragment` is required in the `jsp:invoke` action. If neither `var` nor `varReader` is set, then the `JspFragment` will be directed to the default `JspWriter`. If one of `var` or `varReader` is set, but `scope` isn't, then the fragment's scope will be set to `page`.

Listing 17.10 presents an example tag file that uses the `jsp:invoke` action. First, it specifies a required attribute named `frag`, whose value will be contained in a `JspFragment`. Then, depending to the value of `proc`, it returns the fragment to the JSP as a `String` variable.

Listing 17.10 `invokeaction.tag`

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<%@ attribute name="frag" required="true" fragment="true"%>
<%@ attribute name="proc" required="true" %>
<c:if test='${proc == "yes"}'>
  <jsp:invoke fragment="frag"/>
</c:if>
```

Listing 17.11 presents the JSP needed to test this tag file. First, it incorporates the tag and sets its `proc` attribute to `yes`. Then, using the `<jsp:attribute>` action, it specifies a line of JSP code to serve as the value of the `frag` attribute.

Listing 17.11 `invokeaction.jsp`

```
<%@ taglib prefix="inv" tagdir="/WEB-INF/tags" %>
<html><body>
  <inv:invokeaction proc="yes">
    <jsp:attribute name="frag">
      Two + two = ${2+2}
    </jsp:attribute>
  </inv:invokeaction >
</body></html>
```

So far, you've seen all there is to know about setting and processing tag file attributes. Now, let's see how tag files make use of the information between the tags. To enable you to process this body content, tag files provide the `<jsp:doBody>` action.

Processing body content with the `jsp:doBody` action

The `jsp:doBody` action works like `jsp:invoke`, but it receives the tag's body instead of a fragment attribute. It contains the same attributes as `jsp:invoke`, except `fragment`. So, when a tag file receives body content, it can manipulate it in three ways: display it with the default `JspWriter`, send it to a variable with the `var` attribute, or store it as a `Reader` object with the `varReader` attribute.

The tag file in listing 17.12 processes the tag's body content according to the `att` attribute. When `att` equals "var," it will be stored within a variable, and when `att` equals "reader," it will be stored in a `Reader` object. If `att` isn't specified, the default `JspWriter` will display its output.

Listing 17.12 `bodyaction.tag`

```
<%@ taglib uri="http://java.sun.com/jstl/core_rt" prefix="c" %>
<%@ attribute name="att" required="true" %>
<c:choose>
  <c:when test='${att == "var"}'>
    <jsp:doBody var="bodyvar" scope="application"/>
  </c:when>
  <c:when test='${att == "reader"}'>
    <jsp:doBody varReader="bodyReader" />
  </c:when>
  <c:otherwise >
    <jsp:doBody />
  </c:otherwise>
</c:choose>
```

The JSP in listing 17.13 performs two tasks. First, it accesses the tag file and sets the `att` attribute to “var.” Then, using EL, it displays the variable containing the tag’s body content.

Listing 17.13 `bodyaction.jsp`

```
<%@ taglib prefix="bod" tagdir="/WEB-INF/tags" %>
<html><body>
  <bod:bodyaction att="var">
    This is the tag body.
  </bod:bodyaction >
  ${bodyvar}
</body></html>
```

In an earlier chapter, we showed how JSPs (`*.jsp`) can be converted into JSP documents (`*.jspx`) by using well-formed XML. The process of creating tag file documents (`*.tagx`) is very similar. The main task involves replacing tag file directives, such as `<%@ attribute ... %>`, with XML statements, such as `<jsp:directive.attribute ... />`.

This discussion ends our treatment of JSPs in general and tag library development in particular. As you can see, there are many different methods of creating tag libraries. If you are familiar with Java, you may want to use the simple model, but if you need to incorporate scripts, the classic method may be best. But if you prefer building tags with JSP code, you can’t do better than to use the tag files described here.

Quizlet

- Q:** In what directory should you place unarchived tag files? What directory for archived tag files?
- A:** All unarchived tag files should be placed in `/WEB-INF/tags` or a subdirectory underneath. All archived tag files should be placed in `/META-INF/tags` or a subdirectory.

17.4 SUMMARY

One of Sun’s primary goals in releasing JSP 2.0 was to simplify JSP development. To reduce the amount of Java in JSPs, they created Expression Language. To reduce the amount of code in tag handlers, they came up with the `SimpleTag` interface. Finally, to remove the need for tag handlers and TLDs altogether, they introduced tag files.

The advantages of `SimpleTags` over `Tags`, `BodyTags`, and `IterationTags` stem from their less-complex life cycles. With `SimpleTags`, there are no elaborate flowcharts or multiple event-based tag handler methods. After the web container performs its initialization, it only invokes one method, `doBody()`, which performs all of the `SimpleTag`’s processing. JSP 2.0 also provides the `SimpleTagSupport` adapter class for additional capabilities.

The drawback to `SimpleTag` operation involves its processing of body content. The body content is encapsulated in a `JspFragment` object, which cannot contain script elements such as declarations, expressions, or scriptlets. Further, `JspFragments` have no built-in mechanisms for converting body content into `Strings` or `Readers`.

Tag files are a fascinating addition to the traditional methods of tag library development. By specifying a precise directory for tag file location, the new JSP specification removes the need for tag library descriptors. Tag files still provide all of the information normally contained in TLDs, but they use directives and actions instead. The tag directive resembles the JSP's `page` directive, and `attribute` and `variable` resemble their corresponding TLD elements. Finally, the `jsp:doBody` and `jsp:invoke` actions allow you to process JSP code in the tag body and tag attributes, respectively.

17.5 REVIEW QUESTIONS

1. What method should you use in a `SimpleTag` tag handler to access dynamic variables?

- a `doTag()`
- b `setDynamicAttribute()`
- c `getParent()`
- d `setDynamicParameter()`

2. Which object does a `SimpleTag` tag handler use to access implicit variables?

- a `PageContext`
- b `BodyContent`
- c `JspContext`
- d `SimpleTagSupport`

3. Consider the following TLD excerpt:

```
<body-content>
  empty
</body-content>
<attribute>
  <name>color</name>
  <rtexprvalue>>true</rtexprvalue>
</attribute>
<dynamic-attributes>
  true
</dynamic-attributes>
```

If the name of the tag is `tagname` and its prefix is `pre`, which of the following JSP statements is valid?

- a `<pre:tagname color="yellow" size=${sizenum} />`
- b `<pre:tagname size="18" color="red"> </pre:tagname>`

- c** `<pre:tagname color="{colorname}" size="22" font="verdana"></pre:tagname>`
- d** `<pre:tagname color="green" size="30">font="Times New Roman"</pre:tagname>`
- e** `<pre:tagname color="{colorname}" size="18"></pre>`
4. If placed inside the body of a simple tag, which of the following statements won't produce "9"? (Select one)
- a** `#{3 + 3 + 3}`
- b** "9"
- c** `<c:out value="9">`
- d** `<%= 27/3 %>`
5. Which of the following methods need to be invoked in a SimpleTag to provide iterative processing? (Select one)
- a** `setDynamicAttribute()`
- b** `getParent()`
- c** `getJspBody()`
- d** `doTag()`
- e** `getJspContext()`
6. Which of the following values is invalid inside a SimpleTag's `<body-content>` subelement? (Select one)
- a** JSP
- b** `scriptless`
- c** `tagdependent`
- d** `empty`
7. Which of the following is a valid return value for the SimpleTag's `doTag()` method? (Select one)
- a** `EVAL_BODY_INCLUDE`
- b** `SKIP_BODY`
- c** `void`
- d** `EVAL_PAGE`
- e** `SKIP_PAGE`
8. Which tag file directive makes it possible to process dynamic attributes?
- a** `taglib`
- b** `page`
- c** `tag`
- d** `attribute`

9. Which of the following statements can't be used to access a tag file from a JSP? (Select one)
- a `<%@ taglib prefix="pre" uri="www.mysite.com/dir/" %>`
 - b `<%@ taglib prefix="pre" tagdir="/WEB-INF/tags" %>`
 - c `<%@ taglib prefix="pre" tagdir="/WEB-INF/tagfiles" %>`
 - d `<%@ taglib prefix="pre" tagdir="/WEB-INF/tags/myDirectory" %>`
10. Which tag file action processes `JspFragments` in tag attributes?
- a `taglib`
 - b `jsp:invoke`
 - c `tag`
 - d `jsp:doBody`
 - e `attribute`
11. Which `JspFragment` method is used to process body content in a `SimpleTag`? (Select one)
- a `invoke()`
 - b `getOut()`
 - c `getJspContext()`
 - d `getBodyContent()`
12. Which class provides an implementation of the `doTag()` method? (Select one)
- a `TagSupport`
 - b `BodyTagSupport`
 - c `SimpleTagSupport`
 - d `IterationTagSupport`
 - e `JspTagSupport`
13. In what directory shouldn't you place tag files? (Select one)
- a `/META-INF/tags/tagfiles`
 - b `/WEB-INF/`
 - c `/WEB-INF/tags/tagfiles/tagdir/taglocation`
 - d `/META-INF/tags/`
14. Which type of object is returned by `JspContext.getOut()`? (Select one)
- a `ServletOutputStream`
 - b `HttpServletOutputStream`
 - c `JspWriter`
 - d `BodyContent`

15. Which of the following methods does the web container call first to initiate a `SimpleTag`'s life cycle?
- a** `setJspContext ()`
 - b** `setParent ()`
 - c** `getJspContext ()`
 - d** `getJspBody ()`
 - e** `getParent ()`



CHAPTER 18

Design patterns

- 18.1 Design patterns: a brief history 377
- 18.2 Patterns for the SCWCD exam 382
- 18.3 Summary 400
- 18.4 Review questions 401

EXAM OBJECTIVES

- 11.1** Given a scenario description with a list of issues, select the design pattern that would best solve the issues. The list of patterns you must know are:
- Intercepting Filter
 - Model-View-Controller
 - Front Controller
 - Service Locator
 - Business Delegate
 - Transfer Object
- (Sections 18.1 and 18.2)
- 11.2** Match design patterns with statements describing potential benefits that accrue from the use of the pattern, for any of the following patterns:
- Intercepting Filter
 - Model-View-Controller
 - Front Controller
 - Service Locator
 - Business Delegate
 - Transfer Object
- (Sections 18.1 and 18.2)

INTRODUCTION

In our daily lives as designers and programmers, we are continuously developing our problem-solving skills. With each problem we encounter, we immediately start considering the different ways it can be solved, including successful solutions that we have used in the past for similar problems. Out of many possible solutions, we pick the one that best fits our application. By documenting this solution, we can reuse and share the information that we have learned about the best way to solve the specific problem.

Design patterns address the recurring design problems that arise in particular design situations and propose solutions to them. Design patterns are thus successful solutions to known problems. There are various ways to implement design patterns. These implementation details are called *strategies*.

In this chapter, we will introduce the following J2EE design patterns that are named in the exam objectives: Intercepting Filter, Model-View-Controller, Front Controller, Service Locator, Business Delegate, and Transfer Object.

18.1 DESIGN PATTERNS: A BRIEF HISTORY

A design pattern is an abstraction of a solution at a very high level. Many designers and architects have defined the term *design pattern* in various ways that suit the domain to which they apply the patterns. Further, they have divided the patterns into different categories according to their usage. Let's look at some of them before we go into the details of the J2EE patterns in the next section.

18.1.1 The civil engineering patterns

In the 1960s and '70s, Christopher Alexander, professor of architecture and director of the Center for Environmental Structure, along with his colleagues, wrote a number of books describing and documenting the principles of civil engineering from a layperson's point of view. Of them, one of the most widely known books is *A Pattern Language: Towns, Buildings, Constructions*. It provides practical guidance on how to build houses, construct buildings and parking lots, design a good neighborhood, and so forth. The book examines how these simple designs integrate with each other to create well-planned towns and cities.

As its title suggests, the book describes 253 patterns that are split into three broad categories: towns, buildings, and construction.

18.1.2 The Gang of Four patterns

Software designers extended the idea of design patterns to software development. Since features provided by the object-oriented languages, such as inheritance, abstraction, and encapsulation, allowed them to easily relate programming language entities to real-world entities, designers started applying those features to create common and reusable solutions to recurring problems that exhibited similar patterns.

Around 1994, the now famous Gang of Four (GoF)—Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides—documented 23 such software design

patterns in the book *Elements of Reusable Object-Oriented Software*. They classified the patterns at the very top level into three types of categories based on their purpose: creational, structural, and behavioral, as shown in table 18.1.

Table 18.1 The GoF categories of design patterns

Type of GoF pattern	Description
Creational	Creational patterns deal with the ways to create instances of objects. The objective of these patterns is to abstract the instantiation process and hide the details of how objects are created or initialized.
Structural	Structural patterns describe how classes and objects can be combined to form larger structures and provide new functionality. These aggregated objects can be either simple objects or composite objects themselves.
Behavioral	Behavioral patterns help us define the communication and interaction between the objects of a system. The purpose of these patterns is to reduce coupling between objects.

At the second level, they classified the patterns as falling into one of the two scopes: class or object. Thus, we have six types of patterns, as described in table 18.2.

Table 18.2 The GoF categories and scopes of design patterns

GoF Pattern Category	Brief Description	Examples
Creational		
Creational class	Creational class patterns use inheritance as a mechanism to achieve varying class instantiation.	Factory Method
Creational object	Creational object patterns are more scalable and dynamic compared to the class creational patterns.	Abstract Factory Singleton
Structural		
Structural class	Structural class patterns use inheritance to provide more useful program interfaces by combining the functionality of multiple classes.	Adapter (class)
Structural object	Structural object patterns create composite objects by aggregating individual objects to build larger structures. The composition of the structural object pattern can be changed at runtime, which gives us added flexibility over structural class patterns.	Adapter (object) Facade Bridge Composite
Behavioral		
Behavioral class	Behavioral class patterns use inheritance to distribute behavior between classes.	Interpreter
Behavioral object	Behavioral object patterns allow us to analyze the patterns of communication between interconnected objects, such as the included objects of a composite object.	Iterator Observer Visitor

18.1.3 The distributed design patterns

Though the GoF patterns served well in designing and developing object-oriented systems in both distributed and nondistributed environments, they were not created with the distributed nature of large-scale enterprise systems in mind. As the demands for more distributed enterprise applications grew, the architects felt the need to document the solutions to recurring problems as they experienced the same problem occurring over and over again. They started extending and refining the patterns over a larger scale and with a broader scope. One book that documents some 30 patterns at the architectural level is *CORBA Design Patterns*, by Thomas J. Mowbray and Raphael C. Malveau. Though the book focuses mainly on the Common Object Request Broker Architecture (CORBA), the patterns described are applicable to a wide range of distributed applications, including those that do not use CORBA.

The authors have categorized the design patterns at seven architectural levels:

- Global
- Enterprise
- System
- Application
- Macrocomponent
- Microcomponent
- Object

The authors have two basic arguments. The first is that software design involves making choices, such as which details of an object should be abstracted and what should be exposed, or which aspects of the objects are to be generalized and which one should be specialized. These decisions are based on the facts surrounding the problem at hand. The book terms such facts as *primal forces*, because they influence our choice of a particular pattern. The discussion on forces focuses on issues such as whether a pattern increases performance, whether it aids in enhanced functionality, or whether it helps to reduce complexity within modules.

The authors' second argument is that not all design patterns scale well at all seven architectural levels. Each pattern has a set of applicable levels, which is an important feature that must be considered when using the pattern.

18.1.4 The J2EE patterns

With the advent of the J2EE, a whole new catalog of design patterns cropped up. Since J2EE is an architecture in itself that comprises other architectures, including servlets, JavaServer Pages, Enterprise JavaBeans, and so forth, it deserves its own set of patterns specifically tailored to the various types of enterprise applications that the architecture addresses.

The book *Core J2EE Patterns: Best Practices and Design Strategies*, by Deepak Alur, John Crupi, and Dan Malks, describes the five tiers of the J2EE architecture:

- Client
- Presentation
- Business
- Integration
- Resource

The book then explains 15 J2EE patterns that are divided among three of the tiers: presentation, business, and integration.

The five tiers in J2EE

A *tier* is a logical partition of the components involved in the system. Each tier is loosely coupled with the adjacent tier. It is easier to understand the role of design patterns once we fully grasp the different tiers involved in a J2EE application. The J2EE architecture identifies the five tiers described in table 18.3.

Table 18.3 Tiers in the J2EE architecture

Tier Name	Description
Client	This tier comprises all the types of components that are clients of the enterprise application. Examples of client components are a web browser, a hand held device, or another application that accesses the services of the enterprise application remotely.
Presentation	This tier interfaces with the client tier and encapsulates the presentation logic. It accepts the requests, handles authentication and authorization, manages client sessions, delegates the business processing to the business tier, and presents the clients with the desired response. The components that make up this tier are filters, servlets, JavaBeans, JSP pages, and other utility classes.
Business	This tier is the heart of the enterprise application and implements the core business services. It is normally composed of the Enterprise JavaBeans components that handle all the business processing rules.
Integration	The job of this tier is to seamlessly integrate different types of external resources in the resource tier with the components of the business tier. The components that make up the integration tier use various mechanisms like JDBC, J2EE connector technology, or proprietary middleware to access the resource tier.
Resource	This tier comprises the external resources that provide the actual data to the application. The resources can either be data stores such as relational databases and file-based databases, or systems such as applications running on mainframes, other legacy systems, modern business-to-business (B2B) systems, and third-party services like credit card authorization services.

The J2EE pattern catalog

Table 18.4 lists the 15 patterns of J2EE, with a brief description of each.

Table 18.4 J2EE design patterns

Name (s)	Description
Presentation Tier	
Decorating Filter/ Intercepting Filter	An object that sits between the client and the web components. It pre-processes a request and post-processes the response.
Front Controller/ Front Component	An object that accepts all the requests from the client and dispatches or routes them to appropriate handlers. The Front Controller pattern may divide the above functionality into two different objects: the Front Controller and the Dispatcher. In that case, the Front Controller accepts all the requests from the client and does the authentication, and the Dispatcher dispatches or routes them to the appropriate handlers.
View Helper	A helper object that encapsulates data access logic on behalf of the presentation components. For example, JavaBeans can be used as View Helper patterns for JSP pages.
Composite View	A view object that is made up of an aggregate of other view objects. For example, a JSP page that includes other JSP and HTML pages using the include directive or the include action is a Composite View pattern.
Service To Worker	A kind of Model-View-Controller with the Controller acting as a Front Controller but with one important point: here the Dispatcher (which is a part of the Front Controller) uses View Helpers to a large extent and aids in view management.
Dispatcher View	A kind of Model-View-Controller with the controller acting as a Front Controller but with one important point: here the Dispatcher (which is a part of the Front Controller) does not use View Helpers and does very little work in view management. The view management is handled by the View components themselves.
Business Tier	
Business Delegate	An object that resides on the presentation tier and on behalf of other presentation tier components calls remote methods on the objects in the business tier.
Transfer Object/ Replicate Object	A serializable object for transferring data over the network.
Session Façade/ Session Entity Façade/ Distributed Façade	An object that resides in the business tier, acts as an entry point into the business tier, and manages the workflow of business service objects, such as session beans, entity beans, and Data Access Objects. The Session Facade itself is usually implemented as a session bean.
Aggregate Entity	An object (entity bean) that is made up of or is an aggregate of other entity beans.
Transfer Object Assembler	An object that resides in the business tier and creates Transfer Objects on the fly as and when required.
Value List Handler/ Page-by-Page Iterator/Paged List	An object that manages execution of queries, caching, and processing of results. Usually implemented as a Session Bean, serving a subset of the fetched result set to the client as and when needed.

continued on next page

Table 18.4 J2EE design patterns (continued)

Name (s)	Description
Service Locator	An object that performs the task of locating business services on behalf of other components in the tier. Usually present in the presentation tier, it is used by Business Delegates to look up business service objects.
Integration Tier	
Data Access Object	An object that talks to the actual underlying database and provides other application components. It serves as a clean, simple, and common interface for accessing the data, and for reducing the dependency of other components on the details of using the database.
Service Activator	An object that helps in processing of business methods asynchronously.

Of the design patterns listed in table 18.4, five—Intercepting Filter, Front Controller, Service Locator, Business Delegate, and Transfer Object—are mentioned in the exam objectives. They also cover the Model-View-Controller pattern, which is considered more of an architecture than a J2EE design pattern. In the rest of this chapter, we will examine these design patterns in depth. Discussing the other patterns is beyond the scope of this book, but they are frequently used as possible options in the single-choice questions of the exam. It will be helpful if you study table 18.4 so that during the exam you can eliminate the incorrect choices, thereby making it easier for you to select the right answer.

18.2 PATTERNS FOR THE SCWCD EXAM

Table 18.5 lists the most important J2EE design patterns that you need to know for the exam.

Table 18.5 Patterns required for the SCWCD exam

Name	Other Name(s)	Tier
Intercepting Filter	Decorating Filter	Presentation
Model-View-Controller		Presentation
Front Controller	Front Component	Presentation
Service Locator		Business
Business Delegate		Business
Transfer Object	Transfer Object, Replicate Object	Presentation

But before we discuss each of these, we want to describe the structure behind their documentation. This structure is called a *template*.

18.2.1 The pattern template

During the process of analyzing a problem, designing a solution, creating an action plan, and implementing the ideas, one of the most important tasks is to document

each and every aspect of the process. This facilitates preserving the ideas in a systematic manner so that they can be read, understood, and reused by others.

Design patterns are documented using a pattern template. The template is made up of headings; each heading explains a different aspect of the pattern, such as the cause of the problem, the situation in which the problem can occur, and the possible facts to look for. Different organizations and pattern catalog writers use different sets of headers and, therefore, their templates vary. However, their goal is the same: systematic documentation of the causes of the problems, the solutions provided by the patterns, their consequences and implications, examples, related patterns, and so forth.

So before we start explaining the design patterns required for the exams, let's first look at the template that we have used to explain the design patterns in the following sections. We have kept the template large enough to cover different aspects under different headings, while at the same time keeping it small enough to maintain simplicity. Also, wherever possible, we have avoided long paragraphs of text and used a bulleted-points approach while presenting the facts to make it easier to remember and reference. The following subsections describe the template headers.

Context

This section describes the context or the situation in which the problems can occur and the given pattern that we can apply.

Patterns, if not used wisely or if applied in the wrong context, can turn out to be anti-patterns. Various facts have to be sorted out, and the pros and cons of the pattern have to be evaluated against these facts before applying the pattern. Therefore, it is important to understand the situation in which the pattern may be applicable.

Problem

This section provides a common, day-to-day problem that we face in the given context. Because so many developers and programmers have experienced this problem, designers have developed standard solutions for it.

In many books or tutorials, you may find a separate heading called *Intent* to describe the intent of the pattern. The *Intent* heading is suitable where the primary aim of the pattern is to enhance the design of the overall system rather than attacking a particular problem. In the patterns described below, we have provided a *Problem* heading instead, and the intent of the pattern is to solve the given problem.

Example

This section provides a simple example of the specified problem type.

Facts or forces to consider

Other than the given situation and the given type of problem, there may be various facts that need to be considered before using design patterns. These facts can vary

largely from system to system, depending on the requirements. They are also referred to as *forces* in design pattern jargon, because they influence whether or not a design pattern is applicable in the given scenario. In some cases where more than one pattern can solve the problem, these forces can be helpful in determining the choice of using one pattern over the other. You can think of facts and forces as observations or requirements of the problem that you will take into consideration when you are evaluating the candidate design patterns for this particular situation.

Solution

This section describes a typical solution that has been tried, tested, and proven by experienced developers to solve the problem, taking into account the given facts or forces. The problem-solution pair forms a design pattern. However, the solution is really a high-level design guideline that we can follow. The actual implementation of the solution may vary from context to context, and each implementation of the solution is referred to as a *strategy* of the solution in the J2EE jargon. Thus, each design pattern can be implemented by applying different strategies depending on the context of the problem and the facts or forces surrounding the problem.

We have not provided a separate section of strategies, and we have discussed only the most common of the strategies for the patterns, since describing all the strategies is outside the scope of this book.

Consequences/implications

This section describes the advantages, aftereffects, and hidden problems associated with patterns. Every pattern has inherent side effects, and may solve one problem but generate another. We need to weigh the pros and cons of the pattern and find a balance between the two.

Category

This section indicates which of the three categories the pattern falls into: creational, structural, or behavioral, with a brief description.

Points to remember

These are the points we must take into account when considering a particular pattern as a solution. In the following discussions of the J2EE design patterns, we have provided a list of words, phrases, and terms that are related to the pattern, and which may help you to prepare for the exam. When you read a problem statement, see if you can identify any of the terms or phrases listed in this section in the statement. Another tactic is to restate the problem (without changing the meaning) and find any of these terms that match. This will aid you in selecting or eliminating the options in exam questions.

18.2.2 The Intercepting Filter

Context

Client requests may have many different processing needs, and the system must be able to meet them:

- The system receives requests using multiple protocols, such as HTTP, FTP, or SMTP.
- The system must authorize or authenticate some requests, while handling others directly.
- The system needs to add or remove information from requests or responses before further processing.

Problem

Requests and responses need to be pre- and post-processed in a simple, organized manner. Further, this processing needs to be performed without affecting the operation of the rest of the system, and must use standard interfaces in order to be added and replaced.

Example

If you have to build an enterprise system that doubles as a web portal and corporate intranet, you need to ensure that different client requests are handled appropriately. Not only must the system block requests for secure resources, it must also check for trusted IP addresses. Further, it must ensure that responses and requested resources are compressed/decompressed and encrypted/decrypted as needed.

Facts or forces to consider

In the context and the problem presented above, we observe that there are three characteristics of the desired system, as follows:

- 1 Centralized, standardized pre- and post-processing of incoming requests and outgoing responses.
- 2 Provide authentication, transformation, and conversion of resources as needed.
- 3 Perform function with minimal effect on external code processing

Solution

The solution is to separate request/response processing from the rest of the system by creating filters. These objects, more fully described in chapter 7, provide a standardized, replaceable means of altering requests and responses as needed.

Consequences/implications

- The standard interfaces associated with intercepting filters allows you to stack them for multiple processing tasks. Further, this modularity allows you to test filter processing without the full system present.
- A centralized location for request/response processing is important, but must not create a bottleneck. This may require distributed processing.
- The standard filter mechanisms allow you to dynamically insert or remove filters during application deployment.

Category

Because intercepting filters reduce coupling between request processing and the rest of the application, they are considered part of the behavioral category. They function by managing the information sent from the client to the application layer, and then control the response processing. By providing a single point-of-contact for transforming requests and responses, they reduce the confusion associated with having many different interfaces for different client/server relationships.

Points to remember

The Intercepting Filter design pattern consists of multiple `javax.servlet.Filter` objects that process requests and responses in a modular fashion. These filters are described in greater technical depth in chapter 7.

Pay special attention to the following words, phrases, or terms appearing in the problem statement:

- Central object for processing requests and responses
- Resource management, particularly regarding security, compression, and encryption
- Standard interface allows layered processing
- Intercepting filters deal with request processing, *not* with the application's appearance or dispatching mechanisms

18.2.3 Model-View-Controller (MVC)

Context

In systems involving user interfaces, the following situation typically arises:

- The system has to accept data from the user, update the database, and return the data to the user at a later point in time.
- There are several ways in which the data can be accepted from and presented to the system users.
- Data fed to the system in one form should be retrievable in another form.

Problem

If the system deploys a single component that interacts with the user as well as maintains the database, then a requirement to support a new type of display or view will necessitate the redesign of the component.

Example

Suppose a bank provides online stock trading facilities. When the user is logged into the site, the web application allows the user to view the rates of the stocks over a period of time in various ways, such as a bar graph, a line graph, or a plain table. Here, the same data that represents the rates of the stocks is viewed in multiple ways, but is controlled by a single entity, the web application.

Facts or forces to consider

In the context and the problem presented above, we observe the following facts:

- There are three tasks to be done:
 - 1 Manage the user's interaction with the system.
 - 2 Manage the actual data.
 - 3 Format the data in multiple ways and present it to the user.
- Thus, a single component that does all the tasks can be split into three independent components.
- All three tasks can then be handled by different components.

Solution

The solution is to separate the data presentation from the data maintenance and have a third component that coordinates the first two. These three components are called the Model, the View, and the Controller, and they form the basis of the MVC pattern. Figure 18.1 shows the relationship between the components of the MVC pattern.

Here are the responsibilities of the three MVC components:

- *Model*—The Model is responsible for keeping the data or the state of the application. It also manages the storage and retrieval of the data from the data source. It notifies all the Views that are viewing its data when the data changes.
- *View*—The View contains the presentation logic. It displays the data contained in the Model to the users. It also allows the user to interact with the system and notifies the Controller of the users' actions.
- *Controller*—The Controller manages the whole show. It instantiates the Model and the View and associates the View with the Model. Depending on the application requirements, it may instantiate multiple Views and may associate them with the same Model. It listens for the users' actions and manipulates the Model as dictated by the business rules.

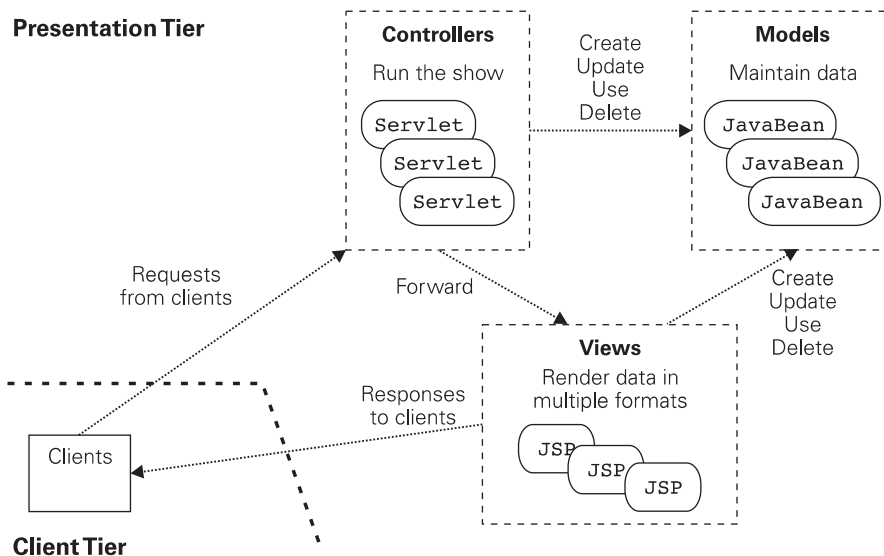


Figure 18.1 The Model-View-Controller pattern

Consequences/implications

- Separating the data representation (Model) from the data presentation (View) allows multiple Views for the same data. Changes can occur in both the Model and View components independently of each other as long as their interfaces remain the same. This increases maintainability and extensibility of the system.
- Separating the application behavior (Controller) from the data presentation (View) allows the controller to create an appropriate View at runtime based on the Model.
- Separating the application behavior (Controller) from the data representation (Model) allows the users' requests to be mapped from the Controller to specific application-level functions in the Model.

Category

Although MVC involves communication between the Model, View, and Controller components, it is not a behavioral pattern because it does not specify how the three components should communicate. The MVC pattern only specifies that the structure of a system of components be such that each individual component take up one of the three roles—Model, View, or Controller—and provide functionality only for its own role. As such, it is a structural pattern.

In the J2EE world, MVC is thought of more as an architecture rather than a design pattern. Though it can be applied in any of the tiers, it is most suitably applied in the presentation tier.

Points to remember

The Model-View-Controller design pattern, consisting of three subobjects—Model, View, and Controller—is applicable in situations where the same data (Model) is to be presented in different formats (Views), but is to be managed centrally by a single controlling entity (Controller).

Pay special attention to the following words, phrases, or terms appearing in the problem statement:

- Separation of data presentation and data representation
- Provide services to different clients: web client, WAP client, etc.
- Multiple views, such as HTML or WML
- Single controller

18.2.4 Front Controller

Context

Managing an application for an individual user requires a number of considerations:

- Controlling the application's view and navigation
- Performing security processing to determine which resources and services are available for the user
- Activating system services according to user selection
- Locating and accessing available resources according to user selection

Problem

In many cases, the functions listed above are performed in a decentralized manner. Without a single point for view management, each different view mechanism will need to access services and resources separately. This combines the system's appearance and business logic, and removes the benefits of modularity and reusability associated with MVC design.

Example

In a web application that accepts credit card information, several steps must be completed:

- Browse the catalog.
- Add items to the shopping cart.
- Confirm the checkout.
- Get the name and shipping address of the receiver.
- Get the name and billing address and the credit card information of the payer.

The different views associated with the store are decentralized. This means that they each control their navigation functions, as well as perform the tasks needed to dispatch the user's request.

Facts or forces to consider

In the context and the problem presented above, we observe the following facts:

- Since the different pages associated with the application rely on common processing tasks, such as database accessing, they will need duplicate code.
- The different navigation controls may result in a combination of content and navigation.
- Each view must have a separate means of activating system services and locating system resources.

Solution

We need a single object to manage view control, resource/service accessing, error handling, and initial request processing. This object acts as a front door to the client and is called a Front Controller, or Front Component. Among the various strategies suggested by J2EE, the two simplest ones are using a servlet or a JSP as front objects. All requests will be sent to the Front Controller and each request will have an action as a parameter.

Consequences/implications

- The control of use cases is centralized. A change in the sequence of steps affects only the Front Controller component.
- Many web applications save the state of the client-server interaction if the user logs out in the middle of a process. When the user logs in at some other time, the previously saved state is reloaded and the process resumes from the point where it was left. In such cases, it is easier to maintain the state information using the Front Controller component because only one component handles the state management.
- Multiple Front Controller objects can be developed; each controller can concentrate on a different business use case.
- The reusability of worker components increases. Since the code that manages navigation across web pages now resides only in the Front Controller, it need not be repeated in worker components. Thus, multiple Front Controllers can reuse worker components.

Category

Since the Front Controller pattern is concerned with communication with other components, it falls in the category of behavioral pattern.

In the J2EE pattern catalog, the Front Controller pattern is kept under the presentation tier because it directly deals with the client's requests and dispatches them to the appropriate handler or worker components.

Points to remember

A Front Controller, or Front Component, is a component that provides a common point of entry for all client requests. In this way, the controller unifies and streamlines authentication and authorization, and dispatches the work to appropriate worker components, thereby facilitating use case management.

Pay special attention to the following words, phrases, or terms appearing in the problem statement:

- Dispatch requests
- Manage workflow of a web application
- Manage the sequence of steps
- Manage use cases

Many times, developers confuse the two design patterns—Intercepting Filter and Front Controller. Remember that while an Intercepting Filter performs pre-processing of requests, the Front Controller actually begins the main process. Further, the Front Controller is much more involved in controlling the application's appearance and activating system services than the Intercepting Filter. As an analogy, if the Intercepting Filter is the security guard outside a building, the Front Controller is the receptionist inside who can direct you to where you need to go.

18.2.5 Service Locator

Context

Many systems depend on distributed processing and communication. In these cases, a number of concerns must be kept in mind:

- The system needs to locate and access resources and services across networks.
- If the services or resources are unavailable, the system must create local implementations.
- New requests to a directory service require more time and resources than cached repetitions of previous requests.

Problem

If distributed objects access the same resources across a network, their operation will cause added traffic and reduced communication efficiency. Also, each object will need updating when the network changes, a process demanding a great deal of code and time.

Example

A large company contains a number of worldwide branches and needs to maintain communication throughout. Many different applications require access to similar external resources, such as personnel databases and inventory listings. In this case, the system uses Java Naming and Directory Interface (JNDI) to help components find one another. But as each object accesses JNDI separately, its responsiveness slows.

Facts or forces to consider

In the context and the problem presented above, we observe the following facts:

- Using distributed objects with external resource connections means greater code and communication traffic.
- Caching external requests becomes more efficient as more requests are made.
- Each distributed object needs to be reformed whenever a resource changes its position in the network.

Solution

The solution is to encapsulate the process of external communication into a Service Locator. With this method, the Locator object manages all of the complexity related to locating and interfacing resources across the network. Further, to reduce the amount of time associated with making requests, they can provide request caching. Finally, if external resources have vendor-specific interfaces, this object will manage the communication complexity and make accessing the resource transparent to the user.

Consequences/implications

- When a resource across the network is altered in any way, only the Service Locator object needs to be updated.
- The Service Locator caching makes it very efficient for multiple objects to access similar resources across a network.
- The Service Locator provides a single, uniform interface to external resources, no matter what the resource's actual API may be. This reduces the amount of code needed for distributed communication, and makes it easier to test network capability.

Category

Because Service Locator objects assist with network communication by reducing the coupling between business logic and external resources, this pattern is classified into the Behavioral category. By adding caching and removing sources of traffic, this object can improve the performance of a distributed application and reduce the amount of code needed to build and maintain it.

In the J2EE pattern catalog, the Service Locator pattern is kept under the business tier because it doesn't directly deal with clients, but instead manages the connectivity of the underlying business logic.

Points to remember

The Service Locator pattern is very simple to understand and identify on the exam, but you should pay special attention to the following words, phrases, or terms appearing in the problem statement:

- Provides access to heterogeneous networks and services (JNDI, RMI, etc)
- A single point-of-contact for managing distributed connections and resources
- Distributed request caching
- Improves ease of business application development by encapsulating interface complexity

18.2.6 Business Delegate

Context

In an enterprise-scaled distributed application, typically the following situation arises:

- There are separate components to interact with the end users and to handle business logic.
- These separate components reside in different subsystems, separated by a network.
- The components that handle the business logic act as server components because they provide business services to the clients by exposing the business service API to the clients.
- The client components that use this API often reside in a remote system separated by a network.
- There is more than one client using the API.
- There is more than one server component providing similar services, but with minor differences in the API.

Problem

Business services implemented by the business-tier components are accessed directly by the presentation-tier components through the exposed API of the services. However, the interfaces of such services keep changing as the requirements evolve. This affects all the components on the presentation tier. Furthermore, all the client-side components have to be aware of the location details of the business services—that is, each component has to use the JNDI lookup service to locate the required remote interfaces.

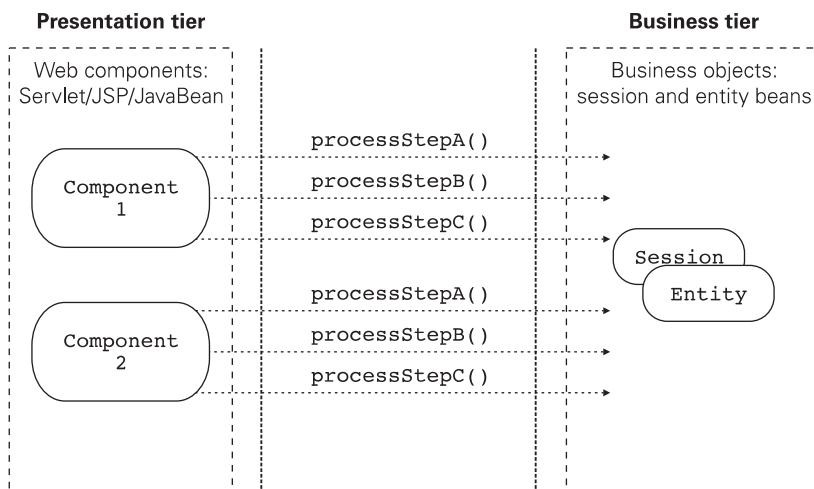


Figure 18.2 Relationship of J2EE components in multiple tiers

Example

In the case of the J2EE architecture, the server components that expose the business service API are the session beans, the API is the remote interface implemented by the session beans, and the client components that use these services are servlets and the JavaBeans used in JSP pages. Figure 18.2 shows this relationship.

Let's look at a real-world example. A company is building a web-based application with JSP pages and servlets that need to access a set of business services. The management has decided not to develop the business services in-house, since they are readily available as off-the-shelf software from various vendors. In addition, the budget for the project is currently tight, so management has decided that they will purchase one of the more economical off-the-shelf solutions initially, and then when the money becomes available in a year, they will replace it with a more elaborate and expensive software solution.

Fact or forces to consider

- The presentation-tier components, which in this case are the web components—servlets, JSP pages, and JavaBeans—perform two main tasks:
- Handling the end user, which involves managing the web application logic, presenting the data, and so forth
- Accessing the business services
- The code that handles the end user should not be dependent on the code that accesses the business services.
- Multiple presentation-tier components can call the same set of remote methods in the same sequence.
- It is expected that the business service APIs will change as business requirements evolve.

Solution

Create a Business Delegate to handle all of the code that accesses the business services in the selected vendor software. When the vendor changes, the only changes that need to be made to the company's application software are changes to the Business Delegate, to access the business services in the new vendor software. The JSP pages and servlets will not have to be modified.

As shown in figure 18.3, we need to separate the code that accesses the remote service from the code that handles the presentation. We can put this service access code in a separate object. This separate object is called a *Business Delegate* object.

The Business Delegate object abstracts the business services API and provides a standard interface to all the client components. It hides the underlying implementation details, such as the lookup mechanism and the API of the business services. This reduces the coupling between the clients and the business services.

The responsibilities of the components participating in this pattern are

- *Client components*—The client components, which are JSP pages and servlets in the presentation tier, delegate the work of locating the business service providers and the work of invoking the business service API methods to the Business Delegate objects.
- *Business Delegate*—The Business Delegate acts as a representative of the client components. It knows how to look up and access the business services. It invokes the appropriate business services methods in the required order. If the API of the business service component changes, only the Business Delegate needs to be modified, without affecting the client components.
- *Business service*—A business service component implements the actual business logic. Some examples of business service components are a stateless session EJB, an entity EJB, a CORBA object, or an RPC server.

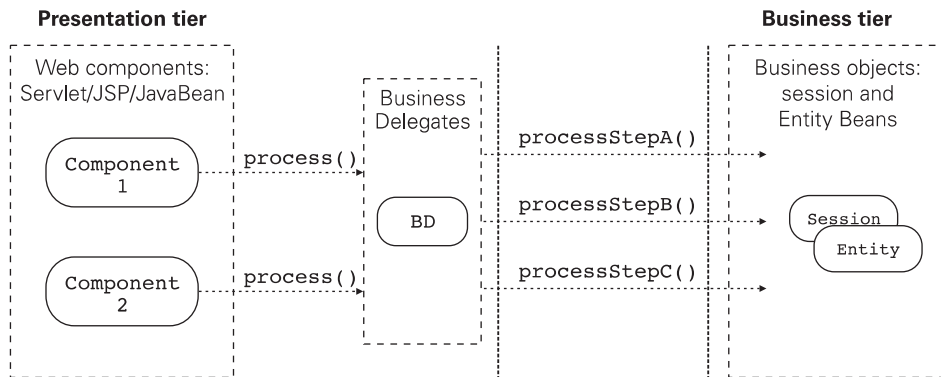


Figure 18.3 The Business Delegate pattern

Consequences/implications

- Repetition of code is avoided. Each component does not have to include the code that will perform the lookup operation on the remote interfaces and invoke the methods.
- The server business logic API is hidden from the client components. Thus, it reduces the number of changes required in the client components when there is a change in the API of the server components.
- The Business Delegate can do all the business service-specific tasks, such as catching exceptions raised by the business services. For instance, it can catch remote exceptions and wrap them into application exceptions that are more user friendly.
- The results of remote invocations may be cached. This significantly improves performance, because it eliminates repetitive and potentially costly remote calls. The cached results may be used by multiple client components, again reducing code repetition and increasing performance.

It is worth mentioning here that the Business Delegate can either locate the business services itself or use another pattern, called the Service Locator pattern, to help locate the business service. In the instances when it uses the Service Locator pattern, the Business Delegate will deal only with the business service API invocation regardless of where the services are located, while the Service Locator does the job of locating the required named services. It is important to know that multiple Business Delegate objects can share a common Service Locator.

Category

Since the Business Delegate pattern is concerned with the communication between two components—the presentation-tier and business-tier components—it falls into the category of a behavioral pattern. It describes how we can reduce the coupling between the two communicating parties by introducing the delegation layer in between, and how we can increase the flexibility of the design.

In the J2EE pattern catalog, the Business Delegate pattern is kept under the business tier category since it is more closely related to the business-tier components. However, note that the objects that implement the Business Delegate pattern reside in the presentation tier.

Points to remember

A Business Delegate is an object that resides on the client side and communicates with the business service components residing on the server side. The client-side components can *delegate* the work of accessing the *business* services exposed by the business service components to the Business Delegate.

Pay special attention to the following words, phrases, or terms appearing in the problem statement:

- Reduce coupling between presentation and business tiers
- Proxy for the client
- Client-side facade
- Cache business service references for presentation-tier components
- Cache business service results for presentation-tier components
- Encapsulate business service lookup
- Encapsulate business service access
- Decouple clients from business service API

18.2.7 Transfer Object

Context

In distributed applications, the following situation typically arises:

- The client-side and the server-side components reside at remote locations and communicate over the network.
- The server handles the database.
- The server provides getter methods to the clients so that the clients can call those getter methods one by one to retrieve database values.
- The server provides setter methods to the clients so that the clients can call those setter methods one by one to update database values.

Problem

Every call between the client and the server is a remote method call with substantial network overhead. If the client application calls the individual getter and setter methods that retrieve or update single attribute values, it will require as many remote calls as there are attributes. These individual calls generate a lot of network traffic and degrade the system performance.

Example

In the J2EE architecture, the business tier accesses the database directly or via the resource tier, and wraps the data access mechanism inside a set of entity beans and session beans. These entity and session beans expose the data via remote interfaces. The servlets and JSP pages in the presentation tier that need to access business data can do so by calling methods on the remote interfaces implemented by the beans.

As a specific example, suppose we maintain the address information in the database of the registered users of our enterprise application. In this case, the address information, which is a summation of four other pieces of data—street, city, state, and zip—is a business entity. The access to this information is encapsulated by the application's business tier with the help of a session bean called `AddressSessionBean`.

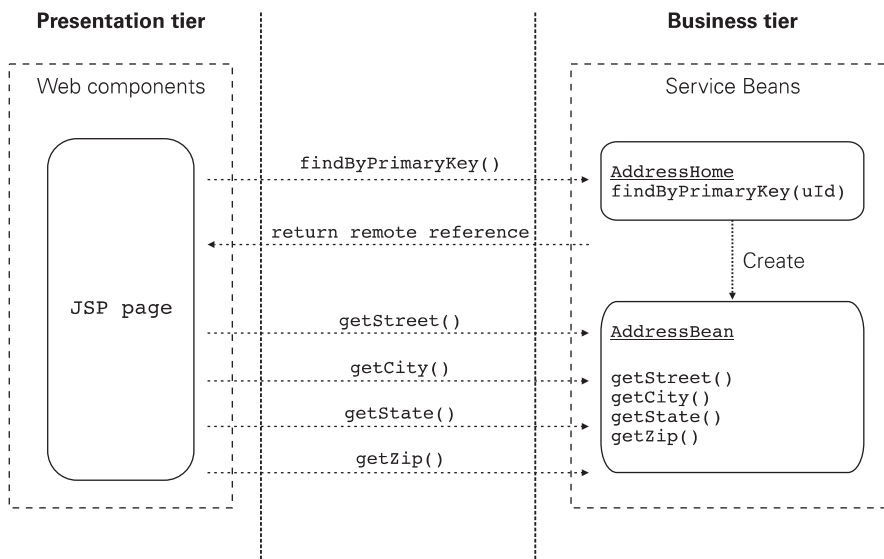


Figure 18.4 Accessing attributes remotely

`AddressSessionBean` exposes methods for the remote clients, such as `getState()`, `setState()`, `getCity()`, and `setCity()`. The servlets and the JSP pages then have to call each of the methods one by one on the remote server, as shown in figure 18.4.

Facts or forces to consider

In the context and the problem presented above, we observe the following facts:

- A single business object has many attributes.
- Most of the time, the client requires values for more than one attribute simultaneously rather than just an individual attribute.
- The rate of retrieving the data (calling the getter methods) is higher than the rate of updating the data (calling the setter methods).

For example, the address consists of the street, state, city, and zip. Each time the user buys a product online, we want to show the full address information for billing purposes. On the other hand, it is unlikely that the user's address changes very often.

Solution

Create an object to encapsulate all of the attribute values that are required by the client application. This object is called the Transfer Object. When the client requests the data from the server, the server-side component gathers the data values and constructs the Transfer Object by setting its data values. This object is then sent to the client by

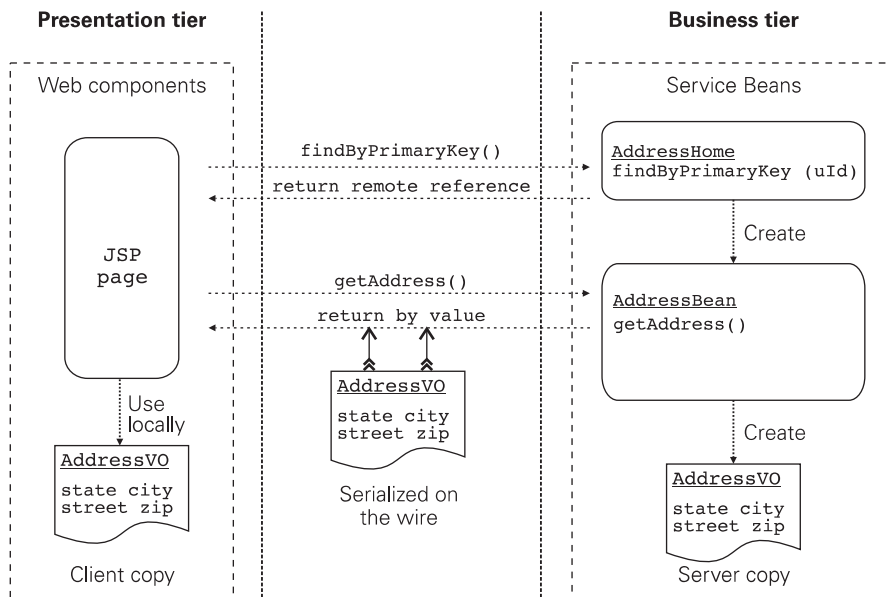


Figure 18.5 Accessing attributes using the Value Object design pattern

value (not by reference), which means that the whole object is serialized and each of its bits is transferred over the network.

The client on the other side reconstructs this object locally with all the values intact. It can then query this local instance for all the attribute values. Because the Transfer Object is local on the client, all of the calls to this object are local and do not incur any network overhead. The Transfer Object on the client serves as a proxy for the properties of the remote object. This scenario is shown in figure 18.5.

Now, instead of making multiple remote calls on the business object, `AddressBean`, to retrieve all the attributes, the client calls a single method, `getAddress ()`, which returns all the attributes structured in an `AddressVO` object.

Consequences/implications

- The remote interfaces are simpler because the multiple methods returning single values are collapsed into one single method returning a group of multiple values.
- Because of the reduced number of calls across the network, the user response time improves.
- If the client wants to update the attribute values, it first updates the values in the local Transfer Object and then sends the updated Transfer Object to the server to persist the new values. This also happens using the pass-by-value mechanism.

- The Transfer Object can become stale—that is, if the client has acquired a Transfer Object for a long time, there is a possibility that the information may have been updated by another client.
- In the case of a mutable Transfer Object, requests for update from two or more clients can result in data conflict.

Category

Since the Transfer Object pattern is concerned with communication between two other components, it falls in the category of a behavioral pattern.

In the J2EE pattern catalog, the Transfer Object is kept under the business tier because it represents the business object on the client side. However, note that even though the object that implements the Transfer Object pattern is created in the business tier, it is transferred to the presentation tier and is actually used in the presentation tier.

Points to remember

A Transfer Object is a small-sized serializable Java *object* that is used for carrying grouped data (*values*) over the network from one component residing in one tier to another component residing in another tier of a multitier distributed application. Its purpose is to reduce communication overhead by reducing the number of remote calls between the distributed components.

Pay special attention to the following words, phrases, or terms appearing in the problem statement:

- Small object
- Grouped information
- Read-only data
- Reduce network traffic
- Improve response time
- Transfer data across networked tiers

18.3 SUMMARY

Design patterns induce abstraction, division of labor, and reusability in software systems. Consistent use of design patterns results in scalable and maintainable systems. We briefly looked at design patterns of the J2EE architecture, classified into three tiers: presentation, business, and integration. Then we examined in depth the six design patterns that are specified by the exam objectives: Intercepting Filter, Model-View-Controller, Front Controller, Service Locator, Business Delegate, and Transfer Object.

If you are interested in learning more about design patterns, various books and articles on patterns related to all kinds of domains are available. There are books on design

patterns in compiler writing, design patterns on parallel computing, design patterns specifically for the Java programming language, and so forth. During your explorations, you may soon find yourself lost in a jungle of patterns, as you discover that some of them do the same thing but have different names, and that many of them provide similar solutions but different implementations.

If you are planning to study for the Sun Certified Enterprise Architect (SCEA) exam in addition to becoming a Sun Certified Web Component Developer (SCWCD), then we suggest you become familiar with all the J2EE patterns and the various specifications that are part of the J2EE family of architectures.

As for the SCWCD, with the end of this chapter, you are now ready to answer the questions based on the important design patterns: Intercepting Filter, Model-View-Controller, Front Controller, Service Locator, Business Delegate, and Transfer Object.

18.4 REVIEW QUESTIONS

1. What are the benefits of using the Transfer Object pattern? (Select two)
 - a The type of the actual data source can be specified at deployment time.
 - b The data clients are independent of the data source vendor API.
 - c It increases the performance of data-accessing routines.
 - d It allows the clients to access the data source through EJBs.
 - e It allows resource locking in an efficient way.
2. Which design pattern allows you to decouple the business logic, data representation, and data presentation? (Select one)
 - a Model-View-Controller
 - b Transfer Object
 - c Bimodal Data Access
 - d Business Delegate
3. Which of the following are the benefits of using the Transfer Object design pattern? (Select two)
 - a It improves the response time for data access.
 - b It improves the efficiency of object operations.
 - c It reduces the network traffic.
 - d It reduces the coupling between the data access module and the database.
4. Which of the following statements are correct? (Select two)
 - a The Transfer Object pattern ensures that the data is not stale at the time of use.
 - b It is wise to make the Transfer Object immutable if the Transfer Object represents read-only data.
 - c Applying the Transfer Object pattern on EJBs helps to reduce the load on enterprise beans.
 - d A Transfer Object exists only on the server side.

5. What are the benefits of using the Business Delegate pattern? (Select three)
- a It implements the business service functionality locally to improve performance.
 - b It shields the clients from the details of the access mechanism, such as CORBA or RMI, of the business services.
 - c It shields the clients from changes in the implementation of the business services.
 - d It provides the clients with a uniform interface to the business services.
 - e It reduces the number of remote calls and reduces network overhead.
6. You are designing an application that is required to display the data to users through HTML interfaces. It also has to feed the same data to other systems through XML as well as WAP interfaces. Which design pattern would be appropriate in this situation? (Select one)
- a Interface Factory
 - b Session Facade
 - c Transfer Object
 - d Model-View-Controller
 - e Factory
7. You are automating a computer parts ordering business. For this purpose, your web application requires a controller component that would receive the requests and dispatch them to appropriate JSP pages. It would also coordinate the request processing among the JSP pages, thereby managing the workflow. Finally, the behavior of the controller component is to be loaded at runtime as needed. Which design pattern would be appropriate in this situation? (Select one)
- a Front Controller
 - b Session Facade
 - c Transfer Object
 - d Model-View-Controller
 - e Data Access Object
8. You are building the server side of an application and you are finalizing the interfaces that you will provide to the presentation layer. However, you have not yet finalized the access details of the business services. Which design pattern should you use to mitigate this concern? (Select one)
- a Model-View-Controller
 - b Data Access Object
 - c Business Delegate
 - d Facade
 - e Transfer Object



A P P E N D I X A

Installing Tomcat 5.0.25

You'll need to install Tomcat to test the sample code that we have developed in the chapters. This appendix will help you install Tomcat.

A.1 PREREQUISITES

You should have JDK 1.4 or higher installed and working on your machine.

A.2 GETTING TOMCAT

You can download the latest version of Tomcat from <http://jakarta.apache.org/tomcat>. The version currently available is Tomcat 5.0.25.

A.3 INSTALLATION

Installing Tomcat is a straightforward process. The following sections will help you install and set up Tomcat on Windows 98/NT/2000.

A.3.1 Extracting files

Double-click on `jakarta-tomcat-5.0.25.exe` to start installing the files. An installation dialog box will appear, asking for configuration information. During the installation procedure, we recommend that you make two changes to the defaults:

- Change the destination folder to `C:\jakarta-tomcat-5.0.25`. This directory will be referred to as `CATALINA_HOME`.
- Change the HTTP Connector/1.1 Port from 8080 to 80 (unless you're already running a web server).

The first change makes it simpler to find and remember the location of `CATALINA_HOME`. The second makes it possible to invoke servlets and JSPs without setting the port number.

A.3.2 Setting environment variables

To run Tomcat, you need to set two environment variables: `CATALINA_HOME` and `JAVA_HOME`. `CATALINA_HOME` refers to the Tomcat installation directory, and `JAVA_HOME` refers to the JDK 1.4 installation directory.

To be able to compile your servlets, you need to add the proper API classes to your `CLASSPATH` environment variable. These classes are contained within two files: `c:\jakarta-tomcat-5.0.25\common\lib\servlet-api.jar` and `c:\jakarta-tomcat-5.0.25\common\lib\jsp-api.jar`.

On Windows 98, open `c:\autoexec.bat` and add the following:

```
SET CATALINA_HOME=c:\jakarta-tomcat-5.0.25
SET JAVA_HOME=c:\jdk1.4.2
SET CLASSPATH=%CLASSPATH%;.;c:\jakarta-tomcat-5.0.25\common\lib\servlet-
api.jar;c:\jakarta-tomcat-5.0.25\common\lib\jsp-api.jar
```

Be sure to substitute the appropriate directory for `c:\jdk1.4.2`. You'll need to restart your machine.

In Windows 2000/NT, go to Start|Settings|Control Panel|System. Then select the Advanced tab and in the Environment Variables section, set the above two variables as either System variables or User variables. You do not need to restart your machine, but you will need to close all the DOS windows and open a new DOS window to see the newly set variables.

A.4 DIRECTORY STRUCTURE

After you've installed Tomcat, your `jakarta-tomcat-5.0.25` directory should look like the one shown in figure A.1.

There are three important directories under the `jakarta-tomcat-5.0.25` directory. Let's take a look at each.

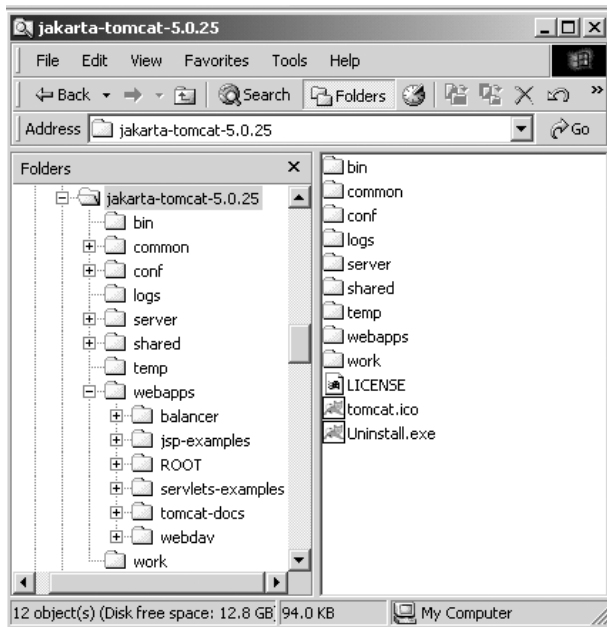


Figure A.1
The Tomcat
directory
structure

The bin directory

This directory contains the executable batch files, such as `startup.bat` and `shutdown.bat`.

The conf directory

This directory contains several configuration files, such as `server.xml` and `web.xml`.

The webapps directory

Tomcat keeps all the web applications in this directory. Each directory that you see under the `c:\jakarta-tomcat-5.0.25\webapps` directory corresponds to a web application. The directory named `ROOT` refers to the default web application.

A.5 RUNNING TOMCAT

Although you can control Tomcat with DOS commands, it's much easier to use desktop shortcuts. If you look in the `c:\jakarta-tomcat-5.0.25\bin` directory, you'll see a number of batch (with the extension `.bat`) files. The two we're interested in are `startup.bat` and `shutdown.bat`, which function as their names imply.

We recommend that you create a shortcut for each by right-clicking, choosing the `Create Shortcut` option, and moving the two shortcuts to your desktop. From there, you can start Tomcat by double-clicking on `startup.bat`, and end its execution by double-clicking on `shutdown.bat`.

To check whether Tomcat is running, go to `http://localhost` from your browser. You'll see the default page saying your installation is successful. Here, the name `localhost` means that you are connecting to your own machine (since Tomcat is running on the same machine). You can also use `127.0.0.1` instead of `localhost`—that is, `http://127.0.0.1`.

In some cases, Windows 98 may complain about low memory. If this happens, right-click on a DOS window's top border and choose Properties|Memory. Set the default memory from Auto to 8192. Alternatively, you can add the following line in `autoexec.bat`:

```
SHELL=C:/windows/command.com /E:8192 /P
```

If you work in an office environment where you use a proxy server to browse the Web, you should set the browser so that it bypasses the proxy for the local addresses. To do this for Microsoft Internet Explorer, choose Tools|Internet Options|Connections|LAN Settings. Then, select the Bypass Proxy For Local Addresses checkbox.

A.6 CREATING A NEW WEB APPLICATION

Each directory in the `webapps` directory of the Tomcat installation represents a web application. So for instance, if you want to create a new web application with the name `helloapp`, you create a directory by that name in the `webapps` directory. You can put static files directly into this directory and then view them from the browser. Be sure to add a `WEB-INF` directory also.

For example, if you put in a file named `myhomepage.html`, you can view it through the following URL:

```
http://localhost/helloapp/myhomepage.html
```

However, before you add servlets and JSP pages, you should read the discussion about the directory structure of a web application in chapter 5, "Structure and deployment."

A.7 SECURITY

By default, Tomcat runs in a nonsecure mode. This means that a web application class can access anything on the system. For example, a malicious servlet class or a JSP page can delete files from your system.

You can prevent this by accessing Tomcat from the DOS prompt and starting it in secure mode:

```
C:\jakarta-tomcat-5.0.25\bin>startup -security
```

The `-security` option forces Tomcat to use a security manager. This security manager uses the policies specified in the `conf/catalina.policy` file. You can customize this policy file by granting rights to specific web applications as per the application requirements. If an application tries to access anything for which it does

not have the appropriate rights, it will be denied access and the server will throw a `java.security.AccessControlException`.

Security is especially important when you want to run a readymade third-party web application under your Tomcat installation; running Tomcat in the secure mode will protect your system. Also, service providers often use a shared Tomcat instance for hosting multiple web applications, in which case running it in the secure mode becomes a necessity.



A P P E N D I X B

A sample web.xml file

The following listing of a sample `web.xml` file illustrates the use of various elements of a deployment descriptor. We have explained the elements that you are required to know for the exam throughout the book. This code listing is just for a quick recap.

```
<?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 web-app_2_4.xsd"
         version="2.4">
  <display-name>Test Webapp</display-name>
  <description>This is a sample deployment descriptor that shows
    the use of important elements.
  </description>
  <!-- Presence of this element indicates that this
    WebApp is distributable. -->
  <distributable/>
  <!-- Defines WebApp initialization parameters.-->
  <context-param>
    <param-name>locale</param-name>
    <param-value>US</param-value>
  </context-param>
```

```

<context-param>
  <param-name>DBName</param-name>
  <param-value>Oracle</param-value>
</context-param>

<!-- Defines filters and specifies filter mapping -->
<filter>
  <filter-name>Test Filter</filter-name>
  <description>Just for test</description>
  <filter-class>filters.TestFilter</filter-class>
  <init-param>
    <param-name>locale</param-name>
    <param-value>US</param-value>
  </init-param>
</filter>

<filter-mapping>
  <filter-name>Test Filter</filter-name>
  <servlet-name>TestServlet</servlet-name>
</filter-mapping>

<!-- Defines application events listeners -->
<listener>
  <listener-class>listeners.MyServletContextListener
</listener-class>
</listener>
<listener>
  <listener-class>listeners.MySessionCumContextListener
</listener-class>
</listener>

<!-- Defines servlets -->
<servlet>
  <servlet-name>TestServlet</servlet-name>
  <description>Just for test</description>
  <servlet-class>servlets.TestServlet</servlet-class>
</servlet>
<servlet>
  <servlet-name>HelloServlet</servlet-name>
  <servlet-class>servlets.HelloServlet</servlet-class>
  <init-param>
    <param-name>locale</param-name>
    <param-value>US</param-value>
  </init-param>
  <load-on-startup>1</load-on-startup>
  <security-role-ref>
    <!-- role-name is used in
      HttpServletRequest.isUserInRole(String role)
      method. -->
    <role-name>manager</role-name>
    <!-- role-link is one of the role-names specified in
      security-role elements. -->
    <role-link>supervisor</role-link>
  </security-role-ref>
</servlet>

```

```

        <security-role-ref>
    </servlet>

<!-- Defines servlet mappings -->
<servlet-mapping>
    <servlet-name>TestServlet</servlet-name>
    <url-pattern>/test/*</url-pattern>
</servlet-mapping>
<servlet-mapping>
    <servlet-name>HelloServlet</servlet-name>
    <url-pattern>*.hello</url-pattern>
</servlet-mapping>

<session-config>
    <!-- specifies session timeout as 30 minutes. -->
    <session-timeout>30</session-timeout>
</session-config>

<mime-mapping>
    <extension>jar</extension>
    <mime-type>application/java-archive</mime-type>
</mime-mapping>
<mime-mapping>
    <extension>conf</extension>
    <mime-type>text/plain</mime-type>
</mime-mapping>

<welcome-file-list>
    <welcome-file>index.html</welcome-file>
    <welcome-file>home.html</welcome-file>
    <welcome-file>welcome.html</welcome-file>
</welcome-file-list>

<error-page>
    <error-code>404</error-code>
    <location>notfoundpage.jsp</location>
</error-page>
<error-page>
    <exception-type>java.sql.SQLException</exception-type>
    <location>sqlexception.jsp</location>
</error-page>

<taglib>
    <taglib-uri>http://abc.com/testlib</taglib-uri>
    <taglib-location>
        /WEB-INF/tlds/testlib.tld
    </taglib-location>
</taglib>
<taglib>
    <taglib-uri>/examplelib</taglib-uri>
    <taglib-location>
        /WEB-INF/tlds/examplelib.tld
    </taglib-location>
</taglib>

```

```

<security-constraint>
  <display-name>Example Security Constraint</display-name>
  <web-resource-collection>
    <web-resource-name>Protected Area</web-resource-name>
    <url-pattern>/test/*</url-pattern>
    <!-- only POST method is protected -->
    <http-method>POST</http-method>
  </web-resource-collection>

  <web-resource-collection>
    <web-resource-name>Another Protected Area</web-resource-name>
    <url-pattern>*.hello</url-pattern>
    <!-- All methods are protected as no http-method is
    specified -->
  </web-resource-collection>

  <auth-constraint>
    <!-- Only the following roles can access the above resources.
    The role must be defined in security-role. -->
    <role-name>supervisor</role-name>
  </auth-constraint>

  <user-data-constraint>
    <!-- Specifies the type of communication
    between the client and the server.
    It can be: NONE, INTEGRAL, or CONFIDENTIAL -->
    <transport-guarantee>INTEGRAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>

<login-config>
  <!-- auth-method can be: BASIC, FORM, DIGEST, or CLIENT-CERT -->
  <auth-method>FORM</auth-method>
  <realm-name>sales</realm-name>
  <form-login-config>
    <form-login-page>/formlogin.html</form-login-page>
    <form-error-page>/formerror.jsp</form-error-page>
  </form-login-config>
</login-config>

<!-- Specifies the roles that are defined in the application
server. For example, Tomcat defines it in
conf\tomcat-users.xml -->
<security-role>
  <role-name>supervisor</role-name>
</security-role>
<security-role>
  <role-name>worker</role-name>
</security-role>
</web-app>

```



A P P E N D I X C

Review Q & A

CHAPTER 4—THE SERVLET MODEL

1. Which method in the `HttpServlet` class services the HTTP POST request? (Select one)
 - a `doPost (ServletRequest, ServletResponse)`
 - b `doPOST (ServletRequest, ServletResponse)`
 - c `servicePost (HttpServletRequest, HttpServletResponse)`
 - d `doPost (HttpServletRequest, HttpServletResponse)`

Answer: d

Explanation

Remember that `HttpServlet` extends `GenericServlet` and provides HTTP-specific functionality. Thus, all its methods take `HttpServletRequest` and `HttpServletResponse` objects as parameters.

Also, the method names follow the standard Java naming convention—for example, the method for processing POST requests is `doPost ()` and not `doPOST ()`.

2. Consider the following HTML page code:

```
<html><body>
<a href="/servlet/HelloServlet">POST</a>
</body></html>
```

Which method of `HelloServlet` will be invoked when the hyperlink displayed by the above page is clicked? (Select one)

- a `doGet`
- b `doPost`
- c `doForm`

- d doHref
- e serviceGet

Answer: a

Explanation

Don't get confused by the text POST displayed by the hyperlink. A click on a hyperlink always generates an HTTP GET request, which is handled by the `doGet()` method. You can generate a POST request through a hyperlink by using JavaScript. For example:

```
<html>
<script lanaguage="JavaScript">
function sendPost()
{
    dummyform.submit();
}
</script>
<body>
<form name="dummyform" action=
        "/servlet/HelloServlet" method="POST">
<input type="text" name="name">
</form>
<a href="" onClick="javascript:sendPost();">POST</a>
</body>
</html>
```

This HTML code executes the JavaScript function `sendPost()` whenever the hyperlink is clicked. This function submits the `dummyform`, causing a POST request to be sent.

3. Consider the following code for the `doGet()` method:

```
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
{
    PrintWriter out = res.getWriter();
    out.println("<html><body>Hello</body></html>");

    //1

    if(req.getParameter("name") == null)
    {
        res.sendError(HttpServletResponse.SC_UNAUTHORIZED);
    }
}
```

Which of the following lines can be inserted at `//1` so that the above code does not throw any exception? (Select one)

- a `if (! res.isSent())`
- b `if (! res.isCommitted())`
- c `if (! res.isDone())`


```
d if ( ! res.isFlushed() )
e if ( ! res.flush() )
```

Answer: b

Explanation

This question is based on the concept that the `HttpServletResponse.send-Error()` method throws an `IllegalStateException` if the response has already been sent to the client. The `ServletRequest.isCommitted()` method checks whether or not the response is committed.

4. Which of the following lines would initialize the `out` variable for sending a Microsoft Word file to the browser? (Select one)

```
a PrintWriter out = response.getServletOutput();
b PrintWriter out = response.getWriter();
c OutputStream out = response.getWriter();
d PrintWriter out = response.getOutputStream();
e OutputStream out = response.getOutputStream();
f ServletOutputStream out = response.getServletOutputStream();
```

Answer: e

Explanation

For sending any data other than text, you need to get the `OutputStream` object. `ServletResponse.getOutputStream()` returns an object of type `ServletOutputStream`, where `ServletOutputStream` extends `OutputStream`.

5. You need to send a GIF file to the browser. Which of the following lines should be called after (or before) a call to `response.getOutputStream()`? (Select one)

```
a response.setContentType("image/gif"); Before
b response.setContentType("image/gif"); After
c response.setContentType("image/gif"); Before
d response.setContentType("image/gif"); After
e response.setContentType("image/gif"); Before
f response.setContentType("image/gif"); After
```

Answer: a

Explanation

You need to set the content type of the response using the `ServletResponse.setContentType()` method *before* calling the `ServletResponse.getOutputStream()` method.

6. Consider the following HTML page code:

```
<html><body>
<form name="data" action="/servlet/DataServlet" method="POST">
<input type="text" name="name">
<input type="submit" name="submit">
</form>
</body></html>
```

Identify the two methods that can be used to retrieve the value of the name parameter when the form is submitted.

- a `getParameter("name");`
- b `getParameterValue("name");`
- c `getParameterValues("name");`
- d `getParameters("name");`
- e `getValue("name");`
- f `getName();`

Answers: a and c

Explanation

`ServletRequest` provides two methods to retrieve input parameters:

- `getParameter("name")`: Returns a `String` or `null`.
- `getParameterValues("name")`: Returns a `String` array containing all the values for the name parameter or `null`.

Besides these two, `ServletRequest` also provides a `getParameterNames()` method that returns an `Enumeration` object of all the parameter names present in the request, or an empty `Enumeration` if the request does not contain any parameter.

7. Which of the following methods would you use to retrieve header values from a request? (Select two)

- a `getHeader()` of `ServletRequest`
- b `getHeaderValue()` of `ServletRequest`
- c `getHeader()` of `HttpServletRequest`
- d `getHeaders()` of `ServletRequest`
- e `getHeaders()` of `HttpServletRequest`

Answers: b and e

Explanation

Headers are a feature of the HTTP protocol. Thus, all the header-specific methods belong to `HttpServletRequest`. `getHeader()` returns a `String` (or `null`), while `getHeaders()` returns an `Enumeration` of all the values for that header (or an empty `Enumeration`).

8. Consider the following code:

```
public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
    throws IOException
{
    if(req.getParameter("switch") == null)
    {
        //1
    }
    else
    {
```

```

        //other code
    }
}

```

Which of the following lines can be inserted at //1 so that the request is redirected to the `collectinfo.html` page? (Select one)

- a `req.sendRedirect("collectinfo.html");`
- b `req.redirect("collectinfo.html");`
- c `res.direct("collectinfo.html");`
- d `res.sendRedirect("collectinfo.html");`
- e `this.sendRedirect("collectinfo.html");`
- f `this.send("collectinfo.html");`

Answer: d

Explanation

You can redirect the client to another resource using the `HttpServletResponse.sendRedirect()` method.

9. Consider the following code:

```

public void doGet(HttpServletRequest req,
                  HttpServletResponse res)
{
    HttpSession session = req.getSession();
    ServletContext ctx = this.getServletContext();

    if(req.getParameter("userid") != null)
    {
        String userid = req.getParameter("userid");
        //1
    }
}

```

You want the `userid` parameter to be available only to the requests that come from the same user. Which of the following lines would you insert at //1? (Select one)

- a `session.setAttribute("userid", userid);`
- b `req.setAttribute("userid", userid);`
- c `ctx.addAttribute("userid", userid);`
- d `session.addAttribute("userid", userid);`
- e `this.addParameter("userid", userid);`
- f `this.setAttribute("userid", userid);`

Answer: a

Explanation

Attributes stored in the session scope are available only for the requests from the same client. Attributes stored in the context scope are available for all the requests to the same web application from all the clients. Attributes stored in the request scope are available only for the request in which it is stored.

10. Which of the following lines would you use to include the output of `DataServlet` into any other servlet? (Select one)

a `RequestDispatcher rd = request.getRequestDispatcher("/servlet/DataServlet"); rd.include(request, response);`
b `RequestDispatcher rd = request.getRequestDispatcher("/servlet/DataServlet"); rd.include(response);`
c `RequestDispatcher rd = request.getRequestDispatcher(); rd.include("/servlet/DataServlet", request, response);`
d `RequestDispatcher rd = request.getRequestDispatcher(); rd.include("/servlet/DataServlet", response);`
e `RequestDispatcher rd = request.getRequestDispatcher(); rd.include("/servlet/DataServlet");`

Answer: a

Explanation

To forward or include a request to another resource, first you need to get a `RequestDispatcher` object from either `ServletRequest` or `ServletContext`. Then you can call `include()` or `forward()` and pass the current request and response objects as parameters.

CHAPTER 5—STRUCTURE AND DEPLOYMENT

1. Which element is used to specify useful information about an initialization parameter of a servlet in the deployment descriptor? (Select one)

a `param-description`
b `description`
c `info`
d `param-info`
e `init-param-info`

Answer: b

Explanation

Remember that the `description` element is used for all the elements that can take a description (useful information about that element). This includes `servlet`, `init-param`, and `context-param`, among others. For a complete list of the elements that can take a description, please read the DTD for `web.xml`.

2. Which of the following deployment descriptor snippets correctly associates a servlet implemented by a class named `com.abc.SalesServlet` with the name `SalesServlet`? (Select one)

a `<servlet>
 <servlet-name>com.abc.SalesServlet</servlet-name>
 <servlet-class>SalesServlet</servlet-class>
</servlet>`
b `<servlet>
 <servlet-name>SalesServlet</servlet-name>`

```

    <servlet-package>com.abc.SalesServlet</servlet-package>
  </servlet>
c <servlet>
  <servlet-name>SalesServlet</servlet-name>
  <servlet-class>com.abc.SalesServlet</servlet-class>
</servlet>
d <servlet name="SalesServlet" class="com.abc.SalesServlet">
  <servlet>
  <servlet-class name="SalesServlet">
    com.abc.SalesServlet
  </servlet-class>
</servlet>
e <servlet>
  <servlet-name class="com.abc.SalesServlet">
    SalesServlet
  </servlet-name>
</servlet>

```

Answer: c

Explanation

A servlet is configured using the `servlet` element. Here is the definition of the `servlet` element:

```

<!ELEMENT servlet ( icon?, servlet-name, display-name?,
                    description?,
                    (servlet-class|jsp-file), init-param*,
                    load-on-startup?,
                    run-as?, security-role-ref* ) >

```

3. A web application is located in a directory named `sales`. Where should its deployment descriptor be located? (Select one)
- a `sales`
 - b `sales/deployment`
 - c `sales/WEB`
 - d `sales/WEB-INF`
 - e `WEB-INF/sales`
 - f `WEB-INF`
 - g `WEB/sales`

Answer: d

Explanation

The deployment descriptor is always located in the `WEB-INF` directory of the web application.

4. What file is the deployment descriptor of a web application named `BankApp` stored in? (Select one)
- a `BankApp.xml`
 - b `bankapp.xml`

- c server.xml
- d deployment.xml
- e WebApp.xml
- f web.xml

Answer: f

Explanation

The deployment descriptor of a web application is always kept in a file named `web.xml`, no matter what the name of the web application is.

5. Your servlet class depends on a utility class named `com.abc.TaxUtil`. Where would you keep the `TaxUtil.class` file? (Select one)
- a WEB-INF
 - b WEB-INF/classes
 - c WEB-INF/lib
 - d WEB-INF/jars
 - e WEB-INF/classes/com/abc

Answer: e

Explanation

All the classes that are not packaged in a JAR file must be kept in the `WEB-INF/classes` directory with its complete directory structure, as per the package of the class. The servlet container automatically adds this directory to the classpath of the web application.

6. Your web application, named `simpletax`, depends on a third-party JAR file named `taxpackage.jar`. Where would you keep this file? (Select one)
- a simpletax
 - b simpletax/WEB-INF
 - c simpletax/WEB-INF/classes
 - d simpletax/WEB-INF/lib
 - e simpletax/WEB-INF/jars
 - f simpletax/WEB-INF/thirdparty

Answer: d

Explanation

All the classes that are packaged in a JAR file must be kept in the `WEB-INF/lib` directory. The servlet container automatically adds all the classes in all the JAR files kept in this directory to the classpath of the web application.

7. Which of the following deployment descriptor elements is used to specify the initialization parameters for a servlet named `TestServlet`? (Select one)
- a No element is needed because initialization parameters are specified as attributes of the `<servlet>` element.
 - b `<servlet-param>`
 - c `<param>`
 - d `<initialization-param>`

- e <init-parameter>
- f <init-param>

Answer: f

Explanation

Each initialization must be specified using a separate <init-param> element:

```
<!ELEMENT init-param (param-name, param-value, description?)>
```

The following is an example of a servlet definition that specifies two initialization parameters:

```
<servlet>
  <servlet-name>TestServlet</servlet-name>
  <servlet-class>com.abc.TestServlet</servlet-class>
  <init-param>
    <param-name>MAX</param-name>
    <param-value>100</param-value>
    <description>maximum limit</description>
  </init-param>
  <init-param>
    <param-name>MIN</param-name>
    <param-value>10</param-value>
  </init-param>
</servlet>
```

8. Assume that the following servlet mapping is defined in the deployment descriptor of a web application:

```
<servlet-mapping>
  <servlet-name>TestServlet</servlet-name>
  <url-pattern>*.asp</url-pattern>
</servlet-mapping>
```

Which of the following requests will not be serviced by TestServlet?
(Select one)

- a /hello.asp
- b /gui/hello.asp
- c /gui/hello.asp/bye.asp
- d /gui/*.asp
- e /gui/sales/hello.asp
- f /gui/asp

Answer: f

Explanation

Here, any request that ends with .asp will be directed to TestServlet. Thus, only answer f will not be serviced by TestServlet. We suggest that you identify the context path, servlet path, and path info for all the above options according to the rules given in section 5.2.4.

CHAPTER 6—THE SERVLET CONTAINER MODEL

1. Which of the following methods will be invoked when a `ServletContext` is destroyed? (Select one)
 - a `contextDestroyed()` of `javax.servlet.ServletContextListener`
 - b `contextDestroyed()` of `javax.servlet.HttpServletContextListener`
 - c `contextDestroyed()` of `javax.servlet.http.ServletContextListener`
 - d `contextDestroyed()` of `javax.servlet.http.HttpServletContextListener`

Answer: a

Explanation

Remember that the concept of servlet context applies to all the servlets and not just `HttpServlets`. Therefore, interfaces related to servlet context belong to the `javax.servlet` package.

2. Which of the following methods will be invoked when a `ServletContext` is created? (Select one)
 - a `contextInstantiated()` of `javax.servlet.ServletContextListener`
 - b `contextInitialized()` of `javax.servlet.ServletContextListener`
 - c `contextInited()` of `javax.servlet.ServletContextListener`
 - d `contextCreated()` of `javax.servlet.ServletContextListener`

Answer: b

Explanation

On the exam, you will be asked questions that require you to know the method names for all the methods of the servlet API. As in this question, the options may be very confusing.

3. Consider the following class:

```
import javax.servlet.*;
public class MyListener implements ServletContextAttributeListener
{
    public void attributeAdded(ServletContextAttributeEvent scab)
    {
        System.out.println("attribute added");
    }

    public void attributeRemoved(ServletContextAttributeEvent scab)
    {
        System.out.println("attribute removed");
    }
}
```

Which of the following statements about the above class is correct? (Select one)

- a This class will compile as is.
- b This class will compile only if the `attributeReplaced()` method is added to it.

- c This class will compile only if the `attributeUpdated()` method is added to it.
- d This class will compile only if the `attributeChanged()` method is added to it.

Answer: b

Explanation

`ServletContextAttributeListener` also declares the public void `attributeReplaced(ServletContextAttributeEvent scab)` method, which is called when an existing attribute is replaced by another one.

4. Which method is used to retrieve an attribute from a `ServletContext`? (Select one)
- a `String getAttribute(int index)`
 - b `String getObject(int index)`
 - c `Object getAttribute(int index)`
 - d `Object getObject(int index)`
 - e `Object getAttribute(String name)`
 - f `String getAttribute(String name)`
 - g `String getObject(String name)`

Answer: e

Explanation

Since we can store any type of object in a servlet context, the `getAttribute()` method returns an object. You can then cast the returned object to whatever type you expect it to be.

5. Which method is used to retrieve an initialization parameter from a `ServletContext`? (Select one)
- a `Object getInitParameter(int index)`
 - b `Object getParameter(int index)`
 - c `Object getInitParameter(String name)`
 - d `String getInitParameter(String name)`
 - e `String getParameter(String name)`

Answer: d

Explanation

Initialization parameters are specified in the deployment descriptor. Since we can only specify strings in the deployment descriptor, the `getInitParameter()` method returns a `String`.

6. Which deployment descriptor element is used to specify a `ServletContextListener`? (Select one)
- a `<context-listener>`
 - b `<listener>`
 - c `<servlet-context-listener>`
 - d `<servletcontextlistener>`
 - e `<servletcontext-listener>`

Answer: b

Explanation

All the listeners that are specified in the deployment descriptor are specified using the `<listener>` element:

```
<listener>
  <listener-class>com.abc.MyServletContextListener</listener-class>
</listener>
```

The servlet container automatically figures out the type of interface that the specified class implements.

7. Which of the following `web.xml` snippets correctly specify an initialization parameter for a servlet context? (Select one)

- a `<context-param>`
 `<name>country</name>`
 `<value>USA</value>`
 `<context-param>`
- b `<context-param>`
 `<param name="country" value="USA" />`
 `<context-param>`
- c `<context>`
 `<param name="country" value="USA" />`
 `<context>`
- d `<context-param>`
 `<param-name>country</param-name>`
 `<param-value>USA</param-value>`
 `<context-param>`

Answer: d

Explanation

Initialization parameters for the servlet context are specified using the `<context-param>` element, which contains exactly one `<param-name>` and exactly one `<param-value>` element.

8. Which of the following is not a requirement of a distributable web application? (Select one)

- a It cannot depend on the notification events generated due to changes in the `ServletContext` attribute list.
- b It cannot depend on the notification events generated due to changes in the session attribute list.
- c It cannot depend on the notification events generated when a session is activated or passivated.
- d It cannot depend on the notification events generated when `ServletContext` is created or destroyed.
- e It cannot depend on the notification events generated when a session is created or destroyed.

Answer: c

Explanation

A servlet container may not propagate `ServletContextEvents` (generated when a context is created or destroyed) and `ServletContextAttributeEvents` (generated when the attribute list of a context changes) to listeners residing in other JVMs. This means that your web application cannot depend on these notifications. The same is true for events generated when a session is created or destroyed and when the attribute list of a session changes.

A session resides in only one JVM at a time. So, all the session attributes that implement `HttpSessionActivationListener` receive notifications when the session is activated or passivated.

9. Which of the following is a requirement of a distributable web application? (Select one)
- a It cannot depend on `ServletContext` for sharing information.
 - b It cannot depend on the `sendRedirect()` method.
 - c It cannot depend on the `include()` and `forward()` methods of the `RequestDispatcher` class.
 - d It cannot depend on cookies for session management.

Answer: a

Explanation

Since each JVM has a separate instance of a servlet context for each web application (except the default one), the attribute set in a `ServletContext` on one JVM will not be visible in the `ServletContext` for the same application on another JVM.

CHAPTER 7—USING FILTERS

1. Which elements are allowed in the `<filter-mapping>` element of the deployment descriptor? (Select three)
- a `<servlet-name>`
 - b `<filter-class>`
 - c `<dispatcher>`
 - d `<url-pattern>`
 - e `<filter-chain>`

Answers: a, c, and d

Explanation

Answer a is correct because you can map filters to named servlets, as well as URL patterns. Answer c will control under which dispatching mechanism the filter is invoked. Answer d allows you to map the filter to an URL pattern. Answer b is a legitimate element, but it belongs in the `<filter>` element. Answer e is a nonexistent element.

2. What is wrong with the following code?

```

public void doFilter(ServletRequest req, ServletResponse res,
    FilterChain chain)
    throws ServletException, IOException {

    chain.doFilter(req, res);
    HttpServletRequest request = (HttpServletRequest)req;
    HttpSession session = request.getSession();
    if (session.getAttribute("login") == null) {
        session.setAttribute("login", new Login());
    }
}

```

- a The doFilter() method signature is incorrect; it should take HttpServletRequest and HttpServletResponse.
- b The doFilter() method should also throw ServletException.
- c The call to chain.doFilter(req, res) should be this.doFilter(req, res, chain).
- d Accessing the request after chain.doFilter() results in an IllegalStateException.
- e Nothing is wrong with this filter.

Answer: e

Explanation

Answers a and b are wrong; the doFilter() method's signature and thrown exceptions are correct. Answer c, calling this.doFilter(req, res, chain), would result in unwanted recursion. Answer d is incorrect; no code here will throw an IllegalStateException.

3. Given these filter mapping declarations:

```

<filter-mapping>
  <filter-name>FilterOne</filter-name>
  <url-pattern>/admin/*</url-pattern>
  <dispatcher>FORWARD</dispatcher>
</filter-mapping>
<filter-mapping>
  <filter-name>FilterTwo</filter-name>
  <url-pattern>/users/*</url-pattern>
</filter-mapping>
<filter-mapping>
  <filter-name>FilterThree</filter-name>
  <url-pattern>/admin/*</url-pattern>
</filter-mapping>
<filter-mapping>
  <filter-name>FilterTwo</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>

```

in what order are the filters invoked for the following browser request?

/admin/index.jsp

- a FilterOne, FilterThree
- b FilterOne, FilterTwo, FilterThree

- c FilterThree, FilterTwo
- d FilterThree, FilterTwo
- e FilterThree
- f None of these filters are invoked.

Answer: d

Explanation

FilterOne is cannot be invoked by calling `/admin/index.jsp` from the browser but only through a request dispatcher forward, so answers a and b are incorrect. Answer c lists the correct filters but in the wrong order. Answer e names only the first filter that is invoked. Answer d names the correct filters, in the correct order.

CHAPTER 8—SESSION MANAGEMENT

1. Which of the following interfaces or classes is used to retrieve the session associated with a user? (Select one)
 - a `GenericServlet`
 - b `ServletConfig`
 - c `ServletContext`
 - d `HttpServlet`
 - e `HttpServletRequest`
 - f `HttpServletResponse`

Answer: e

Explanation

The session associated with a user can only be retrieved using the `HttpServletRequest.getSession()` method.

2. Which of the following code snippets, when inserted in the `doGet()` method, will correctly count the number of GET requests made by a user? (Select one)
 - a

```
HttpSession session = request.getSession();
int count = session.getAttribute("count");
session.setAttribute("count", count++);
```
 - b

```
HttpSession session = request.getSession();
int count = (int) session.getAttribute("count");
session.setAttribute("count", count++);
```
 - c

```
HttpSession session = request.getSession();
int count = ((Integer) session.getAttribute("count")).intValue();
session.setAttribute("count", count++);
```
 - d

```
HttpSession session = request.getSession();
int count = ((Integer) session.getAttribute("count")).intValue();
session.setAttribute("count", new Integer(count++));
```

Answer: d

Explanation

Remember that the `setAttribute()` and `getAttribute()` methods only work with objects and not with primitive data types. The `getAttribute()` method returns an object and so you need to cast the returned value to the actual type (`Integer`, in this case). Similarly, you need to wrap the count variable into an `Integer` object and pass it to the `setAttribute()` method.

3. Which of the following methods will be invoked on a session attribute that implements `HttpSessionBindingListener` when the session is invalidated? (Select one)
- a `sessionDestroyed`
 - b `valueUnbound`
 - c `attributeRemoved`
 - d `sessionInvalidated`

Answer: b

Explanation

When a session is invalidated, all the session attributes are unbound from the session. In the process, if an attribute implements `HttpSessionBindingListener`, the `valueUnbound()` method will be called on the attribute.

4. Which of the following methods will be invoked on a session attribute that implements appropriate interfaces when the session is invalidated? (Select one)
- a `sessionDestroyed` of `HttpSessionListener`
 - b `attributeRemoved` of `HttpSessionAttributeListener`
 - c `valueUnbound` of `HttpSessionBindingListener`
 - d `sessionWillPassivate` of `HttpSessionActivationListener`

Answer: c

Explanation

Only `HttpSessionListeners` and `HttpSessionAttributeListeners` that are configured in the deployment descriptor will receive notification of the events related to each of those interfaces. Therefore, even if a session attribute implements these interfaces, the `sessionDestroyed()` and `attributeRemoved()` methods will not be called on that attribute. `sessionWillPassivate()` is not called when a session is invalidated. The correct answer is therefore c.

5. Which of the following methods will expunge a session object? (Select one)
- a `session.invalidate();`
 - b `session.expunge();`
 - c `session.destroy();`
 - d `session.end();`
 - e `session.close();`

Answer: a

Explanation

The `invalidate()` method of `HttpSession` invalidates (or expunges) the session object.

6. Which of the following method calls will ensure that a session will never be expunged by the servlet container? (Select one)

- a `session.setTimeout(0);`
- b `session.setTimeout(-1);`
- c `session.setTimeout(Integer.MAX_VALUE);`
- d `session.setTimeout(Integer.MIN_VALUE);`
- e None of these.

Answer: e

Explanation

The correct method is `HttpSession.setMaxInactiveInterval(int seconds);`. A negative value (for example, `setMaxInactiveInterval(-1)`) ensures that the session is never invalidated. However, calling this method affects only the session on which it is called. All other sessions behave normally.

7. How can you make sure that none of the sessions associated with a web application will ever be expunged by the servlet container? (Select one)

- a `session.setMaxInactiveInterval(-1);`
- b Set the session timeout in the deployment descriptor to `-1`.
- c Set the session timeout in the deployment descriptor to `0` or `-1`.
- d Set the session timeout in the deployment descriptor to `65535`.
- e You have to change the timeout value of all the sessions explicitly as soon as they are created.

Answer: c

Explanation

The `setMaxInactiveInterval(-1)` method will only affect the session on which it is called. The `<session-config>` element of `web.xml` affects all the sessions of the web application. A value of `0` or less ensures that the sessions are never invalidated.

```
<web-app>
...
  <session-config>
    <session-timeout>0</session-timeout>
  </session-config>
...
</web-app>
```

8. In which of the following situations will a session be invalidated? (Select two)

- a No request is received from the client for longer than the session timeout period.
- b The client sends a `KILL_SESSION` request.
- c The servlet container decides to invalidate a session due to overload.
- d The servlet explicitly invalidates the session.

- e A user closes the active browser window.
- f A user closes all of the browser windows.

Answers: a and d

Explanation

Sessions will be invalidated only in two cases: when no request comes from the client for more than the session timeout period or when you call the `session.invalidate()` method on a session. Closing the browser windows does not actually invalidate the session. Even if you close all the browser windows, the session will still be active on the server. The servlet container will only invalidate the session after the timeout period of the session expires.

9. Which method is required for using the URL rewriting mechanism of implementing session support? (Select one)
- a `HttpServletRequest.encodeURL()`
 - b `HttpServletRequest.rewriteURL()`
 - c `HttpServletResponse.encodeURL()`
 - d `HttpServletResponse.rewriteURL()`

Answer: c

Explanation

In URL rewriting, the session ID has to be appended to all the URLs. The `encodeURL(String url)` method of `HttpServletResponse` does that.

10. The users of your web application do not accept cookies. Which of the following statements are correct? (Select one)
- a You cannot maintain client state.
 - b URLs displayed by static HTML pages may not work properly.
 - c You cannot use URL rewriting.
 - d You cannot set session timeout explicitly.

Answer: b

Explanation

If cookies are not supported, you can maintain the state using URL rewriting. Thus, answers a and c are incorrect. URL rewriting requires the session ID to be appended to all the URLs; however, static HTML pages will not have any session ID in the URLs that they display, so they may not work properly. Thus, answer b is correct.

Once the session is available, it does not matter whether it is maintained using cookies or URL rewriting. You can call all the methods as you would normally would, including `session.setMaxInactiveInterval()`. Therefore, answer d is wrong.

CHAPTER 9—DEVELOPING SECURE WEB APPLICATIONS

1. Which of the following correctly defines data integrity? (Select one)
 - a It guarantees that information is accessible only to certain users.
 - b It guarantees that the information is kept in encrypted form on the server.
 - c It guarantees that unintended parties cannot read the information during transmission between the client and the server.
 - d It guarantees that the information is not altered during transmission between the client and the server.

Answer: d

Explanation

Answers a and c describe authorization and confidentiality. Encrypting data kept on the server may be part of some security plans, but is not covered by the servlet specification.

2. What is the term for determining whether a user has access to a particular resource? (Select one)
 - a Authorization
 - b Authentication
 - c Confidentiality
 - d Secrecy

Answer: a

Explanation

Authentication is the process of identifying a user. Confidentiality ensures that third parties cannot eavesdrop on client-server communication. Encrypting communications between the client and server can prevent secrecy attacks.

3. Which one of the following must be done before authorization takes place? (Select one)
 - a Data validation
 - b User authentication
 - c Data encryption
 - d Data compression

Answer: b

Explanation

First, a user is authenticated. Once the identity of the user is determined using any of the authentication mechanisms, authorization is determined on a per-resource basis.

4. Which of the following actions would you take to prevent your web site from being attacked? (Select three)
 - a Block network traffic at all the ports except the HTTP port.
 - b Audit the usage pattern of your server.
 - c Audit the Servlet/JSP code.

- d Use HTTPS instead of HTTP.
- e Design and develop your web application using a software engineering methodology.
- f Use design patterns.

Answers: a, c, and d

Explanation

Answer a is correct because this will prevent network congestion and will close all possible entry points to the server except HTTP. Answer b seems correct, but it is wrong because auditing the usage pattern will help you in finding out the culprits only *after* the site has been attacked—it will not prevent an attack. Answer c is correct because auditing the Servlet/JSP code will ensure that no malicious code exists inside your server that can open a backdoor for hackers. Answer d is correct because HTTPS will prevent hackers from sniffing the communication between the clients and the server, thereby preventing the leakage of sensitive information such as usernames and passwords. Answers e and f are good for developing an industrial-strength system but are not meant for making a system attack proof.

5. Identify the authentication mechanisms that are built into the HTTP specification. (Select two)
- a Basic
 - b Client-Cert
 - c FORM
 - d Digest
 - e Client-Digest
 - f HTTPS

Answers: a and d

Explanation

The HTTP specification only defines Basic and Digest authentication mechanisms.

6. Which of the following deployment descriptor elements is used for specifying the authentication mechanism for a web application? (Select one)
- a security-constraint
 - b auth-constraint
 - c login-config
 - d web-resource-collection

Answer: c

Explanation

The authentication mechanism is specified using the `login-config` element; for example:

```
<login-config>
  <auth-method>FORM</auth-method>
  <realm-name>sales</realm-name>
```

```

    <form-login-config>
      <form-login-page>/formlogin.html</form-login-page>
      <form-error-page>/formerror.html</form-error-page>
    </form-login-config>
  </login-config>

```

The `security-constraint`, `auth-constraint`, and `web-resource-collection` elements are used for specifying the authorization details of the resources.

7. Which of the following elements are used for defining a security constraint? Choose only those elements that come directly under the `security-constraint` element. (Select three)
- a `login-config`
 - b `role-name`
 - c `role`
 - d `transport-guarantee`
 - e `user-data-constraint`
 - f `auth-constraint`
 - g `authorization-constraint`
 - h `web-resource-collection`

Answers: e, f, and h

Explanation

Remember that, logically, you need three things to define a security constraint: a collection of resources (i.e., `web-resource-collection`), a list of roles who are authorized to access the collection of resources (i.e., `auth-constraint`), and finally, the way the application data has to be transmitted between the clients and the server (i.e., `user-data-constraint`).

The following is the definition of the `security-constraint` element:

```

<!ELEMENT security-constraint (display-name?,
                               web-resource-collection+,
                               auth-constraint?, user-data-constraint?)>

```

8. Which of the following `web.xml` snippets correctly identifies all HTML files under the `sales` directory? (Select two)
- a

```
<web-resource-collection>
  <web-resource-name>reports</web-resource-name>
  <url-pattern>/sales/*.html</url-pattern>
</web-resource-collection>
```
 - b

```
<resource-collection>
  <web-resource-name>reports</web-resource-name>
  <url-pattern>/sales/*.html</url-pattern>
</resource-collection>
```
 - c

```
<resource-collection>
  <resource-name>reports</resource-name>
  <url-pattern>/sales/*.html</url-pattern>
</resource-collection>
```

```

d <web-resource-collection>
  <web-resource-name>reports</web-resource-name>
  <url-pattern>/sales/*.html</url-pattern>
  <http-method>GET</http-method>
</web-resource-collection>

```

Answers: a and d

Explanation

A collection of web resources is defined using the `web-resource-collection` element, which is defined as follows:

```

<!ELEMENT web-resource-collection (web-resource-name, description?,
                                   url-pattern*, http-method*)>

```

Observe that `http-method` is optional. The absence of the `http-method` element is equivalent to specifying all HTTP methods.

9. You want your `PerformanceReportServlet` to be accessible only to managers. This servlet generates a performance report in the `doPost()` method based on a FORM submitted by a user. Which of the following correctly defines a security constraint for this purpose? (Select one)

a <security-constraint>

```

<web-resource-collection>
<web-resource-name>performance report</web-resource-name>
<url-pattern>/servlet/PerformanceReportServlet</url-pattern>
<http-method>GET</http-method>
</web-resource-collection>

```

```

<auth-constraint>

```

```

<role-name>manager</role-name>
</auth-constraint>

```

```

<user-data-constraint>
<transport-guarantee>NONE</transport-guarantee>
</user-data-constraint>

```

```

</security-constraint>

```

b <security-constraint>

```

<web-resource-collection>
<web-resource-name>performance report</web-resource-name>
<url-pattern>/servlet/PerformanceReportServlet</url-pattern>

```

```

<http-method>*</http-method>
</web-resource-collection>

```

```

<accessibility>
<role-name>manager</role-name>
</accessibility>

```

```

    <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    </user-data-constraint>

</security-constraint>

c <security-constraint>

    <web-resource-collection>
    <web-resource-name>performance report</web-resource-name>
    <url-pattern>/servlet/PerformanceReportServlet</url-pattern>
    <http-method>POST</http-method>
    </web-resource-collection>

    <accessibility>
    <role-name>manager</role-name>
    </accessibility>

    <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    </user-data-constraint>

</security-constraint>

d <security-constraint>

    <web-resource-collection>
    <web-resource-name>performance report</web-resource-name>
    <url-pattern>/servlet/PerformanceReportServlet</url-pattern>
    <http-method>POST</http-method>
    </web-resource-collection>

    <auth-constraint>
    <role-name>manager</role-name>
    </auth-constraint>

```

Answer: d

Explanation

Since the question states that the servlet generates the report in the `doPost()` method, either the `<http-method>` must specify POST or there should be no `<http-method>` (which means the restriction applies to all the methods). Thus, answer a is incorrect. Further, the question states that the report should be accessible only to managers. This needs to be specified using the `<auth-constraint>` element. There is no such element as `<accessibility>`. Therefore, answers b and c are incorrect. Answer d is correct because both of the above requirements are satisfied. The question does not say anything about the `<user-data-constraint>` element, which is optional anyway.

10. Which of the following statements regarding authentication mechanisms are correct? (Select two)
- a The HTTP Basic mechanism transmits the username/password "in the open."
 - b The HTTP Basic mechanism uses HTML FORMs to collect usernames/passwords.
 - c The transmission method in the Basic and FORM mechanisms is the same.

- d The method of capturing the usernames/passwords in the Basic and FORM mechanisms is the same.

Answers: a and c

Explanation

The HTTP Basic mechanism uses a browser-specific way (usually a dialog box) to capture the username and password, while the FORM mechanism uses an HTML FORM to do the same. However, both mechanisms transmit the captured values in clear text without any encryption. Therefore, answers a and c are correct.

11. Which of the following statements are correct for an unauthenticated user? (Select two)

- a `HttpServletRequest.getUserPrincipal()` returns null.
- b `HttpServletRequest.getUserPrincipal()` throws `SecurityException`.
- c `HttpServletRequest.isUserInRole()` returns false.
- d `HttpServletRequest.getRemoteUser()` throws a `SecurityException`.

Answers: a and c

Explanation

None of the three methods—`getUserPrincipal()`, `isUserInRole()`, and `getRemoteUser()`—throws an exception. We suggest you read the description of these methods in the JavaDocs.

CHAPTER 10—THE JSP TECHNOLOGY MODEL—THE BASICS

1. Consider the following code and select the correct statement about it from the options below. (Select one)

```
<html><body>
  <%! int aNum=5 %>
  The value of aNum is <%= aNum %>
</body></html>
```

- a It will print "The value of aNum is 5" to the output.
- b It will flag a compile-time error because of an incorrect declaration.
- c It will throw a runtime exception while executing the expression.
- d It will not flag any compile time or runtime errors and will not print anything to the output.

Answer: b

Explanation

It will flag a compile-time error because the variable declaration `<%! int aNum=5 %>` is missing a `;` at the end. It should be

```
<%! int aNum=5; %>
```

2. Which of the following tags can you use to print the value of an expression to the output stream? (Select two)

- a `<%@ %>`
- b `<%! %>`
- c `<% %>`
- d `<%= %>`
- e `<%-- --%>`

Answers: c and d

Explanation

You can use a JSP expression to print the value of an expression to the output stream. For example, if the expression is `x+3`, you can write `<%= x+3 %>`. Answer d is a JSP expression and is therefore the correct answer. But you can also use a scriptlet to print the value of an expression to the output stream as `<% out.print(x+3); %>`. Answer c is a scriptlet and is therefore also correct. If the exam asks you to select one correct option, then select the expression syntax, as in answer d. But if the exam asks for two correct answers, then select the scriptlet syntax as well.

3. Which of the following methods is defined by the JSP engine? (Select one)
- a `jspInit()`
 - b `_jspService()`
 - c `_jspService(ServletRequest, ServletResponse)`
 - d `_jspService(HttpServletRequest, HttpServletResponse)`
 - e `jspDestroy()`

Answer: d

Explanation

The `_jspService()` method of the `javax.servlet.jsp.HttpJspPage` class is defined by the JSP engine. `HttpJspPage` is meant to serve HTTP requests, and therefore the `_jspService()` method accepts the `javax.servlet.http.HttpServletRequest` and `javax.servlet.http.HttpServletResponse` parameters.

4. Which of the following exceptions may be thrown by the `_jspService()` method? (Select one)
- a `javax.servlet.ServletException`
 - b `javax.servlet.jsp.JSPException`
 - c `javax.servlet.ServletException` and `javax.servlet.jsp.JSPException`
 - d `javax.servlet.ServletException` and `java.io.IOException`
 - e `javax.servlet.jsp.JSPException` and `java.io.IOException`

Answer: d

Explanation

The `_jspService()` method may throw a `javax.servlet.ServletException`, a `java.io.IOException`, or a subclass of these two exception classes. Note that the `_jspService()` method does not define `javax.servlet.jsp.JspException` in its throws clause.

5. Write the name of the method that you can use to initialize variables declared in a JSP declaration in the space provided. (Write only the name of the method. Do not write the return type, parameters, or parentheses.)

a [_____]

Answer: `jspInit`

Explanation

The `jspInit()` method is the first method called by the JSP engine on a JSP page. It is called only once to allow the page to initialize itself. You can use this method to initialize variables declared in JSP declarations (`<%! %>`).

6. Which of the following correctly declares that the current page is an error page and also enables it to take part in a session? (Select one)

- a `<%@ page pageType="errorPage" session="required" %>`
- b `<%@ page isErrorPage="true" session="mandatory" %>`
- c `<%@ page errorPage="true" session="true" %>`
- d `<%@ page isErrorPage="true" session="true" %>`
- e None of the above.

Answer: *d*

Explanation

The `isErrorPage` attribute accepts a Boolean value (`true` or `false`) and indicates whether the current page is capable of handling errors. The `session` attribute accepts a Boolean value (`true` or `false`) and indicates whether the current page must take part in a session. Therefore, answer *d* is correct. Since the `pageType` attribute is not a valid attribute for a `page` directive, answer *a* is not correct. The `mandatory` value is not a valid value for the `session` attribute, which means answer *b* is not correct. The `errorPage` attribute is a valid attribute, but it is used for specifying another page as an error handler for the current page. Therefore, answer *c* is also incorrect.

CHAPTER 11—THE JSP TECHNOLOGY MODEL—ADVANCED TOPICS

1. What will be the output of the following code? (Select one)

```
<html><body>
  <% x=3; %>
  <% int x=5; %>
  <%! int x=7; %>
  x = <%=x%>, <%=this.x%>
</body></html>
```

- a `x = 3, 5`
- b `x = 3, 7`
- c `x = 5, 3`
- d `x = 5, 7`
- e Compilation error

Answer: *c*

Explanation

The above code will translate to servlet code similar to the following:

```
public class ...
{
    int x = 7;

    public void _jspService(...)
    {
        ...
        out.print("<html><body>");
        x = 3;
        int x = 5;
        out.write("x = "); out.print(x);
        out.write(","); out.print(this.x);
        out.print("</body></html>");
    }
}
```

The declaration will create a member variable `x` and initialize it to 7. The first scriptlet, `x=3`, will change its value to 3. Then, the second scriptlet will declare a local variable `x` and initialize it to 5. The first expression refers to the local variable `x` and will therefore print 5. The second expression uses the keyword `this` to refer to the member or instance variable `x`, which was set to 3. Thus, the correct answer is `c`, `x = 5, 3`.

2. What will be the output of the following code? (Select one)

```
<html><body>
    The value is <%= ""%>
</body></html>
```

- a Compilation error
- b Runtime error
- c The value is
- d The value is null

Answer: c

Explanation

The expression is converted to

```
out.print("");
```

Thus, the correct answer is `c`.

3. Which of the following implicit objects is not available to a JSP page by default? (Select one)

- a application
- b session
- c exception
- d config

Answer: c

Explanation

The implicit variables `application` and `config` are always available to a JSP page. The implicit variable `session` is available if the value of the page directive's `session` attribute is set to `true`. Since it is set to `true` by default, the implicit variable `session` is also available by default. The implicit variable `exception` is available only if the value of the page directive's `isErrorPage` attribute is set to `true`. It is set to `false` by default, so the implicit variable `exception` is not available by default. We have to explicitly set it to `true`:

```
<%@ page isErrorPage="true" %>
```

The correct answer, therefore, is c.

4. Which of the following implicit objects can you use to store attributes that need to be accessed from all the sessions of a web application? (Select two)
- a `application`
 - b `session`
 - c `request`
 - d `page`
 - e `pageContext`

Answers: a and e

Explanation

To store attributes that are accessible from all the sessions of a web application, we have to put them in the application scope. To achieve this, we have to use the implicit object `application`. If the exam asks you to select one answer, then select `application`. If the exam asks for two correct answers, then read the question carefully. It says, "Which of the following implicit objects *can you use* to store attributes that need to be accessed from all the sessions of a web application?" We can also use `pageContext` to store objects in the application scope as `pageContext.setAttribute("name", object, PageContext.APPLICATION_SCOPE);` and `pageContext.getAttribute("name", PageContext.APPLICATION_SCOPE);`.

5. The implicit variable `config` in a JSP page refers to an object of type: (Select one)
- a `javax.servlet.PageConfig`
 - b `javax.servlet.jsp.PageConfig`
 - c `javax.servlet.ServletConfig`
 - d `javax.servlet.ServletContext`

Answer: c

The implicit variable `config` in a JSP page refers to an object of type `javax.servlet.ServletConfig`.

6. A JSP page can receive context initialization parameters through the deployment descriptor of the web application.

- a True
- b False

Answer: a

Explanation

Context initialization parameters are specified by the `<context-param>` tags in `web.xml`. These parameters are for the whole web application and not specific to any servlet or JSP page. Thus, all components of a web application can access context initialization parameters.

7. Which of the following will evaluate to `true`? (Select two)

- a `page == this`
- b `pageContext == this`
- c `out instanceof ServletOutputStream`
- d `application instanceof ServletContext`

Answers: a and d

Explanation

The implicit variable `page` refers to the current servlet, and therefore answer a will evaluate to `true`. The `application` object refers to an object of type `ServletContext`, which means answer d will also evaluate to `true`. The `pageContext` object refers to an object of type `PageContext` and not to the servlet, which means answer b will evaluate to `false`. The `out` implicit variable refers to an instance of `javax.servlet.jsp.JspWriter` and not to an instance of `javax.servlet.ServletOutputStream`, so answer c evaluates to `false`. Note that `JspWriter` is derived from `java.io.Writer`, while `ServletOutputStream` is derived from `java.io.OutputStream`.

8. Select the correct statement about the following code. (Select one)

```
<%@ page language="java" %>
<html><body>
out.print("Hello ");
out.print("World ");
</body></html>
```

- a It will print `Hello World` in the output.
- b It will generate compile-time errors.
- c It will throw runtime exceptions.
- d It will only print `Hello`.
- e None of above.

Answer: e

Explanation

The lines `out.print("Hello ")` and `out.print("World ")` are not contained in a scriptlet (`<%...%>`). The JSP engine assumes they are a part of the template text and sends them to the browser without executing them on the server. Therefore, it will print the two statements in the browser window:

```
out.print("Hello ");out.print("World ");
```

9. Select the correct statement about the following code. (Select one)

```
<%@ page language="java" %>
<html><body>
<%
    response.getOutputStream().print ("Hello ");
    out.print ("World");
%>
</body></html>
```

- a It will print Hello World in the output.
- b It will generate compile-time errors.
- c It will throw runtime exceptions.
- d It will only print Hello.
- e None of above.

Answer: c

Explanation

As explained in chapter 4, “The Servlet model,” the `OutputStream` of a response object is used for sending binary data to the client while the `Writer` object is used for sending character data. However, we cannot use both on the same response object. Since the JSP engine automatically gets the `JspWriter` from the response object to output the content of the JSP as character data, the call to `getOutputStream()` throws a `java.lang.IllegalStateException`. Thus, the correct answer is c.

10. Which of the following implicit objects does not represent a scope container? (Select one)

- a application
- b session
- c request
- d page
- e pageContext

Answer: d

Explanation

The implicit objects `application`, `session`, and `request` represent the containers for the scopes, *application*, *session*, and *request*, respectively. The implicit object `page` refers to the generated Servlet and does not represent any scope container. The implicit object `pageContext` represents the page scope container, so the correct answer is d.

11. What is the output of the following code? (Select one)

```
<html><body>
<% int i = 10 ;%>
<% while(--i>=0) { %>
```

```
        out.print(i);
    <% } %>
</body></html>
```

- a 9876543210
- b 9
- c 0
- d None of the above.

Answer: d

Explanation

The statement `out.print(i)` is not inside a scriptlet and is part of the template text. The above JSP page will print

```
out.print(i);out.print(i);out.print(i);.....
```

ten times.

When in doubt, always convert a JSP code to its equivalent servlet code step by step:

```
out.write("<html><body>");
int i = 10;
while (--i>=) {
    out.write("out.print(i); ");
}
out.write("<html><body>");
```

12. Which of the following is not a valid XML-based JSP tag? (Select one)

- a `<jsp:directive.page />`
- b `<jsp:directive.include />`
- c `<jsp:directive.taglib />`
- d `<jsp:declaration></jsp:declaration>`
- e `<jsp:scriptlet></jsp:scriptlet>`
- f `<jsp:expression></jsp:expression>`

Answer: c

Explanation

The tag `<jsp:directive.taglib>` is not a valid XML-based tag. Remember that tag library information is provided in the `<jsp:root>` element.

13. Which of the following XML syntax format tags do not have an equivalent in JSP syntax format? (Select two)

- a `<jsp:directive.page/>`
- b `<jsp:directive.include/>`
- c `<jsp:text></jsp:text>`
- d `<jsp:root></jsp:root>`
- e `<jsp:param/>`

Answers: c and d

Explanation

The equivalent of `<jsp:directive.page/>` is `<%@ page %>`. The equivalent of `<jsp:directive.include/>` is `<%@ include %>`. The `<jsp:param/>` tag is the same for both the syntax formats. Thus, the correct answers are c and d. The tags `<jsp:text>` and `<jsp:root>` have no equivalent in the JSP syntax format.

14. Which of the following is a valid construct to declare that the implicit variable `session` should be made available to the JSP page? (Select one)

- a `<jsp:session>true</jsp:session>`
- b `<jsp:session required="true" />`
- c `<jsp:directive.page>`
 `<jsp:attribute name="session" value="true" />`
 `</jsp:directive.page>`
- d `<jsp:directive.page session="true" />`
- e `<jsp:directive.page attribute="session" value="true" />`

Answer: d

Explanation

The correct way to declare that the implicit variable `session` should be made available to the JSP page in XML format is shown in answer d: `<jsp:directive.page session="true" />`.

CHAPTER 12—REUSABLE WEB COMPONENTS

1. Which of the following JSP tags can be used to include the output of another JSP page into the output of the current page at request time? (Select one)

- a `<jsp:insert>`
- b `<jsp:include>`
- c `<jsp:directive.include>`
- d `<jsp:directive:include>`
- e `<%@ include %>`

Answer: b

Explanation

The tags in answers a and d are not valid JSP tags. Answers c and e are valid tags in XML syntax and JSP syntax, respectively, but they are directives and include other JSP pages or HTML/XML files at translation time. Answer b is the right answer because it includes the output of another component, JSP page, or Servlet at request time.

2. Consider the contents of the following two JSP files:

File 1: `test1.jsp`

```
<html><body>
  <% String message = "Hello"; %>

  //1 Insert LOC here.
```

```
        The message is <%= message %>
    </body></html>
```

File 2: test2.jsp

```
<% message = message + " world!"; %>
```

Which of the following lines can be inserted at //1 in test1.jsp so that it prints "The message is Hello world!" when requested? (Select one)

- a <%@ include page="test2.jsp" %>
- b <%@ include file="test2.jsp" />
- c <jsp:include page="test2.jsp" />
- d <jsp:include file="test2.jsp" />

Answer: b

Explanation

Since the test2.jsp file does not declare or define the variable message, it cannot compile on its own. This rules out dynamic inclusion using the <jsp:include> action. The file test1.jsp could print "Hello world" if it statically included test2.jsp. This could be done using the include directive: <%@ include %>. For the include directive, the valid attribute is file and not page, so answer b is correct.

3. Which of the following is a correct way to pass a parameter equivalent to the query string user=mary at request time to an included component? (Select one)

- a <jsp:include page="other.jsp" >
 <jsp:param paramName="user" paramValue="mary" />
 </jsp:include>
- b <jsp:include page="other.jsp" >
 <jsp:param name="mary" value="user" />
 </jsp:include>
- c <jsp:include page="other.jsp" >
 <jsp:param value="mary" name="user" />
 </jsp:include>
- d <jsp:include page="other.jsp" >
 <jsp:param param="user" value="mary"/>
 </jsp:include>
- e <jsp:include page="other.jsp" >
 <jsp:param user="mary" />
 </jsp:include>

Answer: c

Explanation

The only valid attributes that a <jsp:param> tag can have are name and value. This rules out answers a, d, and e. Answer b, <jsp:param name="mary" value="user" />, is equivalent to the query string mary=user. In the included component, a call to request.getParameter("mary"); will return "user".

Answer c, `<jsp:param value="mary" name="user" />`, is equivalent to the query string `user=mary`. In the included component, a call to `request.getParameter("user");` will return "mary". Therefore, answer c is the correct answer.

4. Identify the JSP equivalent of the following code written in a servlet. (Select one)

```
RequestDispatcher rd = request.getRequestDispatcher("world.jsp");
rd.forward(request, response);
```

- a `<jsp:forward page="world.jsp"/>`
- b `<jsp:action.forward page="world.jsp"/>`
- c `<jsp:directive.forward page="world.jsp"/>`
- d `<%@ forward file="world.jsp"%>`
- e `<%@ forward page="world.jsp"%>`

Answer: a

Explanation

The action tags in answers b through e are all invalid JSP tags. Answer a, `<jsp:forward page="relativeURL" />`, is the only valid way to write a forward action.

5. Consider the contents of two JSP files:

File 1: test1.jsp

```
<html><body>
  <% pageContext.setAttribute("ninetyNine", new Integer(99)); %>
  //1
</body></html>
```

File 2: test2.jsp

```
The number is <%= pageContext.getAttribute("ninetyNine") %>
```

Which of the following, when placed at line //1 in the test1.jsp file, will allow the test2.jsp file to print the value of the attribute when test1.jsp is requested? (Select one)

- a `<jsp:include page="test2.jsp" />`
- b `<jsp:forward page="test2.jsp" />`
- c `<%@ include file="test2.jsp" %>`
- d None of the above because objects placed in pageContext have the page scope and cannot be shared with other components.

Answer: c

Explanation

Objects placed in pageContext have the page scope and are accessible within a single translation unit. Since files included statically using the include directive become an integral part of the same translation unit as the including file, they can

share objects in the page scope via the `PageContext` container. So the correct answer is c: `<%@include file="test2.jsp" %>`.

6. Consider the contents of two JSP files:

File 1: `this.jsp`

```
<html><body><pre>
  <jsp:include page="that.jsp" >
    <jsp:param name="color" value="red" />
    <jsp:param name="color" value="green" />
  </jsp:include>
</pre></body></html>
```

File 2: `that.jsp`

```
<%
  String colors[] = request.getParameterValues("color");
  for (int i=0; i<colors.length; i++)
  {
    out.print(colors[i] + " ");
  }
%>
```

What will be the output of accessing the `this.jsp` file via the following URL?
(Select one)

`http://localhost:8080/chapter12/this.jsp?color=blue`

- a blue
- b red green
- c red green blue
- d blue red green
- e blue green red

Answer: c

Explanation

The parameters passed via the `<jsp:param>` tag to an included component tag take precedence over the parameters already present in the request object of the including component. Also, the order of values passed via the `<jsp:param>` tag is the same as the order in which the tags appear. Thus, the correct answer is c. The output will be red green blue.

7. Consider the contents of two JSP files:

File 1: `this.jsp`

```
<html><body>
  <%= request.getParameter("color") %>
  <jsp:include page="that.jsp" >
    <jsp:param name="color" value="red" />
  </jsp:include>
```

```
    <%= request.getParameter("color") %>
</body></html>
```

File 2: that.jsp

```
    <%= request.getParameter("color") %>
```

What will be the output of accessing the `this.jsp` file via the following URL?
(Select one)

```
http://localhost:8080/chapter12/this.jsp?color=blue
```

- a blue red blue
- b blue red red
- c blue blue red
- d blue red null

Answer: a

Explanation

The first call to `request.getParameter("color")` in the `this.jsp` file returns `blue`. This file then includes the `that.jsp` file and passes a value of `red` for the `color` attribute. Since the values passed via `<jsp:param>` take precedence over the original values, a call to `request.getParameter("color")` in `that.jsp` returns `red`. However, this new value exists and is available only within the included component—that is, `that.jsp`. So, after the `that.jsp` page finishes processing, a call to `request.getParameter("color")` in the `this.jsp` file again returns `blue`. Thus, the correct answer is a. The output will be `blue red blue`.

8. Consider the contents of three JSP files:

File 1: one.jsp

```
<html><body><pre>
    <jsp:include page="two.jsp" >
        <jsp:param name="color" value="red" />
    </jsp:include>
</pre></body></html>
```

File 2: two.jsp

```
<jsp:include page="three.jsp" >
    <jsp:param name="color" value="green" />
</jsp:include>
```

File 3: three.jsp

```
<%= request.getParameter("color") %>
```

What will be the output of accessing the `one.jsp` file via the following URL?
(Select one)

```
http://localhost:8080/chapter12/one.jsp?color=blue
```

- a red
- b green
- c blue
- d The answer cannot be determined.

Answer: b

Explanation

The output is generated by the `three.jsp` file. Since the `two.jsp` file calls `three.jsp` and provides a value of `green`, this value takes precedence over all the previous values passed to `one.jsp` and `two.jsp`. Thus, the correct answer is `b`. The output will be `green`.

CHAPTER 13—CREATING JSPs WITH THE EXPRESSION LANGUAGE (EL)

1. Consider the following code:

```
<html><body>
  ${ (5 + 3 + a > 0) ? 1 : 2 }
</body></html>
```

Select the correct statement from the options below:

- a It will print 1 because the statement is valid.
- b It will print 2 because the statement is valid.
- c It will throw an exception because `a` is undefined.
- d It will throw an exception because the expression's syntax is invalid.

Answer: a

Explanation

Using a letter in an EL statement is valid, and so is the statement's syntax. So answers `c` and `d` are incorrect. Since the condition evaluates to `true`, the first of the two results will be displayed, and `a` is the correct answer.

2. Which statement best expresses the purpose of a tag library descriptor (TLD) *in an EL function*?

- a It contains the Java code that will be compiled.
- b It invokes the Java method as part of the JSP.
- c It matches the tag library with a URI.
- d It matches function names to tags that can be used in the JSP.

Answer: d

Explanation

The `web.xml` file matches tag libraries with URIs, so answer `c` is incorrect. Java code is contained in a `.java` file and methods are invoked from within a JSP (`.jsp`). Thus, answers `a` and `b` are incorrect. A tag library descriptor (`.tld`) matches function names with names that can be used in JSPs, so answer `d` is correct. To be more specific, it matches a Java method signature with a class name for use in JSPs.

3. Which of the following variables is not available for use in EL expressions?
- a param
 - b cookie
 - c header
 - d pageContext
 - e contextScope

Answer: e

Explanation

These implicit variables are similar to others used in JSPs. The `param`, `header`, and `cookie` variables access the same objects as those used in request/response access. The `pageContext` variable provides access to scope variables. So answers a through d are incorrect. But there's no such thing as a `contextScope` variable—use `applicationScope` to access context-related information.

4. Which tags tell the web container where to find your TLD file in your filesystem?
- a `<taglib-directory></taglib-directory>`
 - b `<taglib-uri></taglib-uri>`
 - c `<taglib-location></taglib-location>`
 - d `<tld-directory></tld-directory>`
 - e `<taglib-name></taglib-name>`

Answer: c

Explanation

The `<taglib-name>` and `<taglib-uri>` elements provide the library's name and URI, but not its place in filesystem, so b and e are incorrect. There are no `<taglib-directory>` or `<tld-directory tags>`, which means answers a and d are incorrect. The `<taglib-location>` elements specify the directory containing the TLD, starting with `/WEB-INF/`. Therefore, the answer is c.

5. Which two of the following expressions won't return the header's `accept` field?
- a `${header.accept}`
 - b `${header[accept]}`
 - c `${header['accept']}`
 - d `${header["accept']}`
 - e `${header.'accept']}`

Answers: b and e

Explanation

EL arrays can contain strings as long as they are surrounded by single or double quotes. So c and d will return the `accept` field, and b won't. Similarly, the EL property operator (`.`) can contain a string as long as it isn't surrounded by quotes. So a will return a value and e will cause an error.

6. When writing a TLD, which tags would you use to surround `fnName(int num)`, a Java method declared in a separate class?

- a `<function-signature></function-signature>`
- b `<function-name></function-name>`
- c `<method-class></method-class>`
- d `<method-signature></method-signature>`
- e `<function-class></function-class>`

Answer: a

Explanation

Since EL describes these structures as functions, c and d are incorrect. Also, the function's name is contained in `<name>` tags, not `<function-name>`, so b is incorrect. Since `<function-signature>` declares the Java method, a is the correct answer.

7. Which of the following method signatures is usable in EL functions?
- a `public static expFun(void)`
 - b `expFun(void)`
 - c `private expFun(void)`
 - d `public expFun(void)`
 - e `public native expFun(void)`

Answer: a

Explanation

EL functions must refer to methods with modifiers of `public` and `static`. Therefore, a is the only acceptable answer. Similarly, the class containing the method must be `public`.

CHAPTER 14—USING JAVABEANS

1. Which of the following is a valid use of the `<jsp:useBean>` action? (Select one)
- a `<jsp:useBean id="address" class="AddressBean" />`
 - b `<jsp:useBean name="address" class="AddressBean"/>`
 - c `<jsp:useBean bean="address" class="AddressBean" />`
 - d `<jsp:useBean beanName="address" class="AddressBean" />`

Answer: a

Explanation

`name` and `bean` are not valid attributes for a `<jsp:useBean>` tag. Thus, answers b and c are incorrect. The `beanName` and `class` attributes cannot be used together, which means answer d is incorrect. Further, the `id` attribute is the mandatory attribute of `<jsp:useBean>` and is missing from answers b, c, and d. Therefore, only answer a is correct.

2. Which of the following is a valid way of getting a bean's property? (Select one)
- a `<jsp:useBean action="get" id="address" property="city" />`
 - b `<jsp:getProperty id="address" property="city" />`
 - c `<jsp:getProperty name="address" property="city" />`
 - d `<jsp:getProperty bean="address" property="*" />`

Answer: c

Explanation

The `<jsp:getProperty>` action has only two attributes—`name` and `property`—and both are mandatory. Therefore, answer c is correct.

3. Which of the following are valid uses of the `<jsp:useBean>` action? (Select two)

- a `<jsp:useBean id="address" class="AddressBean" name="address" />`
- b `<jsp:useBean id="address" class="AddressBean" type="AddressBean" />`
- c `<jsp:useBean id="address" beanName="AddressBean" class="AddressBean" />`
- d `<jsp:useBean id="address" beanName="AddressBean" type="AddressBean" />`

Answers: b and d

Explanation

Answer a is not correct because `name` is not a valid attribute in `<jsp:useBean>`. Answer c is not correct because `beanName` and `class` cannot be used together.

4. Which of the following gets or sets the bean in the `ServletContext` container object? (Select one)

- a `<jsp:useBean id="address" class="AddressBean" />`
- b `<jsp:useBean id="address" class="AddressBean" scope="application" />`
- c `<jsp:useBean id="address" class="AddressBean" scope="servlet" />`
- d `<jsp:useBean id="address" class="AddressBean" scope="session" />`
- e None of the above.

Answer: b

Explanation

The correct answer is b, because the `ServletContext` container represents the application scope. Answer a is not correct because if the scope is not specified, then the page scope is assumed by default.

5. Consider the following code:

```
<html><body>
<jsp:useBean id="address" class="AddressBean" scope="session" />
state = <jsp:getProperty name="address" property="state" />
</body></html>
```

Which of the following are equivalent to the third line above? (Select three)

- a `<% state = address.getState(); %>`
- b `<% out.write("state = "); out.print(address.getState()); %>`
- c `<% out.write("state = "); out.print(address.getstate()); %>`
- d `<% out.print("state = " + address.getState()); %>`
- e `state = <%= address.getState() %>`
- f `state = <%! address.getState(); %>`

Answers: b, d, and e

Explanation

The third line in the code prints "state = " followed by the actual value of the property in the output HTML. Answer a is incorrect, because it is inside a scriptlet and does not print any output. Answer c is incorrect because the standard convention that the beans follow is to capitalize the first character of the property's name. Therefore, it should be `getState()` and not `getstate()`. Answer f is incorrect because the method call `address.getState()` is in a declaration instead of an expression. Answers b, d, and e all do the same thing and are all equivalent to the third line of the code.

6. Which of the options locate the bean equivalent to the following action? (Select three)

```
<jsp:useBean id="address" class="AddressBean" scope="request" />
a request.getAttribute("address");
b request.getParameter("address");
c getServletContext().getRequestAttribute("address");
d pageContext.getAttribute("address", PageContext.REQUEST_SCOPE);
e pageContext.getRequest().getAttribute("address");
f pageContext.getRequestAttribute("address");
g pageContext.getRequestParameter("address");
```

Answers: a, d, and e

Explanation

Answer b is not correct because beans cannot be *get* or *set* as request parameters. They can be *get* and *set* as request attributes. Answer c is incorrect because `getServletContext()` returns an object of type `ServletContext` and `ServletContext` has nothing to do with the request scope. Answers f and g are incorrect because the methods `getRequestAttribute()` and `getRequestParameter()` do not exist in `PageContext`. Answer a is the simplest way to do that. Answers d and e achieve the same result using the `PageContext` object. `PageContext` is explained in chapter 11, “The JSP technology model—advanced topics.”

7. Consider the following code for `address.jsp`:

```
<html><body>
<jsp:useBean id="address" class="AddressBean" />
<jsp:setProperty name="address" property="city" value="LosAngeles" />
<jsp:setProperty name="address" property="city" />
<jsp:getProperty name="address" property="city" />
</body></html>
```

What is the output if the above page is accessed via the URL

```
http://localhost:8080/chap14/address.jsp?city=Chicago&city=Miami
```

Assume that the `city` property is not an indexed property. (Select one)

- a LosAngeles
- b Chicago

- c Miami
- d ChicagoMiami
- e LosAngelesChicagoMaimi
- f It will not print anything because the value will be null or "".

Answer: b

Explanation

The first `<jsp:setProperty>` action sets the value of the `city` property explicitly to `LosAngeles` using the `value` attribute. The second `<jsp:setProperty>` action overwrites this value with the value from the request parameter. Since it is not an indexed property, only the first value from the parameter is used. Thus, the correct answer is b, Chicago.

8. Consider the following code:

```
<html><body>

<{%>
<jsp:useBean id="address" class="AddressBean" scope="session" />
<%}%>

//1

</body></html>
```

Which of the following can be placed at line `//1` above to print the value of the `street` property? (Select one)

- a `<jsp:getProperty name="address" property="street" />`
- b `<% out.print(address.getStreet()); %>`
- c `<%= address.getStreet() %>`
- d `<%= ((AddressBean)session.getAttribute("address")).getStreet() %>`
- e None of the above; the bean is nonexistent at this point.

Answer: d

Explanation

The `<jsp:useBean>` declaration puts an object of type `AddressBean` in the session scope. The pair of curly braces (`{` and `}`) marks the scope of the variable `address` and, therefore, we cannot use answers a, b, and c. However, the object is still existent and is available in the session scope. This means we can use the implicit variable `session` to get the `address` object, typecast the returned value to `AddressBean`, and call its `getStreet()` method to print the `street` property. Therefore, answer d is correct.

9. Consider the following code:

```
<html><body>

<{%>
<jsp:useBean id="address" class="AddressBean" scope="session" />
<%}%>
```



```
<jsp:useBean id="address" class="AddressBean" scope="session" />
<jsp:getProperty name="address" property="street" />

</body></html>
```

Which of the following is true about the above code? (Select one)

- a It will give translation-time errors.
- b It will give compile-time errors.
- c It may throw runtime exceptions.
- d It will print the value of the street property.

Answer: a

Explanation

A translation time error will occur if we use the same ID more than once in the same translation unit.

10. Consider the following servlet code:

```
//...

public void service (HttpServletRequest request,
                    HttpServletResponse response)
    throws IOException, ServletException
{
    //1
}
```

Which of the following can be used at //1 to retrieve a JavaBean named address present in the application scope? (Select one)

- a `getServletContext().getAttribute("address");`
- b `application.getAttribute("address");`
- c `request.getAttribute("address",APPLICATION_SCOPE);`
- d `pageContext.getAttribute("address",APPLICATION_SCOPE);`

Answer: a

Explanation

The implicit variables are automatically available in the `_jspService()` method of a JSP page but are not defined automatically in a Servlet class. We cannot use the implicit variables `application` and `pageContext` because the code is part of a servlet's `service()` method and not of a JSP page. So, answers b and d are not correct. There is no such method as `getAttribute(String, int)` in `HttpServletRequest` that accepts an integer to identify scopes. Therefore, answer c is also incorrect.

11. Consider the following code, contained in a file called `this.jsp`:

```
<html><body>
<jsp:useBean id="address" class="AddressBean" />
<jsp:setProperty name="address" property="*" />
<jsp:include page="that.jsp" />
</body></html>
```

Which of the following is true about the `AddressBean` instance declared in this code? (Select one)

- a The bean instance will not be available in `that.jsp`.
- b The bean instance may or may not be available in `that.jsp`, depending on the threading model implemented by `that.jsp`.
- c The bean instance will be available in `that.jsp`, and the `that.jsp` page can print the values of the beans properties using `<jsp:getProperty />`.
- d The bean instance will be available in `that.jsp` and the `that.jsp` page can print the values of the bean's properties using `<jsp:getProperty />` only if `that.jsp` also contains a `<jsp:useBean />` declaration identical to the one in `this.jsp` and before using `<jsp:getProperty />`.

Answer: a

Explanation

By default, the scope is page, so the bean is not available in `that.jsp`. If it were any other scope, the answer would have been d.

12. Consider the following code contained in a file named `this.jsp` (the same as above, except the fourth line):

```
<html><body>
<jsp:useBean id="address" class="AddressBean" />
<jsp:setProperty name="address" property="*" />
<%@ include file="that.jsp" %>
</body></html>
```

Which of the following is true about the `AddressBean` instance declared in the above code? (Select one)

- a The bean instance will not be available in `that.jsp`.
- b The bean instance may or may not be available in `that.jsp`, depending on the threading model implemented by `that.jsp`.
- c The bean instance will be available in `that.jsp`, and the `that.jsp` page can print the values of the bean's properties using `<jsp:getProperty />`.
- d The bean instance will be available in `that.jsp`, and the `that.jsp` page can print the values of the bean's properties using `<jsp:getProperty />` only if `that.jsp` also contains a `<jsp:useBean />` declaration identical to the one in `this.jsp` and before using `<jsp:getProperty />`.

Answer: c

Explanation

The `that.jsp` page is included using a directive. Thus, it is a static inclusion, and the two pages form a single translation unit.

CHAPTER 15—USING CUSTOM TAGS

1. Which of the following elements are required for a valid `<taglib>` element in `web.xml`? (Select two)
- a `uri`
 - b `taglib-uri`

- c tagliburi
- d tag-uri
- e location
- f taglib-location
- g tag-location
- h tagliblocation

Answers: b and f

Explanation

The <taglib> element is defined as follows:

```
<!ELEMENT taglib (taglib-uri, taglib-location)>
```

As you can see, both taglib-uri and taglib-location are required elements.

2. Which of the following web.xml snippets correctly defines the use of a tag library? (Select one)

- a <taglib>
 <uri>http://www.abc.com/sample.tld</uri>
 <location>/WEB-INF/sample.tld</location>
 </taglib>
- b <tag-lib>
 <taglib-uri>http://www.abc.com/sample.tld</taglib-uri>
 <taglib-location>/WEB-INF/sample.tld</taglib-location>
 </tag-lib>
- c <taglib>
 <taglib-uri>http://www.abc.com/sample.tld</taglib-uri>
 <taglib-location>/WEB-INF/sample.tld</taglib-location>
 </taglib>
- d <tag-lib>
 <taglib>http://www.abc.com/sample.tld</taglib-uri>
 <taglib>/WEB-INF/sample.tld</taglib-location>
 </tag-lib>

Answer: c

Explanation

The use of a tag library is defined using the <taglib> element:

```
<!ELEMENT taglib (taglib-uri, taglib-location)>
```

3. Which of the following is a valid taglib directive? (Select one)

- a <% taglib uri="/stats" prefix="stats" %>
- b <%@ taglib uri="/stats" prefix="stats" %>
- c <%! taglib uri="/stats" prefix="stats" %>
- d <%@ taglib name="/stats" prefix="stats" %>
- e <%@ taglib name="/stats" value="stats" %>

Answer: b

Explanation

A directive starts with `<%@`, so answers a and c are invalid. A `taglib` directive requires `uri` and `prefix` attributes, so only answer b is correct.

4. Which of the following is a valid `taglib` directive? (Select one)
- a `<%@ taglib prefix="java" uri="sunlib"%>`
 - b `<%@ taglib prefix="jsp" uri="sunlib"%>`
 - c `<%@ taglib prefix="jsp" uri="sunlib"%>`
 - d `<%@ taglib prefix="servlet" uri="sunlib"%>`
 - e `<%@ taglib prefix="sunw" uri="sunlib"%>`
 - f `<%@ taglib prefix="suned" uri="sunlib"%>`

Answer: f

Explanation

The JSP specification does not allow us to use the names `jsp`, `java`, `javax`, `servlet`, `sun`, and `sunw` as a value for the `prefix` attribute. Therefore, only answer f is valid.

5. Consider the following `<taglib>` element, which appears in a deployment descriptor of a web application:

```
<taglib>
  <taglib-uri>/accounting</taglib-uri>
  <taglib-location>/WEB-INF/tlds/SmartAccount.tld</taglib-location>
</taglib>
```

Which of the following correctly specifies the use of the above tag library in a JSP page? (Select one)

- a `<%@ taglib uri="/accounting" prefix="acc"%>`
- b `<%@ taglib uri="/acc" prefix="/accounting"%>`
- c `<%@ taglib name="/accounting" prefix="acc"%>`
- d `<%@ taglib library="/accounting" prefix="acc"%>`
- e `<%@ taglib name="/acc" prefix="/accounting"%>`

Answer: a

Explanation

The `taglib` directive contains two attributes, `uri` and `prefix`:

- `uri`: The value of the `uri` attribute in a JSP page must be the same as the value of the `<taglib-uri>` subelement of the `<taglib>` element in `web.xml`. If the entries in `web.xml` are not used, then the value of the `uri` attribute in a JSP page can directly point to the TLD file using a root-relative URI, as in `uri="/WEB-INF/tlds/SmartAccount.tld"`.
- `prefix`: This can be any string allowed by the XML naming specification. It is similar to an alias and is used in the rest of the page to refer to this tag library.

6. You are given a tag library that has a tag named `printReport`. This tag may accept an attribute, `department`, which cannot take a dynamic value. Which of the following are correct uses of this tag? (Select two)

a `<mylib:printReport/>`
b `<mylib:printReport department="finance"/>`
c `<mylib:printReport attribute="department" value="finance"/>`
d `<mylib:printReport attribute="department" attribute-value="finance"/>`
e `<mylib:printReport>`
 `<jsp:attribute name="department" value="finance" />`
 `</mylib:printReport>`

Answers: a and b

Explanation

Answer a is correct because the `department` attribute is not required. Answer b is syntactically correct, but the rest of the answers are syntactically wrong.

7. You are given a tag library that has a tag named `getMenu`, which requires an attribute, `subject`. This attribute can take a dynamic value. Which of the following are correct uses of this tag? (Select two)

a `<mylib:getMenu />`
b `<mylib:getMenu subject="finance"/>`
c `<% String subject="HR";%>`
 `<mylib:getMenu subject="<%=subject%"/>`
d `<mylib:getMenu>` `<jsp:param subject="finance"/>` `</mylib:getMenu>`
e `<mylib:getMenu>`
 `<jsp:param name="subject" value="finance"/>`
 `</mylib:getMenu>`

Answers: b and c

Explanation

Answer a is wrong because `subject` is a required attribute (as the question states). Answer b is correct because a static value can be specified for an attribute that takes a dynamic value (but the reverse is not true). Answer c is correct because the `subject` attribute takes a dynamic value. Answers d and e do not make sense.

8. Which of the following is a correct way to nest one custom tag inside another? (Select one)

a `<greet:hello>`
 `<greet:world>`
 `</greet:hello>`
 `</greet:world>`
b `<greet:hello>`
 `<greet:world>`
 `</greet:world>`
 `</greet:hello>`

```

c <greet:hello
  <greet:world/>
  />
d <greet:hello>
  </greet:hello>
  <greet:world>
  </greet:world>

```

Answer: b

Explanation

The inner tag should exist completely within the outer tag; therefore, answer a is not valid. We cannot use a tag as an attribute to another tag. Thus, answer c is incorrect. Answer d does not have any kind of nesting at all.

9. Which of the following elements can you use to import a tag library in a JSP document? (Select one)

```

a <jsp:root>
b <jsp:taglib>
c <jsp:directive.taglib>
d <jsp:taglib.directive>
e We cannot use custom tag libraries in XML format.

```

Answer: a

Explanation

In the XML syntax, the tag library information is included in the `<jsp:root>` element:

```

<jsp:root
  xmlns:jsp="http://java.sun.com/JSP/Page"
  xmlns:test="sampleLib.tld"
  version="1.2" >
  ...JSP PAGE...
</jsp:root>

```

10. Using `c` to represent the JSTL library, which of the following produces the same result as `<%= var %>`? (Select one)

```

a <c:set value=var>
b <c:var out=${var}>
c <c:out value=${var}>
d <c:out var="var">
e <c:expr value=var>

```

Answer: c

Explanation

JSTL provides the `out` tag to display a value in the JSP. Therefore, answers a, b, and e are incorrect. If the value corresponds to a variable, then the variable must be enclosed in `${...}`. Answer d is incorrect, and c is the correct choice.

11. Which attribute of `<c:if>` specifies the conditional expression? (Select one)
- a cond
 - b value
 - c check
 - d expr
 - e test

Answer: e

Explanation

The `<c:if>` tag contains only one attribute, and this attribute represents the conditional expression that will be evaluated. This expression must be enclosed with single quotes, and it must be set equal to the `test` attribute. Answers a through d are incorrect, and e is the correct choice.

12. Which of the following JSTL `forEach` tags is valid?
- a `<c:forEach varName="count" begin="1" end="10" step="1">`
 - b `<c:forEach var="count" begin="1" end="10" step="1">`
 - c `<c:forEach test="count" beg="1" end="10" step="1">`
 - d `<c:forEach varName="count" val="1" end="10" inc="1">`
 - e `<c:forEach var="count" start="1" end="10" step="1">`

Answer: b

Explanation

The JSTL `forEach` tag works like the `for` statement in Java. One major difference is the loop variable, which has to be `var`. Therefore, answers a, c, and d are incorrect. When executed, `var` starts at the `begin` value, and proceeds to the `end` value in increments of `step`. Therefore, e is incorrect and b is the correct value.

13. Which tags can be found in a JSTL `choose`? (Select two)
- a case
 - b select
 - c check
 - d when
 - e otherwise

Answers: d and e

Explanation

The JSTL `choose` tag works like the `switch...case` structure in Java. However, each option is represented by `when`, and the default option is given by `otherwise`. Therefore, answers d and e are correct.

CHAPTER 16—DEVELOPING CLASSIC CUSTOM TAG LIBRARIES

1. Which of the following is not a valid subelement of the `<attribute>` element in a TLD? (Select one)
- a `<name>`
 - b `<class>`

- c <required>
- d <type>

Answer: b

Explanation

The valid subelements of the <attribute> element are as shown below:

```
<!ELEMENT attribute (name, required? , rtexprvalue?,
                    type?, description?) >
```

Thus, the correct answer is b. <class> is not a valid subelement of the <attribute> element in a TLD.

2. What is the name of the tag library descriptor element that declares that an attribute can have a request-time expression as its value?

a [_____]

Answer: *rtexprvalue*

Explanation

The <rtexprvalue> element specifies whether or not an attribute can take a request-time expression as its value. Its value can be either true or false. By default, it is false.

3. Consider the following code in a JSP page:

```
<% String message = "Hello "; %>
<test:world>
    How are you?
    <% message = message + "World! " %>
</test:world>
<%= message %>
```

If `doStartTag()` returns `EVAL_BODY_BUFFERED` and `doAfterBody()` clears the buffer by calling `bodyContent.clearBody()`, what will be the output of the above code? (Select one)

- a Hello
- b Hello How are you?
- c Hello How are you? World!
- d Hello World!
- e How are you World!

Answer: d

Explanation

If `doStartTag()` returns `EVAL_BODY_BUFFERED`, then the body is executed and the output is buffered. The text `How are you?` is inserted into the current `JspWriter` buffer and the scriptlet `<% message = message + "World! " %>` assigns the value `Hello World!` to the `message` String. However, the scriptlet

only assigns the new value; it does not print the value using the `out` variable. Therefore, only the text `How are you?` goes into the buffer. The `doAfterBody()` method discards the body contents by calling `bodyContent.clearBody()`. So, the actual output of the JSP page does not contain the text. After the tag is over, the expression `<%=message%>` prints the value of the `String` `message`, which now contains the text `Hello World!`. Thus, the correct answer is `d`, `Hello World!`.

4. Which of the following interfaces are required at a minimum to create a simple custom tag with a body? (Select one)
- a `Tag`
 - b `Tag` and `IterationTag`
 - c `Tag`, `IterationTag`, and `BodyTag`
 - d `TagSupport`
 - e `BodyTagSupport`

Answer: a

Explanation

The `Tag` interface is all you need to create a simple custom tag with a body. You do not have to implement the `BodyTag` interface in order to specify a body for a tag. The `BodyTag` interface is required only when the evaluation of the tag body needs to be buffered. Answers `d` and `e` are classes and not interfaces.

5. At a minimum, which of the following interfaces are required to create an iterative custom tag? (Select one)
- a `Tag`
 - b `Tag` and `IterationTag`
 - c `Tag`, `IterationTag`, and `BodyTag`
 - d `TagSupport`
 - e `BodyTagSupport`

Answer: b

Explanation

To create an iterative tag, we need to implement `IterationTag`. `IterationTag` extends `Tag`. Hence, the correct answer is `b`.

6. Which of the following methods is never called for handler classes that implement only the `Tag` interface? (Select one)
- a `setParent()`
 - b `doStartTag()`
 - c `doAfterbody()`
 - d `doEndTag()`

Answer: c

Explanation

The `doAfterBody()` method is defined in `IterationTag`. The other three methods are defined in the `Tag` interface. Thus, the correct answer is c; `doAfterBody()` is never called on a tag that implements only the `Tag` interface.

7. Which of the following is a valid return value for `doAfterBody()`? (Select one)
- a `EVAL_BODY_INCLUDE`
 - b `SKIP_BODY`
 - c `EVAL_PAGE`
 - d `SKIP_PAGE`

Answer: b

Explanation

The constant in answer a, `EVAL_BODY_INCLUDE`, is valid only for the `doStartTag()` method. The constants in answers c, `EVAL_PAGE`, and d, `SKIP_PAGE`, are valid only for the `doEndTag()` method.

The method `doAfterBody()` can return the following three values:

- `EVAL_BODY_AGAIN`, to reevaluate the body
- `EVAL_BODY_BUFFERED`, to reevaluate the body (only if the tag implements `BodyTag`)
- `SKIP_BODY`, to skip the body

Thus, the correct answer is b, `SKIP_BODY`.

8. Which element would you use in a TLD to indicate the type of body a custom tag expects?
- a [_____]

Answer: `<body-content>`

Explanation

The `<body-content>` element is a subelement of the `<tag>` element in a TLD that indicates the type of body a custom tag expects. The valid values are empty, `JSB`, and `tagdependent`. `JSP` is the default value if `<body-content>` is not specified.

9. If the `doStartTag()` method returns `EVAL_BODY_INCLUDE` one time and the `doAfterBody()` method returns `EVAL_BODY_AGAIN` five times, how many times will the `setBodyContent()` method be called? (Select one)
- a Zero
 - b One
 - c Two
 - d Five
 - e Six

Answer: a

Explanation

If the `doStartTag()` method returns `EVAL_BODY_INCLUDE`, then the evaluation of the body is not buffered. Thus, `setBodyContent()` is never called. The correct answer is a, zero times.

10. If the `doStartTag()` method returns `EVAL_BODY_BUFFERED` one time and the `doAfterBody()` method returns `EVAL_BODY_BUFFERED` five times, how many times will the `setBodyContent()` method be called? Assume that the body of the tag is not empty. (Select one)
- a Zero
 - b One
 - c Two
 - d Five
 - e Six

Answer: b

Explanation

If the `doStartTag()` method returns `EVAL_BODY_BUFFERED`, then the evaluation of the body is buffered. But `setBodyContent()` is called only once—before the first time the body is evaluated. This happens only after the `doStartTag()` method returns `EVAL_BODY_BUFFERED`. It is not called after every call to `doAfterBody()` regardless of the return value of `doAfterBody()`. The correct answer is b, one time.

11. How is the `SKIP_PAGE` constant used? (Select one)
- a `doStartTag()` can return it to skip the evaluation until the end of the current page.
 - b `doAfterBody()` can return it to skip the evaluation until the end of the current page.
 - c `doEndTag()` can return it to skip the evaluation until the end of the current page.
 - d It is passed as a parameter to `doEndTag()` as an indication to skip the evaluation until the end of the current page.

Answer: c

Explanation

The `SKIP_PAGE` constant is defined in the `Tag` interface as a return value for the `doEndTag()` method. It indicates that evaluation of the page from the end of the current tag until the end of the current page must be skipped. The correct answer is c.

12. Which of the following can you use to achieve the same functionality as provided by `findAncestorWithClass()`? (Select one)
- a `getParent()`
 - b `getParentWithClass()`
 - c `getAncestor()`
 - d `getAncestorWithClass()`
 - e `findAncestor()`

Answer: a

Explanation

`findAncestorWithClass(Tag currentTag, Class class)` is a convenient way to get a reference to an outer tag that is closest to the specified `Class` object. This can also be achieved by calling `getParent()` on the current tag to get its immediate parent and then again calling `getParent()` on the returned reference, working our way up the nested hierarchy until we find the tag with the desired `Class` object. Thus, the correct answer is a, `getParent()`. The methods shown in all the other answers are not valid methods of any interface.

13. Consider the following code in a tag handler class that extends `TagSupport`:

```
public int doStartTag()
{
    //1
}
```

Which of the following can you use at `//1` to get an attribute from the application scope? (Select one)

- a `getServletContext().getAttribute("name");`
- b `getApplication().getAttribute("name");`
- c `pageContext.getAttribute("name", PageContext.APPLICATION_SCOPE);`
- d `bodyContent.getApplicationAttribute("name");`

Answer: c

Explanation

The only implicit object made available to tag handler classes by the engine is `PageContext` via the `setPageContext()` method. To access all other implicit objects and user-defined objects in other scopes, we must save the reference to `pageContext` passed in as a parameter to the `setPageContext()` method and use the saved object in other methods. The `TagSupport` class implements `setPageContext()` and maintains this reference in a protected member named `pageContext`. We can use this member in the methods of the classes that extend `TagSupport`. Thus, the correct answer is c, `pageContext.getAttribute("name", PageContext.APPLICATION_SCOPE);`.

14. Which types of objects can be returned by `PageContext.getOut()`? (Select two)

- a An object of type `ServletOutputStream`
- b An object of type `HttpServletOutputStream`
- c An object of type `JspWriter`
- d An object of type `HttpJspWriter`
- e An object of type `BodyContent`

Answers: c and e

Explanation

The return type of `pageContext.getOut()` is `JspWriter`. If the exam asks you to select only one correct option, select `JspWriter`. However, `BodyContent`

extends `JspWriter`, and the return value of the `pageContext.getOut()` is an object of type `BodyContent` if `doStartTag()` returns `EVAL_BODY_BUFFERED`. Thus, if the exam asks you to select two correct options, then select `BodyContent` as well.

15. We can use the directive `<%@ page buffer="8kb" %>` to specify the size of the buffer when returning `EVAL_BODY_BUFFERED` from `doStartTag()`.
- a True
 - b False

Answer: b

Explanation

The page directive attribute `buffer` specifies the size of the buffer to be maintained by the actual output stream that is meant for the whole page. It has nothing to do with the buffer maintained by `BodyContent`, which is meant for a particular tag when the `doStartTag()` method returns `EVAL_BODY_BUFFERED`. The buffer size of a `BodyContent` object is unlimited. Therefore, the correct answer is b, false.

CHAPTER 17—DEVELOPING “SIMPLE” CUSTOM TAG LIBRARIES

1. What method should you use in a `SimpleTag` tag handler to access dynamic variables?
- a `doTag()`
 - b `setDynamicAttribute()`
 - c `getParent()`
 - d `getDynamicAttribute()`

Answer: b

Explanation

Dynamic attributes are one of the important advantages of the new standard. These allow you to add process variables beyond those explicitly described in the TLD. To use them, you need to set the `<dynamic-attributes>` element to `true` in the TLD, and add the method `setDynamicAttribute()` to the Java class.

2. Which object does a `SimpleTag` tag handler use to access implicit variables?
- a `PageContext`
 - b `BodyContent`
 - c `JspContext`
 - d `SimpleTagSupport`

Answer: c

Explanation

The `PageContext` and `BodyContent` classes are only available for classic tags, so answers a and b are incorrect. Simple tag handlers extend `SimpleTagSupport`,

but this doesn't provide a means of accessing implicit variables. Answer d is incorrect. Instead, a tag handler needs a `JspContext` object, and answer c is correct.

3. Consider the following TLD excerpt:

```
<body-content>
  empty
</body-content>
<attribute>
  <name>color</name>
  <rtexprvalue>>true</rtexprvalue>
</attribute>
<dynamic-attributes>
  true
</dynamic-attributes>
```

If the name of the tag is `tagname` and its prefix is `pre`, which of the following JSP statements is valid?

- a `<pre:tagname color="yellow" size=${sizenum} />`
- b `<pre:tagname size="18" color="red"> </pre:tagname>`
- c `<pre:tagname color="${colorname}" size="22" font="verdana"></pre:tagname>`
- d `<pre:tagname color="green" size="30">font="Times New Roman"</pre:tagname>`
- e `<pre:tagname color="${colorname}" size="18"></pre>`

Answer: c

Explanation

Answer e is incorrect because the concluding tag isn't `</pre:tagname>`. Answer d is incorrect because its body content contradicts the `<body-content>` element of the TLD. The `<dynamic-attributes>` element is set to `true`, so JSP statements can incorporate variables beyond those mentioned in the TLD. If there are static and dynamic values in a tag, the static values need to come first, so answer b is incorrect. Also, EL can only be used in static variables, so a is incorrect. Because answer c doesn't make any of these errors, it is the correct choice.

4. If placed inside the body of a custom tag, which of the following statements won't produce "9"? (Select one)

- a `${3 + 3 + 3}`
- b `"9"`
- c `<c:out value="9">`
- d `<%= 27/3 %>`

Answer: d

Explanation

You can never include scripts (declarations, expressions, or scriptlets) within the body content of simple tags. Because answer d contains an expression, it won't produce "9" and is the correct answer.

5. Which of the following methods need to be invoked in a `SimpleTag` to provide iterative processing? (Select one)

- a `setDynamicAttribute()`
- b `getParent()`
- c `getJspBody()`
- d `doTag()`
- e `getJspContext()`

Answer: d

Explanation

In classic tag handlers, you need many different methods for many different capabilities. But with simple tags, you just need `doTag()`. This takes care of body processing and iteration. The correct answer is d.

6. Which of the following values is invalid inside a `SimpleTag`'s `<bodycontent>` subelement? (Select one)

- a `JSP`
- b `scriptless`
- c `tagdependent`
- d `empty`

Answer: a

Explanation

The `<bodycontent>` element in a `SimpleTag` TLD is the same as for the classic model, with one exception. Since you can't incorporate scripts in `SimpleTags`, the `JSP` value is invalid. Therefore, a is the correct answer.

7. Which of the following is a valid return value for the `SimpleTag`'s `doTag()` method? (Select one)

- a `EVAL_BODY_INCLUDE`
- b `SKIP_BODY`
- c `void`
- d `EVAL_PAGE`
- e `SKIP_PAGE`

Answer: c

Explanation

Unlike the methods used in classical tag handlers, the `doTag()` method doesn't return values to control tag processing. Since the rest of the values are only used for classic tags, the correct answer is c.

8. Which tag file directive makes it possible to process dynamic attributes?

- a `taglib`
- b `page`
- c `tag`
- d `attribute`

Answer: c

Explanation

Like the `<attribute>` element in TLDs, the `attribute` directive in tag files can only declare static variables. To add dynamic attributes to a tag file, you need to use the `tag` directive. Then, within this directive, you can add the `dynamic-attributes` attribute. The correct answer is c.

9. Which of the following statements can't be used to access a tag file from a JSP? (Select one)

a `<%@ taglib prefix="pre" uri="www.mysite.com/dir/" %>`
b `<%@ taglib prefix="pre" tagdir="/WEB-INF/tags" %>`
c `<%@ taglib prefix="pre" tagdir="/WEB-INF/tagfiles" %>`
d `<%@ taglib prefix="pre" tagdir="/WEB-INF/tags/myDirectory" %>`

Answer: c

Explanation

When using tag files in a JSP, you can specify their location with the `uri` or `tagdir` attribute. But the files must be located in `/WEB-INF/tags` or a subdirectory. Since `/WEB-INF/tagfiles` is neither, c is the correct answer.

10. Which tag file action processes `JspFragments` in tag attributes?

a `taglib`
b `jsp:invoke`
c `tag`
d `jsp:doBody`
e `attribute`

Answer: b

Explanation

Answers a, c, and e contain tag file directives, and therefore are incorrect. The two new actions used in tag files are `jsp:doBody` and `jsp:invoke`. The first processes body content and the second processes attributes that have been declared as `JspFragments`. Therefore, b is the correct answer.

11. Which `JspFragment` method is used to process body content in a `SimpleTag`? (Select one)

a `invoke()`
b `getOut()`
c `getJspContext()`
d `getBodyContent()`

Answer: a

Explanation

The `getBodyContent()` method processes body content in classic tags, but not in the simple tag model, so answer d is incorrect. The `getJspContext()` method returns a `JspContext` object, and `getJspContext().getOut()` returns a `JspWriter`, but neither processes body content. Therefore, answers b and c are

incorrect. The only method in the simple tag model that processes body content is `invoke()`, which processes the body content obtained through the `getJspBody()` method; thus a is correct.

12. Which class provides an implementation of the `doTag()` method? (Select one)
- a `TagSupport`
 - b `BodyTagSupport`
 - c `SimpleTagSupport`
 - d `IterationTagSupport`
 - e `JspTagSupport`

Answer: c

Explanation

`TagSupport`, `BodyTagSupport`, and `IterationTagSupport` are implementation classes in the classic tag model, and answers a, b, and d are incorrect. At present, there is no `JspTagSupport` class, so e is incorrect. The primary implementation class in the simple tag model is `SimpleTagSupport`, which provides an implementation of `doTag()`. Therefore, c is the correct answer.

13. In what directory shouldn't you place tag files? (Select one)
- a `/META-INF/tags/tagfiles`
 - b `/WEB-INF/`
 - c `/WEB-INF/tags/tagfiles/tagdir/taglocation`
 - d `/META-INF/tags/`

Answer: b

Explanation

When tag files are incorporated into a web application (`/WEB-INF/`) or a web archive (`/META-INF/`), the tag files must be placed in a `tags` folder under these directories or a subfolder beneath. Therefore, answers a, c, and d are incorrect. Because answer b involves placing tag files directly under `/WEB-INF/`, it is correct.

14. Which type of object is returned by `JspContext.getOut()`? (Select one)
- a `ServletOutputStream`
 - b `HttpServletOutputStream`
 - c `JspWriter`
 - d `BodyContent`

Answer: c

Explanation

Body content is returned by the `getJspBody()` method, so answer d is incorrect. Both `ServletOutputStreams` and `HttpServletOutputStreams` are acquired with `getOutputStream()`, so answers a and b are incorrect. The `JspContext.getOut()` method returns a `JspWriter`, so answer c is correct. This object can be used to display data within the JSP.

15. Which of the following methods does the web container call first to initiate a `SimpleTag`'s life cycle?
- a `setJspContext ()`
 - b `setParent ()`
 - c `getJspContext ()`
 - d `getJspBody ()`
 - e `getParent ()`

Answer: c

Explanation

To enable the Java class to access scoped and implicit variables, the web container starts the `SimpleTag` life cycle with `setJspContext ()`. Therefore, a is the correct answer. You can access these variables with `getJspContext ()`. After the context is set, the web container calls `setParent ()` and `setJspBody ()` to continue the life cycle processing.

CHAPTER 18—DESIGN PATTERNS

1. What are the benefits of using the Transfer Object pattern? (Select two)
- a The type of the actual data source can be specified at deployment time.
 - b The data clients are independent of the data source vendor API.
 - c It increases the performance of data-accessing routines.
 - d It allows the clients to access the data source through EJBs.
 - e It allows resource locking in an efficient way.

Answers: a and b

Explanation

This pattern is used to decouple business logic from data access logic. It hides the data access mechanism from the business objects so that the data source can be changed easily and transparently to the business objects.

2. Which design pattern allows you to decouple the business logic, data representation, and data presentation? (Select one)
- a Model-View-Controller
 - b Transfer Object
 - c Bimodal Data Access
 - d Business Delegate

Answer: a

Explanation

In the Model-View-Controller pattern, Model is the data representation, View is the data presentation, and Controller is the implementation of business logic. Therefore, a is the correct answer.

3. Which of the following are the benefits of using the Transfer Object design pattern? (Select two)
- a It improves the response time for data access.
 - b It improves the efficiency of object operations.

- c It reduces the network traffic.
- d It reduces the coupling between the data access module and the database.

Answers: a and c

Explanation

The Transfer Object pattern allows you to retrieve all the data elements in one remote call instead of making multiple remote calls; therefore, it reduces the network traffic and improves the response time since the subsequent calls to the object are local.

4. Which of the following statements are correct? (Select two)
- a The Transfer Object pattern ensures that the data is not stale at the time of use.
 - b It is wise to make the Transfer Object immutable if the Transfer Object represents read-only data.
 - c Applying the Transfer Object pattern on EJBs helps to reduce the load on enterprise beans.
 - d A Transfer Object exists only on the server side.

Answers: b and c

Explanation

Answer a is wrong because just the reverse is true. For instance, this pattern is not used when the attributes of an EJB are volatile, such as stock quotes. Answer b is correct because making the Transfer Object immutable reinforces the idea that the Transfer Object is not a remote object and any changes to its state will not be reflected on the server. Answer c is correct because clients require a fewer number of remote calls to retrieve attributes. Answer d is wrong because a Transfer Object is created on the server and sent to the client.

5. What are the benefits of using the Business Delegate pattern? (Select three)
- a It implements the business service functionality locally to improve performance.
 - b It shields the clients from the details of the access mechanism, such as CORBA or RMI, of the business services.
 - c It shields the clients from changes in the implementation of the business services.
 - d It provides the clients with a uniform interface to the business services.
 - e It reduces the number of remote calls and reduces network overhead.

Answers: b, c, and d

Explanation

Answer a is wrong because a Business Delegate does not implement any business service itself. It calls remote methods on the business services on behalf of the presentation layer. Answer b is correct because the clients delegate the task of calling remote business service methods to the Business Delegate. Thus, they are shielded by the Business Delegate from the access mechanism of the business services. Answer c is correct because the Business Delegate is meant for shielding the clients from the implementation of the business services. Answer d is also correct because this is one of the goals of this pattern. Answer e is not correct because the

Business Delegate does not reduce the number of remote calls. It calls the remote methods on behalf of the client components.

6. You are designing an application that is required to display the data to users through HTML interfaces. It also has to feed the same data to other systems through XML as well as WAP interfaces. Which design pattern would be appropriate in this situation? (Select one)
- a Interface Factory
 - b Session Facade
 - c Transfer Object
 - d Model-View-Controller
 - e Factory

Answer: d

Explanation

The application requires multiple views (HTML, XML, and WAP) for the same data; therefore, MVC is the correct answer.

7. You are automating a computer parts ordering business. For this purpose, your web application requires a controller component that would receive the requests and dispatch them to appropriate JSP pages. It would also coordinate the request processing among the JSP pages, thereby managing the workflow. Finally, the behavior of the controller component is to be loaded at runtime as needed. Which design pattern would be appropriate in this situation? (Select one)
- a Front Controller
 - b Session Facade
 - c Transfer Object
 - d Model-View-Controller
 - e Data Access Object

Answer: a

Explanation

This is a standard situation for the Front Controller pattern. The Front Controller receives all the requests and dispatches them to the appropriate JSP pages. This is not the MVC pattern, because the question only asks about controlling the workflow. You would choose the MVC pattern if it asked about controlling and presenting the data in multiple views.

8. You are building the server side of an application and you are finalizing the interfaces that you will provide to the presentation layer. However, you have not yet finalized the access details of the business services. Which design pattern should you use to mitigate this concern? (Select one)
- a Model-View-Controller
 - b Data Access Object
 - c Business Delegate

- d Facade
- e Transfer Object

Answer: c

Explanation

You already know the services that you have to provide, but the implementation of the service access mechanism for the services has not yet been decided. The Business Delegate pattern gives you the flexibility to implement the access mechanism any way you want. The presentation-tier components—servlets and JSP pages—can use the interface provided by the Business Delegate object to access the business services. Later, when the decision about access mechanism changes, only the Business Delegate object needs to be modified. Other components will remain unaffected.



A P P E N D I X D

Exam Quick Prep

This appendix provides a quick recap of all the important concepts that are covered in the exam objectives. It also notes important points regarding the concepts, which may help you answer the questions correctly during the exam. You should go through this appendix a day before you take the exam.

We have grouped the information according to the exam objectives given by Sun. Therefore, the numbering of the objectives corresponds to the numbering given to the objectives on Sun's web site. The objectives are listed with the chapters in which they are discussed. However, since the first three chapters of this book do not correspond to any exam objectives, the objectives start with chapter 4.

CHAPTER 4—THE SERVLET MODEL

Objectives 1.1–1.4, 3.5

1.1 For each of the HTTP Methods (such as GET, POST, HEAD, and so on) describe the purpose of the method and the technical characteristics of the HTTP Method protocol, list triggers that might cause a Client (usually a Web browser) to use the method; and identify the HttpServlet method that corresponds to the HTTP Method.

Important concepts	Exam tips
✧ For HTTP method XXX, HttpServlet's doXXX(HttpServletRequest, HttpServletResponse) is called.	Servlet container calls the service(ServletRequest, ServletResponse) method of the Servlet interface.
✧ For example, for GET, HttpServlet's doGet() is called.	HttpServlet.service(HttpServletRequest, HttpServletResponse) interprets the HTTP request and calls the appropriate doXXX() method.
✧ Triggers for GET request: <ul style="list-style-type: none">• Clicking on a hyperlink• Browsing through the browser's address field	The default method of the HTML FORM element is GET.
✧ Triggers for POST request: <ul style="list-style-type: none">• Submitting an HTML FORM, <i>only</i> if its method attribute is POST	Parameters sent via GET are visible in the URL. Parameters sent via POST are not visible in the URL, and so it is used to send data such as the user ID/password.
✧ Triggers for HEAD request: <ul style="list-style-type: none">• Clicking a menu option that makes the browser synchronize offline content with the web site	GET supports only text data. POST supports text as well as binary data. The HTTP protocol does not limit the length of the query string, but many browsers and HTTP servers limit it to 255 characters. POST is used to send large amounts of data.

1.2 Using the HttpServletRequest interface, write code to

- Retrieve HTML form parameters from the request
- Retrieve HTTP request header information, or
- Retrieve cookies from the request

Important concepts	Exam tips
✧ HTML FORM parameters or parameters embedded in query string can be retrieved using: <ul style="list-style-type: none">• String ServletRequest.getParameter(String paramName)• String[] ServletRequest.getParameterValues(String param)	ServletConfig allows you to get Init parameters. You cannot set anything into it. sendRedirect is not transparent to the browser.

continued on next page

Important concepts

- ◇ HTTP request headers can be retrieved using:
 - String `HttpServletRequest.getHeader(String headerName)`
 - Enumeration `HttpServletRequest.getHeaderValues(String headerName)`
 - ◇ Cookies can be retrieved from a request using:
 - `Cookie[] HttpServletRequest.getCookies()`
-

1.3 *Using the `HttpServletResponse` interface, write code to*

- *Set an HTTP response header*
 - *Set the content type of the response*
 - *Acquire a text stream for the response*
 - *Acquire a binary stream for the response*
 - *Redirect an HTTP request to another URL, or*
 - *Add cookies to the response*
-

Important concepts

- ◇ HTTP response headers can be set using:
`HttpServletResponse.setHeader(String headerName, String value)`
`HttpServletResponse.addHeader(String headerName, String value)`
 - ◇ Content-Type for HTTP response can be set using:
`ServletResponse.setContentType(String value)`
 - ◇ To acquire a text stream to send character data to the response, use:
`PrintWriter ServletResponse.getWriter()`
 - ◇ To acquire a binary stream to send any data to the response, use:
`ServletOutputStream ServletResponse.getOutputStream()`
 - ◇ To redirect an HTTP request to another URL, use:
`HttpServletResponse.sendRedirect(String newURL)`
 - ◇ Cookies can be added to the response using:
`void HttpServletResponse.addCookie(Cookie cookie)`
-

1.4 Describe the purpose and event sequence of the servlet life cycle:

- *Servlet class loading*
- *Servlet instantiation*
- *Call the init method*
- *Call the service method*
- *Call the destroy method*

Important concepts	Exam tips
◇ To start, the web container looks for the deployment descriptor (web.xml) to determine which servlet class is needed. Then, it loads the servlet class from the local or remote filesystem.	
◇ After loading the class, the web container creates an object based on the class. This is the servlet instantiation.	
◇ <code>init(ServletConfig)</code> : Guaranteed to be called once and only once on a Servlet object by the servlet container before putting the servlet into service.	The <code>init()</code> method is called once per servlet object. You can have multiple instantiations of the same servlet class.
◇ <code>service(ServletRequest, ServletResponse)</code> : Called by the servlet container for each request	
◇ <code>destroy()</code> : Called by the servlet container after it takes the servlet out of service. It is called only once. But it may not be called if the servlet container crashes.	

3.5 Describe the RequestDispatcher mechanism;

- *Write servlet code to create a request dispatcher;*
- *Write servlet code to forward or include the target resource; and*
- *Identify and describe the additional request-scoped attributes provided by the container to the target resource.*

Important concepts	Exam tips
◇ The RequestDispatcher interface provides the following methods to forward the request to another servlet or to include the partial response generated by another servlet: <ul style="list-style-type: none">• <code>RequestDispatcher.forward(ServletRequest, ServletResponse)</code>• <code>RequestDispatcher.include(ServletRequest, ServletResponse)</code>	<code>RequestDispatcher.forward()</code> is transparent to the browser, unlike <code>HttpServletResponse.sendRedirect()</code> . The path string passed to <code>ServletContext.getRequestDispatcher()</code> must start with <code>/</code> . <code>ServletRequest.getRequestDispatcher</code> supports relative paths while <code>ServletContext.getRequestDispatcher</code> does not.
◇ <code>ServletRequest</code> and <code>ServletContext</code> provide the following method to get a Request Dispatcher: <ul style="list-style-type: none">• <code>getRequestDispatcher(String path)</code>	
◇ Additionally, <code>ServletContext</code> provides the following method: <ul style="list-style-type: none">• <code>getNamedDispatcher(String servletName)</code>	

continued on next page

Important concepts	Exam tips
<p>◇ An included servlet can obtain information about the request with the following attributes:</p> <ul style="list-style-type: none">• javax.servlet.include.request_uri• javax.servlet.include.context_path• javax.servlet.include.servlet_path• javax.servlet.include.path_info• javax.servlet.include.query_string	<p>These attributes cannot be set if the RequestDispatcher was acquired with <code>getNamedDispatcher(String servletName)</code>.</p>
<p>◇ A forwarded servlet can obtain information about the request with the following attributes:</p> <ul style="list-style-type: none">• javax.servlet.forward.request_uri• javax.servlet.forward.context_path• javax.servlet.forward.servlet_path• javax.servlet.forward.path_info• javax.servlet.forward.query_string	<p>These attributes cannot be set if the RequestDispatcher was acquired with <code>getNamedDispatcher(String servletName)</code>.</p>

CHAPTER 5—STRUCTURE AND DEPLOYMENT

Objectives 2.1–2.4

2.1 *Construct the file and directory structure of a Web Application that may contain*

- *Static content*
- *JSP pages*
- *Servlet classes*
- *The deployment descriptor*
- *Tag libraries*
- *JAR files, and*
- *Java class files; and*
- *Describe how to protect resource files from HTTP access*

Important concepts	Exam tips
<p>✧ Directory structure of a web application:</p> <ul style="list-style-type: none">webappname<ul style="list-style-type: none"> all html, JSP, static files WEB-INF<ul style="list-style-type: none"> web.xml (deployment descriptor) classes<ul style="list-style-type: none"> servlet classes Java classes lib<ul style="list-style-type: none"> JAR files tag libraries <p>You can protect resource files from HTTP access by placing them in the WEB-INF folder. No file in this directory can be directly accessed by a client.</p>	<p>The name of the deployment descriptor file is web.xml, and it should be in the WEB-INF directory.</p> <p>Class files should be in the WEB-INF/classes directory.</p> <p>JAR files should be in the WEB-INF/lib directory.</p>

2.2 *Describe the purpose and semantics of the deployment descriptor.*

Important concepts	Exam tips
<p>✧ The deployment descriptor (web.xml) provides information to the web container concerning the application's characteristics—its servlets, JSPs, static content, etc. It also specifies where each file can be found.</p>	<p>The web.xml file is written using XML tags, and each pair of tags encloses a specific aspect of the application.</p>

2.3 Construct the correct structure of the deployment descriptor.

Important concepts	Exam tips
<p>✦ A servlet container creates a servlet instance for each <code><servlet></code> element.</p> <p>The <code><servlet-name></code> element is used to give a name to a servlet.</p> <p>The <code><servlet-class></code> element specifies the fully qualified Java class name for the servlet.</p> <p>Each <code><init-param></code> element specifies an initialization parameter.</p> <p>Each <code><servlet-mapping></code> element specifies a URI to the servlet mapping.</p>	<p>You can define two servlets using the same servlet class but different servlet names to provide multiple sets of initialization parameters.</p>
<p>✦ Sample Servlet definition:</p> <pre data-bbox="337 577 736 889"><servlet> <servlet-name> TestServlet </servlet-name> <servlet-class> com.abc.TestServlet </servlet-class> <init-param> <param-name>region</param-name> <param-value>USA</param-value> </init-param> </servlet></pre>	
<p>✦ Sample URL-to-Servlet mapping:</p> <pre data-bbox="337 953 663 1106"><servlet-mapping> <servlet-name> ColorServlet </servlet-name> <url-pattern>*.col</url-pattern> </servlet-mapping></pre>	

2.4 *Explain the purpose of a WAR file, describe the contents of a WAR file, and describe how one may be constructed.*

Important concepts	Exam tips
<p>✧ A WAR (Web ARchive) file is an archived web application. At the time of deployment, the name of the web application is assumed to be the name of the WAR file.</p> <p>The WAR directory structure is given by:</p> <pre>webappname.war -all html, JSP, static files -META-INF -MANIFEST.MF -WEB-INF -web.xml (deployment descriptor) -classes -servlet classes -Java classes -lib -JAR files -tag libraries</pre>	<p>Archiving is performed with regular Java archive (jar) tools. From the top-level application directory, you can create the archive with:</p> <pre>jar cvf webappname.war</pre> <p>The main difference between a WAR directory and a regular web application is the presence of META-INF, which holds information for the archive utility.</p>

CHAPTER 6—THE SERVLET CONTAINER MODEL

Objectives 3.1–3.2, 3.4

3.1 For the `ServletContext` initialization parameters

- Write servlet code to access initialization parameters; and
- Create the deployment descriptor elements for declaring initialization parameters.

Important concepts	Exam tips
<p>✧ <code>ServletContext</code> init parameters can be retrieved using <code>ServletContext.getInitParameter(String)</code>.</p> <p>The following is a <code>web.xml</code> code snippet showing <code>ServletContext</code> init parameters and a listener configuration:</p> <pre><web-app> ... <context-param> <param-name>locale</param-name> <param-value>US</param-value> </context-param> ... </web-app></pre>	<p><code>javax.servlet.GenericServlet</code> implements the <code>ServletConfig</code> interface.</p>

3.2 For the fundamental servlet attribute scopes (request, session, and context):

- Write servlet code to add, retrieve, and remove attributes;
- Given a usage scenario, identify the proper scope for an attribute; and
- Identify multi-threading issues associated with each scope.

Important concepts	Exam tips
<p>✧ For the request, session, and context scopes, attributes are manipulated with the same four methods:</p> <ul style="list-style-type: none">• <code>Object.getAttribute(String name)</code>• <code>Enumeration.getAttributeNames()</code>• <code>void.setAttribute(String name, Object value)</code>• <code>void.removeAttribute(String name)</code>	<p>An attribute name can have only one value.</p>
<p>✧ There is one <code>ServletContext</code> for each web application on each JVM. However, <code>ServletContext</code> for the default web application exists on only one JVM.</p> <p><code>HttpSession</code> exists on only one JVM at a time but may be migrated across the JVMs.</p> <p><code>HttpSession</code> attributes should implement the <code>java.io.Serializable</code> interface; otherwise, <code>setAttribute()</code> may throw an <code>IllegalArgumentException</code>.</p>	<p><code>ServletContext</code> init parameters are available on all the JVMs.</p> <p><code>ServletContext</code> attributes are not visible across the JVMs.</p> <p><code>ServletContextEvent</code>, <code>ServletContextAttributeEvent</code>, and <code>HttpSessionAttributeEvent</code> may not be propagated across the JVMs. Therefore, a distributable web application should not depend on the notification of changes to the attribute list of either <code>ServletContext</code> or <code>HttpSession</code>.</p>

3.4 Describe the Web container life cycle event model for requests, sessions, and web applications;

- Create and configure listener classes for each scope life cycle;
- Create and configure scope attribute listener classes; and
- Given a usage scenario, identify the proper attribute listener to use.

Listener interface	Methods
javax.servlet.ServletContextListener <ul style="list-style-type: none"> • Configured in the deployment descriptor. • Used to listen for creation and destruction of the servlet context. • Should not depend on this in a distributed environment. 	void contextDestroyed(ServletContextEvent sce) Called when the servlet context is about to be destroyed. void contextInitialized(ServletContextEvent sce) Called when the web application is ready to process requests.
javax.servlet.ServletContextAttributeListener <ul style="list-style-type: none"> • Configured in the deployment descriptor. • Used to listen for changes in the attribute list of the servlet context. • Should not depend on this in a distributed environment. 	void attributeAdded(ServletContextAttributeEvent scae) Called when a new attribute is added to the servlet context. void attributeRemoved(ServletContextAttributeEvent scae) Called when an attribute is removed from the servlet context. void attributeReplaced(ServletContextAttributeEvent scae) Called when an attribute on the servlet context is replaced.
javax.servlet.http.HttpSessionListener <ul style="list-style-type: none"> • Configured in the deployment descriptor. • Used to listen for changes in the list of sessions of a web application. • Should not depend on this in a distributed environment. For example, if a session is created on one machine, this may not trigger a call to sessionCreated() on another machine. 	void sessionCreated(HttpSessionEvent se) Called when a session is created. void sessionDestroyed(HttpSessionEvent se) Called when a session is invalidated.
javax.servlet.http.HttpSessionAttributeListener <ul style="list-style-type: none"> • Configured in the deployment descriptor. • Used to listen for changes in the attribute list of sessions of a web application. • Should not depend on this in a distributed environment because the HttpSessionAttributeListener instances will be present on all the JVMs but the session will be present on only one JVM. Thus, the HttpSessionBindingEvents will be delivered only to the listener instance that is present on the session's JVM. 	void attributeAdded(HttpSessionBindingEvent se) Called when an attribute is added to a session. void attributeRemoved(HttpSessionBindingEvent se) Called when an attribute is removed from a session. void attributeReplaced(HttpSessionBindingEvent se) Called when an attribute is replaced in a session.

continued on next page

Listener interface	Methods
javax.servlet.http.HttpSessionBindingListener <ul style="list-style-type: none"> • <i>Not</i> configured in the deployment descriptor. • An attribute should implement this interface if it wants to be notified when it is added or removed from a session. • Can depend on this in a distributed environment, since a session resides on only one machine at a time. 	void valueBound(HttpSessionBindingEvent event) Called on the object when it is being bound to a session. void valueUnbound(HttpSessionBindingEvent event) Called on the object when it is being unbound from a session.
javax.servlet.http.HttpSessionActivationListener <ul style="list-style-type: none"> • <i>Not</i> configured in the deployment descriptor. • An attribute should implement this interface if it wants to be notified when the session is migrated. • Can depend on this in a distributed environment, because a session resides on only one machine at a time. 	void sessionDidActivate(HttpSessionEvent se) Called on all the attributes that implement this interface after the session is activated. void sessionWillPassivate(HttpSessionEvent se) Called on all the attributes that implement this interface just before the session is passivated.
javax.servlet.http.HttpServletRequestAttributeListener <ul style="list-style-type: none"> • Must be configured in the deployment descriptor. • An attribute should implement this interface if it wants to be notified when the request's attribute's change. 	

CHAPTER 7—USING FILTERS

Objective 3.3

3.3 Describe the Web container request processing model;

- Write and configure a filter;
- Create a request or response wrapper; and
- Given a design problem, describe how to apply a filter or a wrapper.

Important concepts

- ◇ All filters implement the `javax.servlet.Filter` interface.
`javax.servlet.FilterConfig` provides initialization parameters to a filter through:
 - `getInitParameter(String name)`
 - `getInitParameterNames()`
- ◇ `FilterConfig` also contains a reference to the `ServletContext` object, which can be retrieved through the `getServletContext()` method.
Filter life-cycle methods:
 - `init(FilterConfig)`: Called by the container during application startup.
 - `doFilter(ServletRequest, ServletResponse)`: Called by the container for each request whose URL is mapped to this filter.
 - `destroy()`: Called by the container during application shutdown.
- ◇ The following is a sample filter declaration:

```
<filter>
  <filter-name>ValidatorFilter</filter-name>
  <description>Validates the requests
  </description>
  <filter-class>
    com.abc.filters.ValidatorFilter
  </filter-class>
  <init-param>
    <param-name>locale</param-name>
    <param-value>USA</param-value>
  </init-param>
</filter>
```
- ◇ The following is a sample filter mapping that associates `ValidatorFilter` with a `*.doc` URL pattern:

```
<filter-mapping>
  <filter-name>ValidatorFilter</filter-name>
  <url-pattern>*.doc</url-pattern>
</filter-mapping>
```
- ◇ The following is a sample filter mapping that associates `ValidatorFilter` with `TestServlet`:

```
<filter-mapping>
  <filter-name>ValidatorFilter</filter-name>
  <servlet-name>TestServlet</servlet-name>
</filter-mapping>
```

continued on next page

Important concepts

- ✧ The `javax.servlet` package defines the `ServletRequestWrapper` and `ServletResponseWrapper` classes.
 - ✧ The `javax.servlet.http` package defines the `HttpServletRequestWrapper` and `HttpServletResponseWrapper` classes.
 - ✧ All four classes delegate the method calls to the underlying request or response object.
-

CHAPTER 8—SESSION MANAGEMENT

Objectives 4.1–4.4

4.1 Write servlet code to store objects into a session object and retrieve objects from a session object.

Important concepts

- ✧ Methods to get and store attributes from and to HttpSession:
 - Object `getAttribute(String name)`;
 - void `setAttribute(String name, Object value)`
-

4.2 Given a scenario

- Describe the APIs used to access the session object,
- Explain when the session object was created, and
- Describe the mechanisms used to destroy the session object, and when it was destroyed.

Important concepts

Exam tips

- | | |
|--|---|
| <ul style="list-style-type: none">✧ Methods to access a session while processing a user request:<ul style="list-style-type: none">• <code>request.getSession(boolean createFlag)</code>• <code>request.getSession()</code>: This is the same as <code>request.getSession(true)</code>✧ Three ways to monitor session creation and usage:<ul style="list-style-type: none">• Cookies• URL rewriting• SSL information✧ Three ways to terminate a session:<ul style="list-style-type: none">• <code>HttpSession.invalidate()</code>• Set a time-out period using <code>HttpSession.setMaxInactiveInterval(int)</code>• Set a time-out period in <code>web.xml</code>:<pre><web-app> ... <session-config> <session-timeout>30</session-timeout> </session-config> ... </web-app></pre>✧ A session is invalidated if no request comes from the client for <code>session-timeout</code> minutes or if <code>session.invalidate()</code> is called. | <p>If the argument is set to true, an HttpSession object will be created if it does not already exist.</p> <p>With cookies, the server sends the session ID to the client. Each user request also contains this cookie, so the server can track the session.</p> <p>Session timeout is specified in minutes. A value of 0 or less means the sessions will never expire (unless explicitly expunged using <code>session.invalidate()</code>).</p> <p>Session timeout for a particular session can be changed using <code>HttpSession.setMaxInactiveInterval(int seconds)</code>. This method takes the interval in number of seconds (unlike the deployment descriptor, which takes minutes). It does not affect other sessions. Any negative value prevents the session from being timed out.</p> |
|--|---|
-

4.3 Using session listeners,

- Write code to respond to an event when an object is added to a session, and
- Write code to respond to an event when a session object migrates from one VM to another.

Important concepts	Exam tips
<ul style="list-style-type: none">◇ HttpSessionBindingListener is used by objects to receive notifications when they are added to or removed from a session. It has two methods:<ul style="list-style-type: none">• void valueBound(HttpSessionBindingEvent e)• void valueUnbound(HttpSessionBindingEvent e)◇ HttpSessionActivationListener is used to receive notifications when any session migrates from one VM to another:<ul style="list-style-type: none">• void sessionWillPassivate(HttpSessionEvent e)• void sessionDidActivate(HttpSessionEvent e)	<p>HttpSessionBindingListener is not configured in the deployment descriptor and is implemented by classes, objects of which are to be stored in a session.</p> <p>HttpSessionBindingEvent extends HttpSessionEvent and adds methods: String getName() Object getValue()</p> <p>HttpSessionEvent has only one method: HttpSession getSession()</p>

4.4 Given a scenario,

- Describe which session management mechanism the Web container could employ,
- How cookies might be used to manage sessions,
- How URL rewriting might be used to manage sessions, and
- Write servlet code to perform URL rewriting.

Important concepts	Exam tips
<ul style="list-style-type: none">◇ With cookies, the server sends the session ID with each response. Each user request also contains this ID, so the server can track the session.◇ To use URL rewriting, all the URLs displayed by the application must have the session ID attached.◇ HttpServletResponse.encodeURL(String url) appends the session ID to the URL only if it is required.◇ HttpServletResponse.encodeRedirectURL(String url) should be used to rewrite URLs that are to be passed to the response.sendRedirect() method.	<p>The name of the cookie must be JSESSIONID (uppercase).</p> <p>None of the static HTML pages that contain URLs can be served directly to the client. They must be parsed in a servlet and the session ID must be attached to the URLs.</p>

CHAPTER 9—DEVELOPING SECURE WEB APPLICATIONS

Objectives 5.1–5.3

5.1 *Based on the servlet specification, compare and contrast the following security mechanisms:*

- *Authentication*
- *Authorization*
- *Data integrity*
- *Confidentiality*

Important concepts

- ◇ Authentication: Verifying that the user is who he claims to be. Performed by asking for the username and password.
 - ◇ Authorization: Verifying that the user has the right to access the requested resource. Performed by checking with an access control list.
 - ◇ Data integrity: Verifying that the data is not modified in transit. Performed using cryptography.
 - ◇ Confidentiality: Making sure that unintended parties cannot make use of the data. Performed using encryption.
-

5.2 *In the deployment descriptor, declare*

- *A security constraint,*
- *A Web resource,*
- *The transport guarantee,*
- *The login confirmation, and*
- *A security role.*

Important concepts	Exam tips
◇ Three things are used to define a security constraint: <ol style="list-style-type: none">1. web-resource-collection (at least one is required)2. auth-constraint (optional)3. user-data-constraint (optional)	Values for transport-guarantee: NONE implies HTTP CONFIDENTIAL, INTEGRAL imply HTTPS Values for auth-method: BASIC, FORM, DIGEST, and CLIENT-CERT.

continued on next page

Important concepts

- ✧ Sample security constraint:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>declarativetest
    </web-resource-name>
    <url-pattern>/servlet/SecureServlet
    </url-pattern>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>supervisor</role-name>
  </auth-constraint>
  <user-data-constraint>
    <transport-guarantee>NONE
    </transport-guarantee>
  </user-data-constraint>
</security-constraint>
```

- ✧ Login config is used to authenticate the users. Sample login configuration:

```
<login-config>
  <auth-method>FORM</auth-method>
  <form-login-config>
    <form-login-page>/formlogin.html
    </form-login-page>
    <form-error-page>/formerror.html
    </form-error-page>
  </form-login-config>
</login-config>
```

- ✧ Sample security role:

```
<security-role>
  <role-name>supervisor</role-name>
</security-role>
```

- ✧ Programmatic security requires role names that are hard-coded in the servlet to be specified in the security-role-ref element. An example:

```
<servlet>
  <servlet-name>SecureServlet</servlet-name>
  <servlet-class>cgscwcd.chapter9.SecureServlet
  </servlet-class>
  <security-role-ref>
    <role-name>manager</role-name>
    <role-link>supervisor</role-link>
  </security-role-ref>
</servlet>
```

In this example, manager will be hard-coded in the servlet while supervisor is the actual role name in the deployment environment.

Exam tips

The <web-resource-collection> tags refer to the resources that are protected by the security constraints. Each collection is named by the <web-resource-name> tags.

The web resource is represented by a <url-pattern> that can be accessed with a specific <http-method>. These methods correspond to the HTTP methods discussed earlier: GET, POST, etc.

The <transport-guarantee> tags specify the protection of the communication layer
NONE implies no protection (HTTP)
INTEGRAL implies error-free data (HTTPS)
CONFIDENTIAL implies protection against eavesdropping (HTTPS)

Values for auth-method: BASIC, FORM, DIGEST, and CLIENT-CERT.

5.3 *Compare and contrast the authentication types (BASIC, DIGEST, FORM, and CLIENT-CERT); describe how the type works; and given a scenario, select an appropriate type.*

Important concepts

- ✧ BASIC: Performed by sending the username and password in Base64 encoding.
Advantages:
 - Very easy to set up
 - Supported by all browsersDisadvantages:
 - It is not secure, since the username and password are not encrypted.
 - You cannot customize the look and feel of the dialog box.

 - ✧ DIGEST: Performed by sending a digest of the password in an encrypted form.
Advantages:
 - SecureDisadvantages:
 - Not supported by all browsers

 - ✧ FORM: Performed by sending username and password in Base64 encoding. The username and password are captured using a customized HTML FORM.
Advantages:
 - Easy to set up
 - Supported by all browsers
 - Customized look and feelDisadvantages:
 - It is not secure, since the username and password are not encrypted unless HTTPS is used.

 - ✧ CLIENT-CERT:
Advantages:
 - Very secure
 - Supported by all browsersDisadvantages:
 - Costly to implement
-

CHAPTER 10—THE JSP TECHNOLOGY MODEL—THE BASICS

Objectives 6.1, 6.2, 6.4

6.1 Identify, describe, or write the JSP code for the following elements:

- *Template text*
- *Scripting elements (comments, directives, declarations, scriptlets, and expressions)*
- *Standard and custom actions, and*
- *Expression language elements.*

Important concepts	Exam tips
<p>✦ Template text contains strings that remain uninterpreted and are sent directly to the JspWriter for display.</p>	Tag names, their attributes, and their values are case sensitive.
<p>✦ Comments block out sections of text in the page. They are coded with two dashes: -- <%-- This won't be processed --%></p>	Unknown attributes or invalid values result in errors and are caught during the translation phase.
<p>✦ Directives specify translation time instructions to the JSP engine. They are coded with @ <%@ page attribute list %> <%@ include file="relativeURL" %> <%@ taglib prefix="" uri="" %></p>	Page directives can be placed anywhere in a page but apply to the entire translation unit.
<p>✦ Declarations declare and define methods and variables. They are coded with ! <%! int count; %> <%! int getCount() { return count; } >%></p>	Variable declarations end with a semicolon; methods do not.
<p>✦ Scriptlets are used for writing free-form Java code. They have no special coding character. <% //some Java code >%></p>	Expressions must not end with a semicolon. <%! String name="SCWCD"; %> is an instance variable declared outside _jspService(). <% String name="SCWCD"; %> is a local variable declared inside _jspService().
<p>✦ Expressions serve as shortcuts to print values in the generated page. They are coded with = <%= request.getParameter("paramName") %></p> <p>Standard actions are covered in Chapter 11 and Expression Language elements are covered in Chapter 13.</p>	

6.2 Write JSP code that uses the directives:

- *'page'* (with attributes *'import'*, *'session'*, *'contentType'*, and *'isELIgnored'*)
- *'include'*, and
- *'taglib'*.

Important concepts	Exam tips
<p>✧ The page directive specifies parameters that apply to the entire JSP. It has four attributes:</p> <ul style="list-style-type: none">• <i>'import'</i> identifies a package that should be imported to the JSP-translated servlet <code><%@ page import="java.io.*" %></code>• <i>'session'</i> tells the web container that the JSP is part of a session and will need to access the HttpSession object <code><%@ page session="true" %></code>• <i>'contentType'</i> identifies the MIME type of the JSP-translated servlet <code><%@ page contentType="MIME type; charset=value" %></code>• <i>'isELIgnored'</i> specifies whether the JSP will include Expression Language elements <code><%@ page isELIgnored="false" %></code>	<p>By default, the web container assumes that the <i>'session'</i> attribute is set to true.</p> <p>This is similar to the <code>setContentType()</code> method of the response</p> <p>In JSPs following the 1.2 standard, the default is <i>'true.'</i> In JSPs following the 2.0 standard, the default is <i>'false.'</i></p>
<p>✧ The include directive adds files to the servlet processing during translation. <code><%@ include file="filename.suf" %></code></p>	
<p>✧ The taglib directive identifies a tag library (*.tld) containing tags that will be used in the JSP. The library-s location is specified by the uri attribute, and the tag prefix is identified by the prefix attribute. <code><%@ taglib prefix="pre" uri="/WEB-INF/lib.tld" %></code></p>	

6.4 Describe the purpose and event sequence of the JSP page life cycle:

- JSP page translation
- JSP page compilation
- Load class
- Create instance
- Call the `jspInit` method
- Call the `_jspService` method
- Call the `_jspDestroy` method

Important concepts

- ◇ The life-cycle phases occur in the following order:
1. Page translation
 2. Page compilation
 3. Load class
 4. Create instance
 5. Call `jspInit()`
 6. Call `_jspService()`
 7. Call `jspDestroy()`

Exam tips

`jspInit()` and `jspDestroy()` are defined in the `javax.servlet.jsp.JspPage` interface.

`_jspService()` is defined in the `javax.servlet.jsp.HttpJspPage` interface.

`jspInit()` and `jspDestroy()` are called only once.

`_jspService()` is called multiple times, once for every request for this page.

JSP declarations are used to declare `jspInit()` and `jspDestroy()`.

We never declare `_jspService()` explicitly. The JSP engine automatically declares it.

The return type of all the three methods is `void`.

`jspInit()` and `jspDestroy()` take no arguments.

`jspInit()` and `jspDestroy()` do not throw any exceptions.

`_jspService()` takes two arguments:

- `HttpServletRequest`
- `HttpServletResponse`

`_jspService()` throws two exceptions:

- `ServletException`
- `IOException`

CHAPTER 11—THE JSP TECHNOLOGY MODEL—ADVANCED TOPICS

Objectives 6.3, 6.5

6.3 Write a JSP document (XML-based document) that uses the correct syntax.

Important concepts	Exam tips
<p>✧ XML syntax for directives:</p> <pre><jsp:directive.page attribute list /> <jsp:directive.include file="relativeURL" /></pre>	<p>The rules for tag names, attributes, and values are the same in XML and JSP formats.</p>
<p>✧ XML syntax for declarations:</p> <pre><jsp:declaration> int count; </jsp:declaration> <jsp:declaration> int getCount(){ return count; } </jsp:declaration></pre>	<p>The semantics of placement of the tags are the same in XML and JSP formats.</p> <p>All XML-based pages should have a root tag named <jsp:root>. Thus, the page ends with </jsp:root>.</p> <p>There is no taglib directive in XML. Taglibs are specified in <jsp:root>.</p>
<p>✧ XML syntax for scriptlets:</p> <pre><jsp:scriptlet> //some Java code </jsp:scriptlet></pre>	
<p>✧ XML syntax for expressions:</p> <pre><jsp:expression> request.getParameter("paramName") </jsp:expression></pre>	

6.5 Given a design goal, write JSP code using the appropriate implicit objects: request, response, out, session, config, application, page, pageContext, and exception.

Implicit object	Purpose, Function, or Uses
request	Object of type javax.servlet.http.HttpServletRequest. Passed in as a parameter to _jspService(). Used for getting HTTP header information, cookies, parameters, etc. Also used for getting and setting attributes into the request scope.
response	Object of type javax.servlet.http.HttpServletResponse. Passed in as a parameter to _jspService(). Used to send a response to the client. Used for setting HTTP header information, cookies, etc.
out	Object of type javax.servlet.jsp.JspWriter; used for writing data to the output stream.
session	Object of type javax.servlet.http.HttpSession. Used for storing and retrieving session related information and sharing objects across multiple requests and pages within the same HTTP session.
config	Object of type javax.servlet.ServletConfig. Used to retrieve initialization parameters for a JSP page.

continued on next page

Implicit object	Purpose, Function, or Uses
application	Object of type <code>javax.servlet.ServletContext</code> . Used for storing and retrieving application related information and sharing objects across multiple sessions, requests, and pages within the same web application.
page	Refers to the generated Servlet class. Not used much because it is declared as of type <code>java.lang.Object</code> .
pageContext	Object of type <code>javax.servlet.jsp.PageContext</code> . Used for storing and retrieving page-related information and sharing objects within the same translation unit and same request. Also used as a convenience class that maintains a table of all the other implicit objects.
exception	Object of type <code>java.lang.Throwable</code> . Only available in pages that have the page directive <code>isErrorPage</code> set to true.

CHAPTER 12—REUSABLE WEB COMPONENTS

Objective 6.7, 8.2

- 6.7 Given a specific design goal for including a JSP segment in another page, write the JSP code that uses the most appropriate inclusion mechanism (the `include` directive or the `jsp:include` standard action).
- 8.2 Given a design goal, create a code snippet using the following standard actions:
- `jsp:include`,
 - `jsp:forward`, and
 - `jsp:param`

Important concepts	Exam tips
<p>✧ To include a component statically, use the <code>include</code> directive:</p> <pre><%@ include file="relativeURL" %></pre>	<p>Points to remember for the <code>include</code> directive:</p> <ul style="list-style-type: none">• There is only one attribute: <i>file</i>.• The <i>file</i> attribute is mandatory.• The <i>file</i> attribute's value is a relative path. We cannot specify a protocol, hostname, or port number.• The attribute cannot point to a servlet.
<p>✧ To include a component dynamically, use the <code>include</code> action:</p> <pre><jsp:include page="relativeURL" flush="true" /></pre>	<p>Points to remember for <code><jsp:include></code>:</p> <ul style="list-style-type: none">• There are two attributes: <i>page</i> and <i>flush</i>.• The <i>page</i> attribute is mandatory.• The <i>flush</i> attribute is optional. The default value of <i>flush</i> is <i>false</i>.• The name of the attribute is <i>page</i>, not <i>url</i> or <i>file</i>.• The value of the <i>page</i> attribute is a relative path. We cannot specify a protocol, hostname, or port number.• Even though the name of the attribute is <i>page</i>, it can point to a servlet.• You can also include an HTML file dynamically.
<p>✧ To forward a request to another component dynamically, use the <code>forward</code> action:</p> <pre><jsp:forward page="relativeURL" /></pre>	<p>Points to remember for <code><jsp:forward></code>:</p> <ul style="list-style-type: none">• There is only one attribute: <i>page</i>. There is no <i>flush</i> attribute in <code><jsp:forward></code>.• The name of the attribute is <i>page</i> and not <i>url</i>.• The attribute's value is a relative path. We cannot specify a protocol, hostname, or port number.• Even though the name of the attribute is <i>page</i>, the value can point to a servlet.• You can also forward to an HTML file dynamically.

Include Directive vs. Include Action

	Include directive	Include action
Syntax	<pre><%@ include file='relativeURL'%></pre> <p>The attribute name is <i>file</i> and not <i>page</i>.</p>	<pre><jsp:include page='relativeURL' flush='true false' /></pre> <p>The attribute name is <i>page</i> and not <i>file</i>.</p>
relativeURL	<p>It can point to any file; JSP, HTML, text, XML, etc. However:</p> <ul style="list-style-type: none"> • It <i>cannot</i> point to a servlet. • It <i>cannot</i> be a request-time expression. 	<p>It can point to any file; JSP, HTML, text, XML, etc. Also:</p> <ul style="list-style-type: none"> • It <i>can</i> point to a servlet. • It <i>can</i> be a request-time expression. <pre><%= expr %></pre>
Parameters	<p>The including JSP file <i>cannot</i> pass new parameters to the included JSP files. The following is <i>not</i> valid:</p> <pre><%@ include file='other.jsp?abc=123'%></pre> <p>However, the included files can access the original parameters available in the request to the including JSP file.</p>	<p>The including JSP page <i>can</i> pass new parameters to the included JSP pages and servlets using <code><jsp:param></code> as</p> <pre><jsp:include page='other.jsp' > <jsp:param name='abc' value='123' /> </jsp:include></pre> <p>The included pages can also access the original parameters present in the request to the including JSP page.</p>
Inclusion	<p>Inclusion is static.</p> <p>Inclusion of the included file happens at translation time.</p>	<p>Inclusion is dynamic.</p> <p>Inclusion of the output of the included component happens each time the including page is requested.</p>
Translation Unit	<p>The including page and the included pages <i>become</i> a part of a single translation unit.</p> <p>The page directives in any of the components including or included JSP pages affect the entire translation unit.</p>	<p>The including page and the including pages <i>do not become</i> a part of a single translation unit.</p> <p>The page directives in the including page do not affect the included components, and vice versa.</p>
If included file changes	<p>The changes are not reflected unless all the including pages are re-translated.</p>	<p>The changes are reflected automatically each time the pages are requested.</p>
Uses	<p>Files that do not change very often are included using directives.</p> <p>Examples include copyright information and navigational bars.</p>	<p>Components that do change often are included using actions.</p> <p>Examples include news headlines and advertisement bars.</p>

<jsp:include> vs. <jsp:forward>

	<jsp:include>	<jsp:forward>
Syntax	<pre><jsp:include page='relativeURL' flush='true false' /></pre>	<pre><jsp:forward page='relativeURL' /></pre>

CHAPTER 13—CREATING JSPTS WITH THE EXPRESSION LANGUAGE (EL)

Objectives 7.1–7.3

7.1 Given a scenario, write EL code that accesses the following implicit variables:

- *pageScope*, *requestScope*, *sessionScope*, and *applicationScope*
- *param* and *paramValues*, *header* and *headerValues*
- *cookie*, *initParam* and *pageContext*

Implicit objects	Purpose, Function, or Uses
<i>pageScope</i> , <i>requestScope</i> , <i>sessionScope</i> , and <i>applicationScope</i>	Returns a Map relating scope attributes and their values. For example, <code>#{sessionScope.attrname}</code> evaluates to the <code>attrname</code> attribute of the session scope.
<i>param</i> and <i>paramValues</i>	Returns a Map containing a named request parameter (<i>param</i>) or a set of parameters (<i>paramValues</i>). Similar to <code>getParameter(String)</code> and <code>getParameterValues(String[])</code> . For example, <code>#{param.paramname}</code> evaluates to the request parameter named <code>paramname</code> .
<i>header</i> and <i>headerValues</i>	Returns a Map containing a named header element (<i>header</i>) or a set of header elements (<i>headerValues</i>). For example, <code>#{header.accept}</code> returns the 'accept' field of the incoming header.
<i>cookie</i>	Returns a Map containing all of the cookies in a single object. Same as <code>getCookies()</code> method. For example, <code>#{cookie.cname}</code> returns the cookie named <code>cname</code> .
<i>initParam</i>	Returns a Map relating the context's initial parameter names and their values. For example, <code>#{initParam.iname}</code> returns the value of the initial parameter named <code>iname</code> .
<i>pageContext</i>	Object of type <code>java.util.Map</code> . Used for storing and retrieving page-related information and sharing objects within the same translation unit and same request. Also used as a convenience class that maintains a table of all the other implicit objects. For example, <code>#{pageContext.out.bufferSize}</code> returns the size of the <code>JspWriter</code> 's buffer.

7.2 Given a scenario, write EL code that uses the following operators:

- property access (the `.` operator) and collection access (the `[]` operator)
- arithmetic operators, relational operators, and logical operators

Important concepts	Exam tips
<p>✦ The property access operator (<code>.</code>) allows you to access properties of a variable. For example, <code>{\$.b}</code> evaluates to the property of <code>a</code> with the identifier, <code>b</code>. If <code>b</code> is an array, you can access its <code>n</code>th element with <code>{\$.b[n]}</code></p>	Don't use quotes with property access operators.
<p>✦ The collection access operator (<code>[]</code>) allows you to access properties of a Map, List, or Array. For example, <code>{\$.[b]}</code> evaluates to the value of <code>a</code> associated with the key or index, <code>b</code>. If <code>b</code> is an array, you can access its <code>n</code>th element with <code>{\$.b[n]}</code></p>	If the key is a string, you should surround its name with single or double quotes.
<p>✦ EL's arithmetic operators are similar to those used in regular Java: Addition: <code>+</code> Subtraction: <code>-</code> Multiplication: <code>*</code> Division: <code>/</code> and <code>div</code> Modular division: <code>*</code> and <code>mod</code> For example, <code>\$(6 + 4)</code> evaluates to 10 <code>\$(26/2)</code> evaluates to 13 <code>\$(34 mod 5)</code> evaluates to 4 <code>\${-20.0/5}</code> evaluates to -4.0</p>	Can use floating point numbers and number in exponential notation (e.g. <code>1e10</code>)
<p>✦ EL's relational operators are also similar to those used in regular Java: Equality: <code>'=='</code> and <code>'eq'</code> Non-equality: <code>'!='</code> and <code>'ne'</code> Less than: <code>'<'</code> or <code>'lt'</code> Greater than: <code>'>'</code> or <code>'gt'</code> Less than or equal: <code>'<='</code> or <code>'le'</code> Greater than: <code>'>='</code> or <code>'ge'</code> For example, <code>\$(3 > 2)</code> evaluates to true <code>\$(4 >= 8)</code> evaluates to false</p>	
<p>✦ EL's logical operators are like those used in regular Java: Conjunction: <code>'&&'</code> and <code>'and'</code> Disjunction: <code>' '</code> and <code>'or'</code> Inversion: <code>'!'</code> and <code>'not'</code> For example, <code>\$(5 > 3) && (6 mod 2 == 0)</code> evaluates to true <code>\$(6 + 7 == 13) (2 div 3 == 1)</code> evaluates to true <code>\$(5 > 3) && !(3 div 4 == 2)</code> evaluates to true <code>\$(2 >= 10) (18 mod 3 == 2)</code> evaluates to false</p>	

7.3 Given a scenario, write EL code that uses a function;

- write code for an EL function; and
- configure the EL function in a tag library descriptor.

Important concepts

- ✧ Coding an EL function means creating a public Java class containing a public static method. This method must return a value that can be evaluated inside a JSP. The class must be stored in the /WEB-INF/classes directory.

For example,

```
public Example
{
    public static String upper( String x)
    {
        return x.toUpperCase();
    }
}
```

- ✧ To tell the JSP how to find and access the function, you must update a tag library descriptor (*.tld). This declares the function corresponding to the Java method. It requires <function> tags that enclose <function-class> and <function-signature> subelements.

For example,

```
<taglib...>
...
<function>
  <name>Upper</name>
  <function-class>Example</function-class>
  <function-signature>
    java.lang.String Upper(java.lang.String)
  </function-signature>
</function>
...
</taglib>
```

- ✧ To access the function from within the JSP, you must insert a taglib directive that matches the tag library descriptor to a prefix.

For example,

```
<%@ taglib prefix="up" uri="http://example.com/taglib" %>
```

Then, you could use the function, Upper, as

```
#{up:Upper(stringvar)}
```

continued on next page

Important concepts

- ✧ To match the TLD to the URI, you need to update the deployment descriptor:

```
<taglib>
  <taglib-uri>
    http://example.com/taglib
  </taglib-uri>
  <taglib-location>
    /WEB-INF/example.tld
  </taglib-location>
</taglib>
```

CHAPTER 14—USING JAVABEANS

Objective 8.1

8.1 Given a design goal, create a code snippet using the following standard actions:

- `jsp:useBean` (with attributes: 'id', 'scope', 'type', and 'class'),
- `jsp:getProperty`, and
- `jsp:setProperty` (with all attribute combinations).

Important concepts	Exam tips
<p>✧ To declare a JavaBean component, use the <code><jsp:useBean></code> action:</p> <pre><jsp:useBean attribute-list /></pre> <p>Valid attributes: id, scope, class, type, and bean-Name.</p> <p>The attribute id is mandatory, while scope is optional. At least one of the following combinations of class, type, and bean Name must be present: class, type, class and type, bean Name and type.</p> <p>The body of the <code><jsp:useBean></code> action can be used to initialize its properties. Example:</p> <pre><jsp:useBean id="address" scope="session" class="AddressBean" > <jsp:setProperty name="address" property="street" value="123 Main" /> </jsp:useBean></pre>	<p>The default value of the scope attribute is page.</p> <p>The beanName attribute can also be a JSP expression.</p> <p>beanName can be used to instantiate a class or a serialized object whereas class is only for a class.</p> <p>class uses the new keyword to instantiate a class and beanName uses <code>java.beans.Beans.instantiate()</code>.</p> <p>Using beanName and class together is illegal.</p>

continued on next page

Important concepts**Exam tips**

- ✧ To set a bean's property, use `<jsp:setProperty>`:
- ```
<jsp:setProperty name="address" property="city"
 value="Albany" />
```
- Valid attributes: name, property, param, and value.
- The attributes name and property are mandatory, while others are optional. The name attribute must refer to a bean that is already declared using a use-Bean action.
- Examples:
- ```
<jsp:setProperty name="aName" property="*" />:
```
- Sets all the properties for which there is a matching parameter in the request.
- ```
<jsp:setProperty name="aName" property="aProp" /> :
```
- Sets aProp using the parameter of the same name in the request.
- ```
<jsp:setProperty name="aName" property="aProp"
                 param="aParam"/>:
```
- Sets aProp using the parameter named aParam in the request.
- ```
<jsp:setProperty name="aName" property="aProp"
 value="aValue"/>:
```
- Sets aProp to the value aValue.
- ```
<jsp:setProperty name="aName" property="aProp"
                 value="<%= JSPExpression %>"/>:
```
- Sets aProp to the value returned by the expression.
- ✧ To access a JavaBean property, use `<jsp:getProperty>`.
- Valid attributes: name and property.
- Both are mandatory.

Using param and value together is illegal.

`<jsp:getProperty>` prints out the value.

continued on next page

Important concepts

Once the use of a JavaBean is declared using a `<jsp:useBean>` action, a variable by the given name is automatically declared in the servlet code. Therefore, besides the `setProperty` and `getProperty` actions, the bean can also be accessed through this variable in the scripting elements.

Example:

```
<jsp:useBean id="user"
             class="UserBean"
             scope="session" />

<%
//The bean is used in a scriptlet here.
//You can call methods on the object
//referred to by the user variable.
user.initialize();
out.println(user.getName());
%>
```

CHAPTER 15—USING CUSTOM TAGS

Objectives 6.6, 9.1–9.3

6.6 *Configure the deployment descriptor to declare one or more tag libraries, deactivate the evaluation language, and deactivate the scripting language*

Important concepts	Exam tips
<p>✧ The deployment descriptor (web.xml) can declare the presence of tag libraries for the web application.</p> <p>The format of this declaration is</p> <pre data-bbox="337 471 599 679"><taglib> <taglib-uri> www.uri.com/library </taglib-uri> <taglib-location> /WEB-INF/libs/library.tld </taglib-location> </taglib></pre>	<p>This subject is treated in greater depth in chapters 13, 16, and 17</p> <p>In this case, the URI is matched to the tag library (library.tld) at the specified location.</p> <p>This URI can be used in JSPs throughout the application. For example,</p> <pre data-bbox="781 565 1201 616"><%@ taglib prefix="lib" uri="www.uri.com/library" %></pre> <p>will direct the web container to the location of the tag library descriptor, library.tld.</p>
<p>✧ To de-activate the Expression Language, you need to create <jsp-property-group> elements with <url-pattern> and <el-ignored> subelements.</p> <p>For example,</p> <pre data-bbox="337 850 659 953"><jsp-property-group> <url-pattern>*.jsp</url-pattern> <el-ignored>>true</el-ignored> </jsp-property-group ></pre> <p>will de-activate EL in every JSP in the application.</p>	
<p>✧ De-activating scripting is similar, except you need <scripting-invalid> sub-elements.</p> <pre data-bbox="337 1089 746 1192"><jsp-property-group> <url-pattern>*.jsp</url-pattern> <scripting-invalid>>true</scripting-invalid> </jsp-property-group></pre>	

9.1 *For a custom tag library or a library of Tag Files, create the 'taglib' directive for a JSP page.*

Important concepts	Exam tips
<p>✧ Before creating a 'taglib' directive in a JSP, it is recommended to create a URI in the deployment descriptor.</p> <p>This is given by:</p> <pre><web-app> <!--...other stuff --> <taglib> <taglib-uri> http://www.manning.com/sampleLib </taglib-uri> <taglib-location> /WEB-INF/sampleLib.tld </taglib-location> </taglib> <!--...other stuff --> </web-app></pre>	<p>Remember the following points:</p> <ul style="list-style-type: none">• Each taglib element maps one URI to one location.• Understand the syntax of <taglib> thoroughly. There is no hyphen in <taglib>, but there is a hyphen in <taglib-uri> and <taglib-location>.• The value of <taglib-uri> can be an absolute URI, a root-relative URI, or a non-root-relative URI.• The value of <taglib-location> can be either a root-relative URI or a non-root-relative URI. It cannot be an absolute URI.• The value of <taglib-uri> must be unique in the deployment descriptor.• The value of <taglib-location> must point to a valid TLD resource path. It can be either a TLD file or a JAR file containing the TLD file at location META-INF/taglib.tld.
<p>✧ The syntax of the taglib directive is:</p> <pre><%@ taglib prefix="test" uri="http://www.manning.com/sampleLib" %></pre>	<p>Understand the syntax of a taglib directive thoroughly.</p>

9.2 Given a design goal, create the custom tag structure in a JSP page to support that goal.

Important concepts	Exam tips
<ul style="list-style-type: none">✦ An empty tag: <code><test:required /></code> or <code><test:required></test:required></code>	You cannot nest a custom tag in the attribute list of another custom tag like this: <code><test:tag1 name="<test:tag2 />" /></code>
<ul style="list-style-type: none">✦ A tag with attributes: <code><test:greet user="john"/></code>	In this case, the JSP engine may assume " <code><test:tag2/"></code> " as a String value passed to the name attribute. It will not execute tag2.
<ul style="list-style-type: none">✦ A tag that surrounds JSP code: <code><test:debug></code> Some code here <code></test:debug></code>	
<ul style="list-style-type: none">✦ Nested custom tags: <code><test:switch></code> <code><test:case></code> Some JSP code Here <code></test:case></code> <code><test:default></code> Some JSP code Here too <code></test:default></code> <code></test:switch></code>	

9.3 Given a design goal, use an appropriate JSP Standard Tag Library (JSTL v1.1) tag from the "core" tag library.

Important concepts	Exam tips
<ul style="list-style-type: none">✦ JSTL's core library contains a number of tags that can be directly inserted into JSPs using the prefix, c. <p>General-purpose JSTL tags include <code><c:catch></code>, which catches exceptions within a JSP, and <code><c:out></code>, which displays the result of its value parameter.</p> <p>For example,</p> <pre><c:catch var="e"> Java code </c:catch></pre> <p>will hold the exception within the variable, e.</p> <pre><c:out value="{expr}" /></pre> <p>will send the value of expr to the JspWriter for display.</p>	To use JSTL tags, you need to insert the appropriate TLD in the <code>/WEB-INF/libs</code> directory and declare it with: <pre><%@ taglib uri="http://www.sun.com/jstl/core_rt" prefix="c" %></pre>

continued on next page

Important concepts

JSTL provides variable support with `<c:set>`, assigns a value to a variable, and `<c:remove>`, which removes a variable from the specified scope.

For example,

```
<c:set var="x" value="12" />
```

will set the value of x equal to 12.

```
<c:remove var="x" scope="session" />
```

will remove x from the session scope.

- ❖ Four JSTL tags enable flow control within JSPs:

`<c:if>` performs the tasks within its body if the condition set by test is true. For example,

```
<c:if test="{num == '5'}">
  ${x}
</c:if>
```

will display the value of x if it equals 5.

`<c:choose>` contains a number of `<c:when>` tags that examine the condition set by the test parameter. For example,

```
<c:choose>
  <c:when test="{letter == 'a'}">
    task1
  </c:when>
  <c:when test="{letter == 'b'}">
    task2
  </c:when>
  <c:otherwise>
    task3
  </c:otherwise>
</c:choose>
```

will perform task1 if letter equals a, task2 if letter equals b, and task3 if neither of the previous conditions was met.

`<c:forEach>` iterates through the values of the var parameter, similar to Java's for loop. For example,

```
<c:forEach var="num" begin="0" end="9">
  ${num}
</c:forEach>
```

will display each value of num as it passes from 0 to 9.

continued on next page

Important concepts

`<c:forTokens>` functions similarly to `<c:forEach>`, but instead iterates through the tokens contained in the items string.

For example,

```
<set var="n" value="one, two, three" />
<c:forTokens var="x" items="{n}"
delims=",">
  ${x}
</c:forTokens>
```

will display each value of x as cycles through the comma-delimited list of tokens contained within n.

- ❖ JSPs can access URL information with three different JSTL tags:

`<c:url>` rewrites the URL specified by the value parameter and contains it within the var variable. For example,

```
<c:url value="/index.html" var="page" />
```

contains the encoded URL within page.

`<c:import>` enables you to access and include content from outside the web application. For example,

```
<c:import url="/index.html" />
```

includes the content identified by the url parameter.

`<c:redirect>` functions similarly to the `sendRedirect()` method. For example,

```
<c:redirect url="/newpage.html" />
```

redirects processing to the page identified by the url parameter.

CHAPTER 16—DEVELOPING CLASSIC CUSTOM TAG LIBRARIES

Objectives 10.1–10.3

10.1 Describe the semantics of the “Classic” custom tag event model when each event method (*doStartTag*, *doAfterBody*, and *doEndTag*) is executed, and explain what the return value for each event method means; and write a tag handler class.

Important concepts	Exam tips
<p>✧ <i>doStartTag()</i> is called after the JSP engine completely parses the opening tag. Before calling <i>doStartTag()</i>, the JSP engines calls the following methods:</p> <ol style="list-style-type: none">1. <i>setPageContext()</i>2. <i>setParent()</i>3. setter methods for attributes <p><i>doStartTag()</i> can return three values:</p> <ol style="list-style-type: none">1. SKIP_BODY—Do not process the content of the body. Ignore it completely.2. EVAL_BODY_INCLUDE—Process the contents of the body as with the normal JSP code.3. EVAL_BODY_BUFFERED—Process the contents of the body as with the normal JSP code, but the output should be buffered and not sent to the client. The JSP engine uses a stack of <code>javax.servlet.jsp.tagext.BodyContent</code> objects for buffering.	<p>The Tag interface defines <i>doStartTag()</i>. Implementation classes of Tag interface can return only two values in <i>doStartTag()</i>: SKIP_BODY or EVAL_BODY_INCLUDE</p> <p>The IterationTag interface extends the Tag interface, but does not add any new return values for <i>doStartTag()</i>. Implementation classes of the IterationTag interface can return only one of the two values in <i>doStartTag()</i>: SKIP_BODY or EVAL_BODY_INCLUDE.</p> <p>The BodyTag interface extends the IterationTag interface, and adds a new return value, EVAL_BODY_BUFFERED, for <i>doStartTag()</i>. Implementation classes of BodyTag interface can return any one of the three values in <i>doStartTag()</i>: SKIP_BODY, EVAL_BODY_INCLUDE, or EVAL_BODY_BUFFERED.</p> <p>Thus, <i>doStartTag()</i> can return EVAL_BODY_BUFFERED <i>only if</i> the handler class implements the BodyTag interface.</p>

continued on next page

Important concepts**Exam tips**

✧ doAfterBody() is called after the JSP engine completely evaluates the entire body of the tag. This happens the first time only if:

1. The tag implements IterationTag and doStartTag() returns EVAL_BODY_INCLUDE.
OR
2. The tag implements BodyTag and doStartTag() returns EVAL_BODY_INCLUDE.
OR
3. The tag implements BodyTag and doStartTag() returns EVAL_BODY_BUFFERED.

The body of the tag is evaluated again (repeatedly), that is, the doAfterBody() is called repeatedly, only if:

1. The tag implements IterationTag and the previous call to doAfterBody() returns EVAL_BODY_AGAIN.
OR
 2. The tag implements BodyTag and the previous call to doAfterBody() returns EVAL_BODY_AGAIN.
OR
 3. The tag implements BodyTag and the previous call to doAfterBody() returns EVAL_BODY_BUFFERED.
-

doAfterBody() is not called for tags that implement only the Tag interface.

continued on next page

Important concepts	Exam tips
<p>doAfterBody() can return three values:</p> <ol style="list-style-type: none"> 1. EVAL_BODY_AGAIN—Evaluate the body of the tag again. Do not use buffering. 2. EVAL_BODY_BUFFERED—Evaluate the body of the tag again, but the output of the tag should be buffered. 3. SKIP_BODY—Do not process the content of the body again. Ignore it. The loop is over and doAfterBody() is not called again. 	<p>The Tag interface does not have doAfterBody().</p> <p>The IterationTag interface defines doAfterBody().</p> <p>Implementation classes of the IterationTag interface can return only one of two values in doAfter(): EVAL_BODY_AGAIN or SKIP_BODY.</p> <p>The BodyTag interface extends the IterationTag interface, and adds a new return value, EVAL_BODY_TAG, for doAfter Body(). EVAL_BODY_TAG is deprecated.</p> <p>Implementations classes of the BodyTag interface can return three values in doAfterBody(): EVAL_BODY_AGAIN, EVAL_BODY_BUFFERED, or SKIP_BODY.</p> <p>Thus, doAfterBody() can return EVAL_BODY_BUFFERED <i>only if</i> the handler class implements the BodyTag interface.</p>
<p>✧ doEndTag() is always called at the end of processing a tag.</p> <p>doEndTag() can return two values:</p> <ol style="list-style-type: none"> 1. SKIP_PAGE—Do not process the rest of the JSP page. Ignore it completely. 2. EVAL_PAGE—Process the rest of the JSP page as with the normal JSP code. 	<p>The Tag interface defines doEndTag().</p> <p>Implementation classes of the Tag interface can return two values in doEndTag(): SKIP_PAGE or EVAL_PAGE</p> <p>IterationTag and BodyTag inherit doEndTag().</p> <p>They do not add any new return values. Implementation classes of IterationTag and BodyTag also can return one of the two values in doEndTag(): SKIP_PAGE or EVAL_PAGE.</p> <p>doEndTag() is always called at the end of processing a tag regardless of the interfaces implemented and regardless of the return values from doStartTag() and doAfterBody()</p>

10.2 Using the *PageContext* API, write tag handler code to access the JSP implicit variables and access web application attributes.

Getting Implicit Objects in the Tag Handler

Implicit objects	Getting implicit objects	
	Using convenience methods	Using constants
application	pageContext.getServletContext()	pageContext.getAttribute(PageContext.APPLICATION)
session	pageContext.getSession()	pageContext.getAttribute(PageContext.SESSION)
request	pageContext.getRequest()	pageContext.getAttribute(PageContext.REQUEST)
response	pageContext.getResponse()	pageContext.getAttribute(PageContext.RESPONSE)
out	pageContext.getOut()	pageContext.getAttribute(PageContext.OUT)
config	pageContext.getConfig()	pageContext.getAttribute(PageContext.CONFIG)
page	pageContext.getPage()	pageContext.getAttribute(PageContext.PAGE)
pageContext		pageContext.getAttribute(PageContext.PAGECONTEXT)
exception	pageContext.getException()	pageContext.getAttribute(PageContext.EXCEPTION)

Getting Page Attributes in the Tag Handler

Scope	Getting attributes in different scopes	
	Using implicit objects	Using constants
application	pageContext.getServletContext().getAttribute("name")	pageContext.getAttribute("name", PageContext.APPLICATION_SCOPE)
session	pageContext.getSession().getAttribute("name")	pageContext.getAttribute("name", PageContext.SESSION_SCOPE)
request	pageContext.getRequest().getAttribute("name")	pageContext.getAttribute("name", PageContext.REQUEST_SCOPE)
page	pageContext.getAttribute("name")	pageContext.getAttribute("name", PageContext.PAGE_SCOPE)

10.3 Given a scenario, write tag handler code to access the parent tag and an arbitrary tag ancestor.

Important concepts	Exam tips
<p>✧ The Tag interface has two methods:</p> <ul style="list-style-type: none">• <code>public void setParent(Tag parentTag)</code>• <code>public Tag getParent()</code> <p>The container calls <code>setParent()</code> before calling <code>doStartTag()</code>. It is the responsibility of the tag implementation class to save this reference in a private member for later use.</p> <p>When the <code>getParent()</code> method is called, the tag returns its parent tag (the outer handler).</p>	
<p>✧ The TagSupport class implements the Tag interface and provides implementation for the <code>setParent()</code> and <code>getParent()</code> methods.</p> <p>So a class derived from the TagSupport class need not maintain its own parent, nor does it need to implement these methods. It can call <code>getParent()</code> to retrieve the outer handler, and subsequently call <code>getParent()</code> on the returned value to get ancestors.</p>	
<p>✧ The TagSupport class also provides a new convenience method:</p> <pre>public static final Tag findAncestorWithClass(Tag from, java.lang.Class class)</pre> <p>This method works outward within the nested tags and gets the instance of a given class type that is closest to the given tag instance.</p>	<p>The <code>findAncestorWithClass()</code> is a static method. Hence, it is <i>not</i> necessary to subclass TagSupport to use this method. Even simple tag handlers that directly implement the Tag interface can use <code>TagSupport.findAncestorWithClass()</code>.</p>

CHAPTER 17—DEVELOPING “SIMPLE” CUSTOM TAG LIBRARIES

Objectives 10.4–10.5

10.4 Describe the semantics of the “Simple” custom tag event model when the event method (*doTag*) is executed; write a tag handler class; and explain the constraints on the JSP content within the tag.

Important concepts	Exam tips
<ul style="list-style-type: none">✧ When the web container encounters a ‘simple’ custom tag, it performs five main steps:<ol style="list-style-type: none">1. It begins by creating an instance of the tag handler class.2. It invokes <code>setParent()</code> and <code>setJspContext()</code> on the tag handler. This provides a means for it to interact with the application.3. If the simple tag contained attributes, the web container will execute the (setter) methods corresponding to each.4. It invokes <code>setJspBody()</code> to store the tag body for processing by the tag handler.5. It invokes the tag handler’s <code>doTag()</code> method, which performs the tag’s logic.✧ A ‘simple’ tag handler must extend <code>SimpleTagSupport</code> and provide an implementation of the <code>doTag()</code> method. For example,<pre>public class ExampleTag extends SimpleTagSupport { public void doTag() throws JspException, IOException { getJspContext().getOut().write("Example"); } }</pre>will display a string whenever the simple tag is encountered.✧ The body of a simple tag can contain text and EL expressions, but no scripts (declarations, expressions, or scriptlets). Therefore, the <code><body-content></code> element of a simple tag TLD cannot be JSP.	<p>In the ‘classic’ model, the web container doesn’t always create a new tag handler instance. This is an important difference between the ‘classic’ and ‘simple’ methodologies.</p> <p>Unlike ‘classic’ handlers, which contain multiple event methods, a ‘simple’ handler performs all of its event processing with <code>doTag()</code>.</p>

10.5 *Describe the semantics of the Tag File model; describe the web application structure for tag files; write a tag file; and explain the constraints on the JSP content in the body of the tag.*

Important concepts

- ✧ Tag files contain JSP syntax code and must end in .tag or .tagx.
 - ✧ Tag files don't need tag library descriptors. Instead, they are referred to by their directory in a JSP taglib statement with the tagdir attribute.
For example,

```
<%@ taglib prefix="ex" tagdir="/WEB-INF/tags" %>
```

assigns the prefix "ex" for tags in the given directory. To refer to an individual file, use the tag <ex:filename> within the JSP.
 - ✧ Tag files must be placed in the /WEB-INF/tags directory or a subdirectory.
 - ✧ Tag files provide implicit variables and additional directives and actions for processing.
 - ✧ In particular, the <jsp:doBody> enables a tag file to process its body content. This content can be regular text or EL expressions, but it cannot contain script elements (declarations, expressions, and scriptlets).
-

CHAPTER 18—DESIGN PATTERNS

Objectives 11.1 and 11.2

11.1 Given a scenario description with a list of issues, select a pattern that would solve the issues. The list of patterns you must know are:

- *Intercepting Filter*,
- *Model-View-Controller*,
- *Front Controller*,
- *Service Locator*,
- *Business Delegate*
- *Transfer Object*

Issues	Pattern
<ul style="list-style-type: none">• Receive requests before other elements• Pre-process requests using filters• Redirect requests to different resources• Performs necessary post-processing on outgoing responses	Intercepting Filter
<ul style="list-style-type: none">• Flexible design• Allows different designers to focus on different aspects of the application• Provide services to different clients: web client, WAP client, etc.• Multiple views, such as HTML or WML	Model-View-Controller (MVC)
<ul style="list-style-type: none">• Central point of receiving requests• Single resource to enforce security• Prevents having to alter many different resources• Applies policies consistently across application	Front Controller
<ul style="list-style-type: none">• Provides a directory for accessing service resources• Centralizes methods of connection and access• Makes use of Java Naming and Directory Interface (JNDI)	Service Locator
<ul style="list-style-type: none">• Reduces coupling between presentation and business tiers• Proxy for the client• Client-side facade• Caches business service references for presentation-tier components• Caches business service results for presentation-tier components• Encapsulates business service lookup• Encapsulates business service access• Decouples clients from business service API	Business Delegate
<ul style="list-style-type: none">• Previously known as Value Object• Small object• Grouped information• Read-only data• Reduces network traffic• Increases response speed• Transfers data across networked tiers	Transfer Object

11.2 Match design patterns with statements describing potential benefits that accrue from the use of the pattern, for any of the following patterns:

- *Intercepting Filter,*
- *Model-View-Controller,*
- *Front Controller,*
- *Service Locator,*
- *Business Delegate*
- *Transfer Object*

Important concepts	Exam tips
<p>✧ Intercepting Filter The Intercepting Filter design pattern is used where requests and/or responses need to be processed in a consistent manner. This filter wraps around the application, receiving requests as they come in and processing responses before they go out.</p>	<p>The potential benefits of Intercepting Filters are:</p> <ul style="list-style-type: none"> • Apply request pre-processing consistently • Central point of response post-processing • Redirect requests to specific resources
<p>✧ Model-View-Controller (MVC) The Model-View-Controller design pattern is applicable in situations where the same data (Model) is to be presented in different formats (Views), but is to be managed centrally by a single controlling entity (Controller).</p>	<p>The potential benefits of MVC are:</p> <ul style="list-style-type: none"> • Flexible design • Centrally managed data • Multiple ways of presentation
<p>✧ Front Controller The Front Controller design pattern serves as the primary gate for requests entering the application. This object enforces security restrictions and controls the view shown to the client.</p>	<p>The potential benefits of Front Controllers are:</p> <ul style="list-style-type: none"> • Central point for selecting and screening requests • Controls view for incoming requests • Request processing can be changed by altering a single object
<p>✧ Service Locator The Service Locator design pattern provides a central directory for resources to locate services across the enterprise. This system controls the communication methodology to be used (recommended: Java Naming and Directory Interface (JNDI))</p>	<p>The potential benefits of Service Locators are:</p> <ul style="list-style-type: none"> • Centralizes the process of looking up services • Controls connectivity and the means of directory access • Can provide cache for repeated resource requests
<p>✧ Business Delegate A Business Delegate is an object that communicates with the business service components on behalf of the client components.</p> <p>The client-side components delegate the work of accessing the business services to the Business Delegate object.</p>	<p>The potential benefits of Business Delegates are:</p> <ul style="list-style-type: none"> • Reduced coupling between presentation and business tiers • Cached business service results for presentation-tier components. • Business service lookup encapsulated • Business service access encapsulated • Decoupled clients from business service

continued on next page

Important concepts

- ◇ Transfer Object
A Transfer Object is a small-sized serializable Java object that is used for transferring data over the network in a distributed application.

Exam tips

- The potential benefits of Transfer Objects are:
- Less communication overhead
 - Fewer number of remote calls
 - Reduction in network traffic
 - Increased response speed
-

index

Symbols

`%=` delimiter 166, 170
`<%@ include %>`. *See* include directive

A

absolute URI 287
abstraction 377
access control list 140
accessing JavaBean 269
ACL. *See* access control list
actions 19, 166, 171, 364, 366, 368, 372
 custom 286
 forward 225
 include 223
 plugin 171
 standard 172
 syntax 172
active resources 22
Active Server Pages 15
addCookie() 43
addDateHeader() 43
addHeader() 43
addIntHeader() 43
alias 366
Apache
 Jakarta Project 287
 Software Foundation 8
 tag libraries 287
 Tomcat 8
 Web Server 7

application 200
 events 88
 scope 207
 server 22
 state 120
APPLICATION_SCOPE 210
applicationScope 236, 239, 500
architecture
 in JSP 18
 J2EE 15
 Model 1 18
 Model 2 18
 multi-tier 400
attacks 141
attrib 367
attribute 313–314, 357, 359–362, 365–374
attribute scopes 55
 context scope 55
 request scope 55
 session scope 55
attributeAdded() 86–87
attributeRemoved() 87
attributeReplaced() 87
auditing 141
authentication 140
authentication mechanisms 142
 web applications 146
authorization 140
authorization constraint 149
autoFlush attribute 184

B

Base64 145
basic authentication 143
bean
 containers 252
 initialization 265
 variable scope 265
beanName 260
BigDecimal 242, 249
BigInteger 242, 249
body content 313, 316, 353–360, 362–363, 366–374
 empty 316
 JSP 317
 tagdependent 317
BodyContent 319
 getEnclosingWriter() 337
 getString() 337
BodyTag 319, 333, 353–357, 359, 362
 doInitBody() 333–334
 EVAL_BODY_BUFFERED 333–334
 EVAL_BODY_TAG 333
 example 336
 setBodyContent() 333–334
BodyTagSupport 319, 339, 355–357, 374
 methods 339
body-value 363
buffer attribute 184
bufferSize 239

- business delegate 393
- business logic 237
- C**
- caching
 - GET vs. POST 33
 - response page 43
 - results of remote invocations 396
 - static data 47
 - Value List Handler pattern 381
- CGI scripts 5
 - See also* Common Gateway Interface
- class attribute 259
- class files 69
- classes directory 69
 - WEB-INF 69
- Class.forName() 46
- client authentication 145
- code reviews 141
- collection access 236, 241–242, 249, 501
- comments 166, 172
- committed 44
- Common Gateway Interface 4
 - limitations 5
- compilation
 - of JSP class into servlets 176
 - page 175
 - phase 176
 - servlets 9
- components
 - application 18
 - controller 17
 - J2EE 22
 - JavaBeans 251
 - of URI 76
 - reusing JSP 220
 - reusing software 219
 - web 22
- conditional statements 191
- confidentiality 141
- config 204
- configuration web application 90
- CONNECT 25
- containsHeader() 43
- content type, common values 42
- Content-Length 25
- Content-Type 25
- contentType attribute 185
- context
 - path 76
 - scope 55
 - See also* Servlet Context
 - contextDestroyed() 88–89
 - contextInitialized() 88–89
- cookies 43, 121, 132, 236, 239–240, 500
- co-operative tags 343
- CORBA 379
- custom tags 172
 - and JavaBeans 347
 - as custom actions 286
 - attributes 295
 - body content 296
 - buffering the body content 339
 - descriptor 287
 - empty 294
 - hierarchy 322
 - informing JSP engine 288
 - libraries 287
 - nested 297
 - prefix 293
 - usage 293, 298
 - usage in JSP pages 293
 - validation 347
- D**
- data access object 382
- data integrity 141
- data privacy 141
- <declaration> 211, 213
- declarations 168
 - and variable initialization 191
 - JSP 20
 - JSP syntax 166
 - jspDestroy() 178
 - jspInit() 178
 - order in JSP 190
 - order of 190
 - syntax 168
 - translated as 189
 - XML syntax 213
- declarative security 149
 - example 152
- default web application 70–71
- :definition 90
- DELETE 25
- delimiters 17, 170
- denial of service attacks 142
- deployment descriptor 23, 46, 70, 246, 248–249
 - overview 71
 - properties 71, 238, 355–356, 367, 369
 - sample 72
 - servlet 73
- description 312–313, 315
- design patterns 377
 - business delegate 393
 - data access object 382
 - distributed 379
 - front controller 389
 - gang of four 377
 - J2EE 379
 - model-view-controller 18, 116, 385–386, 391
 - page-by-page iterator 381
 - paged-list 381
 - required for SCWCD 382
 - service activator 382
 - service-locator 382
 - tiers 380
 - value list handler 381
 - value object 397
- destroy() 48, 104
- digest authentication 145
- <directive.include> 213
- <directive.page> 213
 - example 211
- directives 167, 288
 - include 167, 220
 - JSP 19, 167
 - JSP syntax 166
 - page 167
 - syntax 168
 - taglib 167, 288
 - translated as 189
 - XML syntax 213

- display-name 312–313
- distributed environment 92
 - HttpSession 93
 - ServletContext 92
 - session migration 129
 - systems 21
 - web applications 92
- div 242
- doAfterBody() 329, 353, 356
- document root 68
- doDelete() 35
- doEndTag() 320, 322, 353, 356
- doFilter() 104–105
- doGet() 35
- doHead() 35
- doInitBody() 333–334
- doOptions() 35
- doPost() 35
- doPut() 35
- doStartTag() 320, 322, 353, 356
- doTag() 353–358, 360–362, 372–374
- doTrace() 35
- doXXX() 35–36
 - parameters 36
- dynamic attributes 359–360, 362, 367–368, 372–373
- dynamic inclusion 223
 - passing parameters 226
 - sharing objects 228
- DynamicValues 359

E

- EJB container 22
- EL 236–249, 356, 361–362, 364, 366–367
- empty tag 294, 324
 - with attribute 326
- empty value 316
- encodeRedirectURL() 134
- encodeURL() 134
- Enterprise JavaBeans
 - components 380
 - support 7
- error conditions 45
- errorPage attribute 182
- escape sequences 194
 - in attributes 195
 - in scripting elements 195
 - in template text 194
- EVAL_BODY 354, 373
- EVAL_BODY_AGAIN 329–330
- EVAL_BODY_BUFFERED 333–334
- EVAL_BODY_INCLUDE 321–322
- EVAL_BODY_TAG 333
- EVAL_PAGE 321, 323, 354, 373
- events listeners 85, 88
 - :example 108, 291, 410
- exception implicit variable 206
- explicit mapping 290, 313
- Expression Language 236–238, 240–241, 245, 249, 355, 371
- <expression> 212–213
- expressions 170, 237–238, 240–242, 247, 249
 - and implicit variable out 203
 - JSP 20, 170
 - JSP syntax 166
 - request-time attribute 194, 213
 - syntax 170
 - translated as 189
 - valid and invalid 170
 - XML syntax 213
- extends attribute 184

F

- fail over 92
- Filter 103
 - destroy() 104
 - doFilter() 104
 - init() 103
- FilterChain 105
 - doFilter() 105
- FilterConfig 105
- filters 98
 - API 102
 - chain 98
 - configuration 106–107
 - example 100
 - in MVC 116
 - threading 116
 - uses 99
- findAncestorWithClass() 338, 345–346, 356
- findAttribute() 210
- forEach 367
- Form-based authentication 146
 - advantages 146
 - disadvantages 146
- forward action, usage 225
- forward() 57, 203
- <forward> 20, 171, 223
- forwarding a resource 57–58
- fragment 362, 368–370
- front controller 389
- FTP 4
- function 236, 241, 244–249
- function-class 245–246
- function-signature 246

G

- Gang of Four 377
- GenericServlet 11
 - getServletContext() 85
 - init() 47
- GET 25, 33
 - features 33
- getAttribute() 56, 123, 210
- getAttributeNames() 56
- getAttributeNamesInScope() 210
- getAttributeScope() 210
- getBodyContent() 339, 357, 362, 374
- getEnclosingWriter() 337
- getExpressionEvaluator() 355
- getFilterName() 105
- getHeaderNames() 39
- getHeaders() 39
- getInitParameter() 50, 84, 105
- getInitParameterNames() 50, 84, 86, 105
- getJspBody() 356–357, 362, 369, 373
- getServletContext() 354, 356–358, 361–363, 373–375, 471
- getNamedDispatcher() 58
- getOut() 354–355, 361–363, 374
- getOutputStream() 41–42
- getParameter() 37–38, 239

- getParameterNames() 37–38
- getParameterValues() 37–38, 156
- getParent() 320, 355, 357, 372–373, 375
- getPreviousOut() 339
- getProperty(), automatic type conversion 278
- <getProperty> 269
- getRealPath() 54, 179
- getRemoteUser() 156
- getRequestDispatcher() 57
- getResource() 53
- getResourceAsStream() 53
 - limitations 54
- getServletContext() 50, 105, 238
- getServletInfo() 189
- getServletName() 50, 52
- getServletSession() 238
- getSession() 123
- getString() 337
- getUserPrincipal() 156
- getValue() 339
- getValues() 339
- getVariableResolver() 355
- getWriter() 42
- GoF. *See* Gang of Four

H

- hashcode 141
- HEAD 25, 34
- header 24, 26, 38–40, 236, 239–241, 500
 - management 43
 - names 43
- headerValues 236, 239–242, 500
- HTML
 - comments 172
 - example Hello User 16
 - files on web server 5, 69
 - FORM and HTTP methods 33
 - FORM for authentication 146
 - MIME type 185
 - tables 361
 - tags and Java code 15
 - template 15
 - URL-rewriting 133
- HTML output
 - from custom tags 294, 324
 - from expressions 170
 - from JSP page 17
 - from scriptlets 169
 - using implicit variable 203
 - using PrintWriter 40
- HTTP 23, 239
 - basic authentication 143
 - advantages 144–146
 - disadvantages 144
 - basics 24
 - error conditions 45
 - GET 25
 - HEAD 25
 - methods 32
 - comparison 33
 - POST 25
 - PUT 26
 - request 24
 - response 26
 - status codes 45
- HTTP Digest authentication 145
 - advantages 145
 - disadvantages 145
- HTTP request, servlets 35
- HttpJspPage 177
- HTTPS 145
- HTTPS client authentication 145
 - advantages 145
 - disadvantages 146
- HttpServlet 12, 35
 - request processing 35
 - service() 35
- HttpServletRequest 12, 36–37
 - getHeader() 39
 - getHeaderNames() 39
 - getHeaders() 39
 - getRemoteUser() 156
 - getUserPrincipal() 156
 - isUserInRole() 156
 - methods to identify users 156
- HttpServletRequestWrapper 110
 - example 112
- HttpServletResponse 12, 43
 - containsHeader() 43
 - encodeRedirectURL() 134
 - encodeURL() 134
- sendRedirect() 44
- setDateHeader() 43
- setHeader() 43
- setIntHeader() 43
- status codes 45
- HttpServletResponseWrapper 110
 - example 112
- HttpSession 121, 239
 - distributed environment 93
 - example 122
 - getAttribute() 123
 - getSession() 123
 - invalidate() 130
 - isNew() 132
 - setAttribute() 123
 - setMaxInactiveInterval() 131
 - usage 122
- HttpSessionActivationListener 94, 129
 - sessionDidActivate() 129
 - sessionWillPassivate() 129
- HttpSessionAttributeListener 86–87, 125
 - attributeAdded() 86
 - attributeRemoved() 87
 - attributeReplaced() 87
- HttpSessionBindingEvent 125
- HttpSessionBindingListener 125, 128
 - example 125
 - valueBound() 125
 - valueUnbound() 125
- HttpSessionListener 126
 - example 127
 - sessionCreated() 126
 - sessionDestroyed() 126

I

- id attribute 259
- IllegalArgumentException 94
- IllegalStateException 42, 44
- implicit mapping 290, 312
- implicit objects 198, 200, 210, 259
 - accessing form custom tags 339
- implicit variables 198, 200, 236, 238, 240–241, 249
 - application 200

- implicit variables (*continued*)
 - config 204
 - declaration 200
 - exception 206
 - out 203
 - page 202
 - pageContext 202
 - request 202
 - response 202
 - session 201
- import attribute 182
- in process servlet container 6
- inactivity of session 122
- include action, usage 223
- include directive 167, 220
 - accessing variables 221
- include() 57, 203
- <include> 171
- including a resource 57–58
- info attribute 185
- init() 46, 103
- init(ServletConfig) 49
- integrity attacks 142
- invalidate() 130
- IP 121
- isELIgnored 367
- isErrorPage attribute 182, 206
- isNew() 132
- ISO-8859-4 42
- isUserInRole() 156
- IterationTag 319, 329, 353–354, 356–357, 359, 362
 - doAfterBody() 330
 - EVAL_BODY_AGAIN 330, 333
 - example 331
- iterative statements 191
- J**
- J2EE pattern catalog 381
- JAR
 - classpath 9
 - content type 41
 - file 69
 - jar command 70
 - location in a web application 69
- sending to browser 41
 - servlet.jar 9
- Java Standard Tag Library 237
- JavaBeans
 - accessibility 259
 - advantages 253
 - and custom tags 347
 - constructors 252
 - containers 252
 - conventions 252
 - declaration 258
 - example 252
 - in JSP actions 258
 - in scripting elements 274
 - in Servlets 271
 - indexed properties 278
 - initializing 265
 - non-string data types 276
 - persisting 255
 - properties 252
 - property types 276
 - requirements for JSP 252
 - scope 259
 - serialized 255
 - support in JSP for 254
 - using serialized beans 255
- java.io.Serializable 93
- java.lang.Math 242
- JavaServer Pages 15, 21, 166, 380
 - comparison with servlets 17
 - example 15
- javax.servlet package 10
- javax.servlet.http package 11
- JMS server 22
- JNDI server 22
- j_password 146
- j_security_check 146
- JSP 237–238, 240–241, 244–249, 317
 - actions 171
 - comments 172
 - directives 167
 - expressions 170
 - forward 171, 225
 - getProperty 171, 269
 - include 171
 - usage 223
 - plugin 171
 - setProperty 171, 266
 - useBean 171, 258
- JSP 2.0 353, 355, 359, 363–364, 371
- JSP architecture models 18
 - model 1 18
 - model 2 18
- JSP life-cycle methods
 - jspDestroy() 178
 - jspInit() 178
 - _jspService() 178
- JSP life-cycle phases
 - compilation 176
 - example 178
 - loading and instantiation 177
 - phases 175
 - translation 176
- JSP Model 2 architecture 116
- JSP page 15
 - life-cycle methods 177
 - XML syntax 211
- JSP page scopes 207
- JSP page translation
 - rules 189
- JSP scriptlet 237
- JSP scriptlets 169
- JSP syntax elements 166
- JspContext 355–356, 358, 361–362, 372, 374
- jspDestroy() 178
- JspException 319
- JspFragment 356–357, 362–363, 369, 372, 374
- jspInit() 178
- JspPage 177
 - _jspService() 178
- JspTag 355–356
- JspTagException 319
- jsp-version 312
- JspWriter 239, 355, 357, 362–363, 369–370, 374
- JSTL 237
- j_username 146
- L**
- language attribute 184
- large-icon 312–313

- lazy loading 47
- lib directory 69
- life-cycle methods
 - JSP 177
 - Servlet 45
- listener configuration 90
- listener interfaces. *See* listeners
- listeners 85, 88, 312
 - HttpSessionActivationListener 94
 - HttpSessionAttributeListener 86–87
 - ServletContextListener 88–89
- load balancing 92
- loading and instantiation JSP
 - phases 177
- logical
 - conjunction 243
 - disjunction 243
 - expressions 243
 - inversion 243
 - operators 236, 241, 501

M

- malicious code 141
- mapping
 - filter 107
 - JSP page to servlet 204
 - URL to servlet 75–76
- :method attribute 35
- MIME type 53
- MIME Type Mapping 71, 238, 355–356
- mime-mapping
 - example 72, 410
- mod 242
- model 1 architecture 18
- model 2 architecture 18
- model-view-controller 116, 385–387, 391
- mutating JavaBeans 266
- MVC. *See* model-view-controller

N

- name 267, 313, 315
- name-from-attribute 366

- nested custom tags 297, 343
- newInstance() 46
- non-empty tag 328
- non-root relative URI 287

O

- operators 236–237, 240–243, 249, 501
- OPTIONS 25
- out 203
- out-of-process servlet container 7
- output stream 41, 199, 237

P

- page directive 167
 - attributes 181
 - autoFlush attribute 184
 - buffer attribute 184
 - contentType attribute 185
 - errorPage attribute 182
 - extends attribute 184
 - import attribute 182
 - info attribute 185
 - isErrorPage attribute 182
 - language attribute 184
 - pageEncoding attribute 185
 - session attribute 182
- page implicit variable 202
- page relative URI 288
- page scopes 207, 209
- PageContext 355–356, 372
 - findAttribute() 210
 - forward() 203
 - getAttribute() 210
 - getAttributeNamesInScope() 210
 - getAttributesScope() 210
 - include() 203
 - removeAttribute() 210
 - setAttribute() 210
- pageContext 202, 236, 238–239, 241
- PageData 347
- page-encoding 367
- pageEncoding attribute 185
- PAGE_SCOPE 210

- pageScope 236, 238, 500
- param 236, 239–240, 248, 267, 500
- <param> 226
- paramValues 236, 239–241, 500
- passive resources 22
- path info 77
- paths, context, servlet info 76
- <plugin> 20, 171
- POST 25, 33
 - features 33
- prefix 359, 361, 364–372, 374
- prefix attribute 293
- preinitializing 47
- preloading. *See* preinitializing
- presentation logic 237
- PrintWriter
 - flush() 44
 - generating HTML 40
 - usage 40
- programmatic security 156
 - example 156
- property access 241
- property attribute 267
- property files 53
- proxy server 121
- public-key cryptography 142
- PUT 26

R

- Reader 362, 369–370
- readObject() 93
- redirecting request 44
- relational expressions 243
- relative path 54
- release() 321, 323
- reload() 33
- removeAttribute() 210
- removeValue() 339
- request headers, retrieving 38–39
- request implicit variable 202
 - redirecting 44
- request scope 55, 208, 210, 236, 239, 241, 500
- request URI, paths 76

- RequestDispatcher 57–58
 - forward() 57
 - include() 57
 - request-time attribute expressions
 - usage 194
 - required attribute 315
 - resource moved permanently 45
 - resource not found 45
 - response header
 - Date 43
 - Expires 43
 - Last-Modified 43
 - names 43
 - Refresh 43
 - setting 43
 - response implicit variable 202
 - sending 40
 - ROOT directory 71
 - root element 212
 - root relative URI 287
 - <root> 212
 - rtexprvalue 315, 360, 362, 368, 372
- S**
- scope attributes 238–239
 - scopes 207, 259
 - application 207
 - bean variable 265
 - page 207, 209
 - request 208
 - session 207
 - scripting elements 168, 171
 - usage 189
 - scriptless 362, 367, 373
 - <scriptlet> 213
 - scriptlets 169, 237
 - conditional and iterative 191
 - JSP 20, 169
 - JSP syntax 166
 - order in JSP 190
 - printing HTML 169
 - translated as 189
 - variable initialization 191
 - XML syntax 213
 - secrecy attacks 142
 - Secure Socket Layer 145
 - sendError() 45
 - sendRedirect() 44
 - Serializable 93
 - serialized beans, usage 262–263
 - server extensions 5
 - server-side includes 15
 - service(), overloading 35
 - Servlet 10
 - destroy() 48
 - init() 46
 - service() 47
 - servlet
 - container 5
 - relationship with Servlet API 10
 - types 5
 - context
 - initialization parameters 85
 - destroyed state 48
 - destroying 48
 - Hello World example 8
 - in deployment descriptor 50
 - initialization parameters 85
 - initialized state 46
 - initializing 46
 - instantiating 46
 - life cycle 45
 - methods 48
 - loaded state 46
 - loading 46
 - mapping 75–76
 - path 77
 - identification 77
 - pre-initializing 47
 - request processing 35
 - servicing state 47
 - state transition 48
 - unloaded state 48
 - unloading 48
 - Servlet API 10, 36
 - advantages and disadvantages 12
 - ServletConfig 50, 85, 204
 - example 51
 - getInitParameter() 50, 74
 - getInitParameterNames() 50
 - getServletContext() 50
 - getServletName() 50
 - methods 50
 - ServletContext 53, 84, 239
 - distributed environment 92
 - getInitParameter() 84
 - getInitParameterNames() 84, 86
 - getNamedDispatcher() 58
 - getRealPath() 54
 - getRequestDispatcher() 57
 - getResource() 53
 - getResourceAsStream() 53
 - initialization 84
 - ServletContextAttributeEvent 92
 - ServletContextAttributeListener
 - attributeAdded() 87
 - attributeRemoved() 87
 - attributeReplaced() 87
 - methods 87
 - ServletContextEvent 89
 - ServletContextListener 88–89
 - contextDestroyed() 88–89
 - contextInitialized() 88–89
 - example 88
 - ServletOutputStream 41
 - ServletRequest 11, 37, 239
 - getParameter() 37
 - getParameterNames() 37
 - getParameterValues() 37
 - getRequestDispatcher() 57
 - use 37
 - ServletRequestWrapper 110
 - ServletResponse 11, 40
 - getOutputStream() 41
 - getWriter() 40
 - setContentType() 42
 - ServletResponseWrapper 110
 - session 120, 201
 - accessibility 124
 - attribute 182
 - cookies 132
 - establishing 121
 - identifier 121
 - implementation 131
 - listener interfaces 124
 - scope 55, 207
 - timeout 122, 130, 132
 - URL rewriting 133

- session ID 121
- sessionCreated() 126
- sessionDestroyed() 126
- sessionDidActivate() 129
- SESSION_SCOPE 210
- sessionScope 236, 239, 500
- sessionWillPassivate() 129
- setAttribute() 56, 123, 210
- setBodyContent() 333–334
- setContentType() 42
- setDateHeader() 43
- setDynamicAttribute() 359–361, 372–373
- setHeader() 43
- setIntHeader() 43
- setJspBody() 355
- setJspContext() 355, 375, 471
- setMaxInactiveInterval() 131
- setPageContext() 321, 355
- setParent() 321, 355, 375, 471
- setProperty()
 - attributes 266
 - automatic type conversion 277
 - name 267
 - param 267
 - property 267
 - setting from request parameters 268
 - value 267
- <setProperty> 266
 - using request parameters 268
 - wild card 269
- setValue() 338
- short-name element 312
- SimpleTag 353–357, 369, 371–375, 471
 - example 359
 - processing body content 362–364
- SimpleTagSupport 353–358, 360, 371–372, 374
- SKIP_BODY 321–322, 354, 373
- SKIP_PAGE 321, 323
- small-icon 312–313, 367
- sniffing 141–142
- spoofing 142
- SSI 15
- SSL 145
- standalone servlet container 6
- state 120
- stateless protocol 23
- static attribute 360–361
- static inclusion 220
- status codes, sending 45

T

- Tag 318, 352–359, 372
 - body content 362–363
 - directives 366–368
 - jsp:invoke 369
 - TLDs 365
- tag 312–313, 320
 - doEndTag() 320, 322
 - doStartTag() 320, 322
 - EVAL_BODY_INCLUDE 321–322
 - EVAL_PAGE 321, 323
 - extension API 318
 - file 352, 364–374
 - files 352–353, 363–371, 374
 - getParent() 320–321
 - handlers 286, 355, 369, 371–372
 - tag files 352–353, 364–365
 - release() 321, 323
 - setPageContext() 321
 - setParent() 321
 - SKIP_BODY 321–322
 - SKIP_PAGE 321, 323
- tag library 287
 - descriptor 236, 245, 247, 249, 287, 310
 - DTD 311
 - example 310
 - location 289
 - resolution 291
- tag-class 313
- tagdependent 317, 362–363, 367, 373
- tagdir 364–368, 370–371, 374
- TagExtraInfo 347
- taglib 245–248, 358–361, 373–374
 - directive 167, 288
 - prefix 293
 - explicit mapping 290
 - location 312
 - map 290–291
 - SimpleTag 354
 - subelements 311–312
 - tag files 364–371
 - uri 312
- taglib-location 290
- TagLibraryValidator 347
- taglib-uri 290
- TagSupport 319, 338, 355–356, 374
 - findAncestorWithClass() 338
 - getValue() 339
 - methods 338
 - removeValue() 339
 - setValue() 338
- TagVariableInfo 347
- tei-class 313
- <text>, example 212
- throwable 206
- TLD 246, 248–249, 353, 358–363, 365–368, 371–372
 - Resource Path 290
 - See also tag library, descriptor
- tlib-version 245–246, 312
- Tomcat 8
 - configuring users 146
 - installation 403, 407
- tomcat-users.xml 146
- TRACE 25
- translation phase 176
- translation units 174
- transport-guarantee 151
 - CONFIDENTIAL 151
 - INTEGRAL 151
 - NONE 151
- trojan horse 141
- TryCatchFinally 347
- type attribute 259, 315
- types of URIs 287

U

- unauthorized access 45
- uniform resource identifier 23

- URI 23, 241, 244, 246–249, 312
 - absolute 287
 - non-root relative 287
 - path 58
 - root relative 287
 - types 287
 - URL 23
 - URL rewriting 121, 133
 - example 135
 - URN 23
 - useBean
 - attributes 258
 - attributes usage 260
 - beanName 260
 - class 259
 - id 259
 - scope 259
 - type 259
 - typecast problem 261
 - <useBean> 258
 - type attribute 264
 - user configuration 146
 - user data constraint 149
 - example 152
- V**
- validator 312
 - value 267
 - value object 397
 - valueBound() 125
 - valueUnbound() 125
 - variable directive 366
 - variable element 313
 - variable initialization 191
 - varReader 369–370
 - virus 141
- W**
- war. *See* web archive
 - web application 22, 90
 - directory structure 68
 - document root 68
 - in distributed environment 92
 - properties 90
 - server 22
 - WEB-INF 69
 - web archive 70
 - creation 70
 - web browser, HTTP methods 32
 - web container 237, 245, 353–354, 371, 375, 471
 - SimpleTag 355–356, 358–361
 - tag files 364–368
 - TLDs 246–247
 - web resource collection 149
 - example 150
 - web server 4–5
 - web site attacks 141
 - availability attacks 142
 - denial of service attacks 142
 - integrity attacks 142
 - secrecy attacks 142
 - webapps directory 68
 - WEB-INF 69
 - classes 69
 - lib 69
 - web.xml 70
 - web.xml 46, 70, 244, 246–248
 - servlet example 50
 - well-known URIs 290
 - worm 141
 - wrapper classes 106
 - usage 110
 - writeObject() 93
- X**
- XML 237, 244–245
 - XML Name Space 213
 - XML syntax for JSP pages 211
 - actions 214
 - comments 214
 - directives 213
 - root 212
 - scripting elements 213
 - text 214
 - xmlns 212
- Z**
- zip file 69

SCWCD Exam Study Kit **SECOND EDITION**

Java Web Component Developer Certification

H. Deshmukh • J. Malavia • M. Scarpino

With the tremendous penetration of J2EE in the enterprise, passing the Sun Certified Web Component Developer exam has become an important qualification for Java and J2EE developers. To pass the SCWCD exam (Number: 310-081) you need to answer 69 questions in 135 minutes and get 62% of them right. You also need \$150 and this (completely updated and newly revised) book.

In its first edition, the *SCWCD Exam Study Kit* was the most popular book used to pass this most desirable web development certification exam. The new edition will help you learn the concepts—large and small—that you need to know. It covers the newest version of the exam and not a single topic is missed.

The SCWCD exam is for Sun Certified Java Programmers who have a certain amount of experience with Servlets and JSPs, but for those who do not, the book starts with three introductory chapters on these topics. Although the *SCWCD Exam Study Kit* has one purpose, to help you get certified, you will find yourself returning to it as a reference after passing the exam.

What's Inside

- Expression Language
- JSP Standard Tag Library (JSTL 1.1)
- Custom tags—‘Classic’ and ‘Simple’
- Session management
- Security
- Design patterns
- Filters
- Example code and the Tomcat servlet container
- All exam objectives, carefully explained
- Review questions and quizlets
- Quick Prep section for last-minute cramming



The authors, *Deshmukh*, *Malavia*, and *Scarpino*, are Sun Certified Web Component Developers who have written a focused and practical book thanks to their extensive background in Java/J2EE design and development. They live, respectively, in Iselin, New Jersey, Ardsley, New York, and Austin, Texas.



Ask the Authors



Ebook edition

www.manning.com/deshmukh2



9 781932 394382

ISBN 1-932394-38-9